“This is not the first time such a view has been put forward. If it could be worked out in detail, so rigorously that not the smallest doubt remained, that, it seems to me, would be a result not entirely without importance.”
- Gottlob Frege

University of Alberta

Compositionality and the Metaphysics of Meaning

by

Jeffery Fedorkiw

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of

Master of Arts

Philosophy

Fall 2011

Edmonton, Alberta

Dedicated to Helen Fedorkiw

Abstract

The principle of compositionality states that the meaning of a complex expression is determined by the meanings of its constituent parts and the way those parts are combined. Jerry Fodor has argued that semantic productivity and systematicity requires compositionality and that compositionality requires atomism about semantic values. Atomism is here the thesis that there are simple meanings which are assigned to grammatical terms completely independent of any other (i.e. regardless of anything like context) and that users can grasp any one of these atomic meanings without grasping any other. This thesis argues against Fodor's claim through a defense of Robert Brandom's holistic semantics. I will argue that Brandom's semantics is able to account for linguistic productivity and systematicity without atomism, even though his molecular ontology of semantic values, whereby complex meanings are prior to simple ones, might at first seem intuitively less plausible than atomism.

Acknowledgements

First, all the examining members of my thesis committee deserve thanks for reading this overly long work and providing a challenging oral examination. Many thanks to my external examiner, Christina Gagné, in this regard.

I definitely would not have been able to tackle this subject matter if I was not able to draw on the work of my supervisor, Jeff Pelletier, who also provided insightful and challenging criticism on several drafts. Bernard Linsky sparked my interest in philosophy of language, and most of what I know about it grew from two undergraduate classes I took with him. Ingo Brigandt went far beyond his mere duties in reading several versions of this thesis and sending me many helpful comments. I did not take a course with him, but I sure wish I had.

I believe the idea to write about compositionality was first suggested to my by Nathalie Schapansky, with whom I took two courses on syntax and semantics as an undergrad. The idea of comparing Fodor and Brandom on this point came from a graduate seminar I took with Wes Cooper. I set out there to show that Brandom's semantics were untenable but ended up defending him. This thesis is my second failed attempt to prove Fodor right. Thanks to Andrei Buleandra for encouraging discussions about Brandom and Fodor early in the process.

I have had many inspiring teachers. Though not involved in this specific project, I would like to thank Amy Schmitter, Adam Morton and Randy Wojtowicz, all of whom challenged me to improve.

Table of Contents

Introduction

Ask anything you like about meaning: perhaps, “What is the meaning of the word ‘meaning’?” or “How is meaningful communication possible?” or even “what is the meaning of ‘brick’?” One distinction we can immediately employ in our answer is between simple and complex meanings. In natural languages these two kinds of meaning seem to be clearly encoded in words and sentences (respectively) or morphemes and utterances (respectively). As a first step we could try to clarify the relation between simple meanings and complex ones, and their relations in turn to simple and complex signs. For example, what is the relation between the complex meaning of “the cat is on the mat” and the simple meanings of ‘the cat’, ‘the mat’? Does the meaning of the sentence depend on the meaning of the parts? Or do we only come to know the true meaning of “cat” and “mat” through their use in such sentences? A broad goal of mine is to work out some of these different answers to such questions which result from two opposing views about the relation between simple and complex meanings.

One option is to hold that the simple meanings of words are building blocks from which the complex meanings of sentences are composed according to rules of composition (a grammar). Simple word-meanings are prior to complex sentence-meanings, as bricks are prior to buildings, and the buildings (sentences) are just bricks (words) arranged in a certain way. Positions like this are called semantic atomism. An incompatible option is to hold that

complex-meanings at least partially determine simple-meanings, as buildings illustrate an important usage of bricks. A brick is a brick and not just a rock partially because it is sometimes used to make buildings. Perhaps we are lead to this by the belief that the building is more than the bricks and how they are arranged. Other factors, perhaps function, are also essential to making it a building rather than just a pile of bricks. In semantic theory this is often put in slogans such as “meaning is determined by use”. Positions like this are called semantic holism.

This problematic of semantic atomism versus holism might strike the reader as inappropriate or uninteresting in a number of ways. It might, for example, seem like a chicken-egg sort of pseudo-problem. Maybe both sides are right about different things. But these two approaches provide incompatible answers on a seemingly straightforward question: What kinds of object can serve as the meaning of “brick”? Atomists tend to think that it is pretty clear that some objects are bricks and that any of them might be meant by “brick”. Holists maintain that the meaning of “brick” must include something in addition to the physical objects which helps us identify the bricks among the non-bricks. Some people claim that the atomist position is the only sensible one and see holism as a non-starter. They charge that the holists are sneaking in confused claims about what meaning is based on skeptical arguments about what we can know. But the holists might

counter that the atomist position amounts to nothing more than dogma if we cannot, in principle, evaluate how it could be demonstrated true or false. One might see no hope that one side could convince the other through reasoned argumentation, since both sides presuppose such disparate methods of evaluation.

Yet debates between these two positions are long-ranging both historically and culturally. The disagreement can almost always be found implicit in any discussion of semantics. Sometimes it is explicitly discussed. Indian philosophers and grammarians, for example, had already in ancient times contributed much argumentation to both sides. It would seem that even then the atomist side constituted the common-sense majority.[1] Indeed, semantic holists sometimes delight in the counter-intuitive shock value of rejecting atomism, as in this following bit of infamously strange reasoning from Chinese antiquity:

> "[t]he term 'horse' does not involve any choice of color and therefore either a yellow horse or a black one may answer. But the term 'white horse' does involve a choice of color. Both the yellow horse and the black one are excluded because of their color. Only a white horse may answer. What does not exclude [color] is not the same as what excludes [color]. Therefore we say that a white horse is not a horse."[2]

1 Deshpande, "Language and testimony in classical Indian philosophy".
2 Gongsun Long, quoted in Chan, "A Source Book in Chinese Philosophy", 236.

Nevertheless, a general goal for this thesis is to motivate semantic holism as a plausible approach for the scientific study of natural languages. For this purpose, it is not desirable that white horses turn out not to be horses.

Explicit debate between semantic atomism and holism has again resurfaced in respectable company. In 2001, Jerry Fodor and Ernest Lepore published “Brandom beleaguered”, in which they allege to refute Robert Brandom's holistic theory. The central thesis that Fodor and Lepore employ against Brandom is that holism is incompatible with the principle of semantic compositionality. This principle says that the meaning of a complex expression depends on the meaning of its constituents and the way they are combined. Fodor claims in many recent works that compositionality provides a decisive refutation of any holistic semantics.

The standard argument for compositionality is that it is the best way to explain the fact that indefinitely many meaningful complex expressions can be produced from a finite set of constituents using a finite set of grammatical rules. This is called linguistic productivity. However, Fodor does not just think that compositionality is simply the best explanation of productivity. Rather, compositionality appears to be the only available explanation and compositionality demands atomism.[3] Compositionality and atomism are, in turn, central tenets of what Fodor says is “still

3 Fodor, LOT 2, 20.

the only game in town", referring to the computational theory of mind.[4] Brandom's holism is not compositional, his notion of concepts is not atomistic, and so his theory is a non-starter according to Fodor.

The specific goal of my thesis is to argue that Fodor is wrong about this. I will argue that Brandom's holism is not refuted by its incompatibility with compositionality because he is able, in principle, to account for the productivity of language in a non-compositional way. If that is successful, the pertinent question then becomes how the two theories compare in light of the fact that both aim at a semantic theory which, among other things, is coherent with the aims of cognitive science. On this point I hope to provide strong considerations for both sides, although I am not completely convinced of either. Fodor's atomism, I contend, is the more intuitively plausible theory. Brandom's holism is the less-intuitive, perhaps less plausible, but nevertheless possible theory which one might accept because of intractable skeptical problems with atomism.

Unsurprisingly, it turns out that Fodor and Brandom are both employing different notions of meaning. If I am right that Brandom's notion is not ruled out a priori by its non-compositionality, then a deciding factor could be whether one notion of meaning is better than the other as an explanation of linguistic

4 Fodor, "Reply to commentators", 112.

productivity. Explaining linguistic productivity has been central to cognitive science ever since, as legend has it, the discipline was birthed by Noam Chomsky's devastating review of Skinner's Verbal Behaviour. Chomsky's focus at that point was on how mastery of syntax could not be explained through behaviourism. However, the productivity of semantics seems it should also be a viable target of scientific explanation. Admittedly, the aims of cognitive science are of far more importance to Fodor than Brandom. However, since any semantic theory must be at least compatible with productivity, so should Brandom's. Fodor and Brandom agree on a number of other points for comparison as well, at least at a highly abstract level.

First of all, both Fodor and Brandom find meanings out in "the world". "The world" is here dressed in quotes because it includes not only physical objects like bricks, buildings, protons, marks on paper and soundwaves, but also fictional objects like Felix the cat and proton torpedos, as well as mental objects such as thoughts and concepts which relate to everything else and themselves. Thus the meaning of symbols can sometimes be other symbols (or the symbols themselves) as is the case when we learn prescriptive grammar, learn to speak a second language, or theorize about language.

Secondly, both Fodor and Brandom take is that mental objects stand between the symbols of natural language and "the

world". For both, "concepts" are the mental objects that roughly correspond to words and "thoughts" are the mental objects that roughly correspond to sentences. Concepts and thoughts can, in turn, refer to "the world" which allows them to be evaluated as true or false, appropriate or inappropriate. So for both theorists, the meaning relation stands between what might be seen as two other relations. On one side is a grammatical relation which takes atomic symbols and combines them into syntactically well-formed expressions. On the other side is subject/object relation of intentionality by which minds are directed toward "the world". The meaning relation performs two functions, which may or may not be separable: it assigns mental objects to the syntactically well-formed formula produced by a grammar; and it assigns references in "the world" to those mental objects. At least, it does so when there is a reference. Even expressions that do not refer seem to point our minds toward outside objects —— Alexius Meinong famously theorized about such objects and created a taxonomy of them based on whether they exist, could exist (like the golden mountain), or even were impossible (like the round square).[5]

Taking semantics to be ultimately concerned with the relation between minds and the outside world is what might be called philosophical semantics or truth-theoretic semantics. I believe both Fodor and Brandom can be broadly construed as taking this

5 Meinong, "The theory of objects".

approach (although the term “truth-theoretic semantics” could describe Brandom's theory only after a long pragmatic revisionism about “truth”). Both theorists treat intentionality as intricately entwined with semantics. For philosophical semantics, a theory of meaning is not a theory of meaning without also being a theory of intentionality. But the notion of intentionality has itself been a philosophical problem since it was re-introduced into contemporary discussion by Franz Brentano.[6] Brentano claimed that this “directedness” of thought towards objects outside of itself was the mark of the mental — all mental activity is intentionally directed to some object. Many people might not agree that all thought is directed outward but perhaps all those thoughts which can be communicated are.

Brentano predicted that intentionality could never be fully explained in purely naturalist terms. Some influential theorists agree, and presumably hold that intentionality cannot, in turn, be an object of study for the natural sciences. Noam Chomsky, for instance, has said that “[i]f ‘cognitive science’ is taken to be concerned with intentional attribution, it may turn out to be an interesting pursuit (as literature is), but it is not likely to provide an explanatory theory or to be integrated into the natural science.”[7] Thus, according to this relatively recent development in Chomsky's thought, the scientific study of language

6 Jacob, “Intentionality”.
7 Chomsky, “Explaining language use”, 209.

should not concern itself with language as an external artifact by which sounds and marks on paper express abstract, external propositions which are in turn made true or false by the way the world is. Rather, the scientific approach should view language as intensional and internalized. It is internalized in being a “system represented in the mind/brain, ultimately in physical mechanisms”.[8] It is intensional in that it is a “specific characterization of a function which assigns a status to a vast range of physical events” which include grammatical utterances, semi-grammatical utterances, utterances of foreign languages and even noises.[9] The goal, then, would seem to be a characterization of a function which maps perceived stimuli onto brain events, even though at this stage we can only describe those brain events in a metaphorical way. Any intentional objects of a vaguely Meinongian nature, like thoughts or propositions, are eventually to be explained away and any concern with the relation between the mental and the physical disappears.

Both Fodor and Brandom are opposed to this reductionism — or at least accept that it could not be done without a lot of philosophical problem solving. Fodor agrees with Chomsky on many points. His language of thought is internal, for instance, but he is still deeply concerned with the problem of naturalizing intentionality. Although he takes it that only matter can think,

8 Chomsky, “Language and problems of knowledge”, 10.
9 Ibid.

thinking involves, at least sometimes, an ability to refer to other objects, some of which are propositions.[10] Brandom, as we will see, clearly takes language to be an external artifact, specifically a social artifact, to be studied and theorized about.

Chomsky seems to want a semantic theory which side-steps the hard problem of intentionality which both Fodor and Brandom confront. It is not clear to me that he can. He himself says that semantics should be the study of the relation between language and the world. He continues, with some amount of disapproval, to say that "much of what is called "semantics" is really the study of the syntax of mental representations".[11] Should we consider the brain events which are mapped to stimuli to be "things in the world" and call the theory semantics, or are these brain events "mental representations" and the resulting theory a syntax? The latter amounts to giving up on a scientific approach to linguistic meaning. But the former might just raise the problem of intentionality is a different guise, if the causal relations between perception and brain events are not just one-way.

Ray Jackendoff draws a distinction between philosophical and psychological semantics which we might interpret as a distinction between literary and scientific semantics respectively. His discussion of this distinction reveals its limited use, however. He characterizes philosophical semantics as epistemolo-

10 Fodor, LOT 2, 196-7.
11 Ibid., 23.

gical and concerned with our relation to the world, and psychological semantics as following from issues in perception and concerned with how the brain functions to enable our successful interactions in the world.[12] He discusses Fodor as someone who is firmly planted in both sides of this divide. Qua cognitive scientist, Fodor is concerned with the scientific issue of how our mental representations allow us to successfully exist in the world. Qua philosopher, Fodor is concerned with things like truth and reference which are ultimately only explained in a literary way as "modes of dasein".[13] For one thing, Jackendoff might be taking Fodor's tongue-in-cheek invocation of "dasein" a little too seriously. For another, I do not think Fodor would see himself as having two feet planted in a scientific/non-scientific divide. Rightly so, because the usefulness of the distinction between philosophical and psychological semantics seems to disappear fairly quickly.

Jackendoff's own position, for example, is to favour the psychological side. But this is done only with the justification that investigation into the psychology of internal representation can reveal the constraints on how we perceive reality, whatever it is.[14] This quasi-Kantian justification betrays a concern with our mental relation to the world. Mental representations need to

12 Jackendoff, Languages of the Mind, 158.
13 Jackendoff, 159.
14 Ibid.

have some connection with the non-mental world even if the only thing we are after is an explanation of evolutionary success. So if there is a divide, Jackendoff also has both feet planted across it, even if he is able to carry on his daily work without using philosophically loaded words like “truth” and “reference”.

I see Fodor as attempting to explain our meaningful connection to “the world” with the vocabulary and methods of natural science, but not in the indirect Kantian manner that Jackendoff suggests. Although Fodor is sometimes forced to make concessions to vaguely Kantian notions of concepts shaping perceptual experience,[15] his semantics aims at a much more direct connection to “the world”. Although concepts can shape preconceptual representation, they could only do so if we they were first causally formed. If I show you the cat, and the mat, then perhaps you can know what I mean by “the cat” and “the mat”. The cat and the mat, as it were, cause representations of themselves in your brain. If you have a further representation of the relation “x is on y”, and the means of expressing that in English, then you can go on to say “the cat is on the mat”.

Although Fodor postulates concepts and thoughts as standing between words and the world, the connections betwixt are straight lines. The meanings of “the cat” and “the mat” are concepts, but the meanings of those just are the physical objects. There is a

15 Fodor, LOT 2, 181.

concept related to the English phrase “x is on y”, but the meaning of that concept is a property of the physical world. With all the apparent simplicity of pointing and naming, our thoughts are directed to the world through finite concepts which grasp finite objects. With these in hand, rules of grammar allow us to compose complex thoughts we may not have entertained before. Those thoughts are true or false depending on whether things in the world are composed in the way we say they are. Given that we are after a “truth-theoretic” semantics which connects us to the world, the advantages of this theory are immediate.

Indeed, I think that Fodor's semantics is a detailed, philosophical defense of an assumption implicit in many approaches to language. For instance, in a critical review of an argument against compositionality, Barbara Partee defends that principle in part by defending the truth-theoretic approach to semantics.[16] The natural assumption she makes, along with almost everyone else, is that any theory which is concerned with our mind's connection to the world, that is, any truly semantic theory, must adhere to compositionality. Fodor works out in detail the further natural assumption that compositionality only works with semantic atomism. It seems like common sense, yet Fodor often portrays himself as the perpetual underdog in philosophy.[17]

16 Partee, “Semantic Facts and Psychological Facts”.
17 See, for example, his “Having concepts: A brief refutation of the 21st century.”

Brandom is similar to Jackendoff in taking an indirect approach inspired by Kant. He argues that the meaning of linguistic expressions cannot come from mental representations which are directly caused by the world. The concepts which natural language relates to us shape our very experience of the world. However, unlike Jackendoff, he does not think a psychology of mental representations is sufficient to characterize these conceptual constraints on perception. The story connecting words with the world that Brandom tells takes an even further detour through the social practices of giving and asking for reasons. Words have their meaning because of the contribution they make to the inferential value in sentences. Sentences have their meaning because of the way they make explicit our commitments to the way the world is. And we discover this meaning by asking and providing each other for the reasons behind our beliefs and implicit in our actions. The difficulties for a holism of this sort begin early in the explanation. Complex meanings are both ontologically and psychologically prior. The mind is intentionally directed to “the world” first by complex thoughts expressed by sentences because it is these through these complex meanings by which we act and explain our actions in the world. If complex meanings are ontologically prior to simple ones, it is natural to conclude that it is the set of all the sentences which is “truly meaningful”. But that is absurd as a psychological explanation of

actual language users. It is surely false that any mortal ever knows the meaning of all the sentences of his or her language. There are indefinitely many such sentences, and many of the complex thoughts they express have not yet occurred to anyone.

Brandom sometimes seems to support holism in this extreme sense. In replying to Fodor's challenge to identify which inferences are the ones which confer semantic content on sentence, he says that it is a viable option to say that it is all the inferences. This approach is inspired by Quine's arguments that meaning is only attributable to an entire theory. He even sees his project in Making it Explicit as taking up this strategy.[18] I do not think it is a viable option. On the one hand, if we interpret this extreme holism as meaning that all the inferences a given speaker is inclined to make as what defines his or her meanings, we will fall into an absurd relativism. We simply could not talk about the same things because by matter of fact we do not make the same inferences. On the other hand, if we interpret this extreme holism as meaning all the inferences everyone who shares a language make, we are committed to an absurd fatalism about meaning. The meaning of a word is just what society tells us it is and it seems we could not disagree about it.

Fortunately, in my opinion, Brandom offers another option. Inspired by Wilfrid Sellars, we can say that the complex-meanings

18 Brandom, "Inferentialism and some of its challenges", 662.

of sentences are conferred by the counter-factually robust inferences which hold between them.[19] The inference from “x is a cat” to “x is mammal” is part of what constitutes the meaning of cat. If Hegel were a cat, he would be a mammal. But the inference from “x is on my mat” to “x is a cat”, which is contingently true whenever Petey the cat sleeping on my mat, does not constitute the meaning of cat whatsoever. If Hegel were on my mat, it would not follow that Hegel is a cat. This approach is brought to the fore in Brandom's Between Saying and Doing, but it is compatible with the specific semantic theory of Making it Explicit. I think that it is this strategy which explains how Brandom's holistic semantics can be productive and it is this strategy I will defend against Fodor.

I think that this strategy allows for a recursive explanation of language productivity because meaning can be conferred by a finite set of inferences. The language user grasps a partial set of complex sentences and corresponding complex meanings. The fact that they confer meaning can be tested against the fact they are counter-factually robust. A grammar then allows him or her to analyze the sentences into component parts, and that in turn allows her to evaluate the meanings of those component parts based on their inferential role in the set. These component meanings might be considered the language user's “theory”. However, even

19 Ibid., 661.

if these component meanings must be reevaluated in light of further evidence, they can still be combined into novel thoughts and be expressed by novel sentences, using the same grammar which allowed us to analyze the initial set. The language user is then able to increase his or her set of complex-meanings, which is to say she is able to produce and understand sentences which were not part of the initial set of grasped meanings.

My thesis is that Brandom's holistic semantics is able to explain language productivity along these lines. It involves the analysis of complex expressions into component meanings but it is not a reduction of the complex meaning to component meanings. It allows for the projection of new expressions, the meanings of which can be shared and communicated amongst speakers. Fodor's alleged refutation is false. However, Fodor also employs whole range of related considerations against what he calls pragmatic theories which do support the initial intuitive appeal of his own atomism. Against this, I hope to motivate holism as a plausible candidate for semantic theory. Here is the plan:

Chapter 1 draws on contemporary work to analyze the concepts of atomism, holism, compositionality and language productivity. I will purposely choose interpretations of these notions which emphasize the incompatibilities between Fodor's atomism and Brandom's holism. Chapter 1 also refers to some recent formal work regarding the alleged incompatibility between

compositionality and holism. I will use this material to argue that Brandom's holism is not formally absurd. That is to say that some kinds of holism, specifically my reading of Brandom's holism sketched above, can, in principle, explain language productivity.

Chapter 2 is a survey of how atomism and holism have been treated in early analytic philosophy of language. The works discussed here were all written before explicit attempts to define and defend the principle of compositionality, but one can find many aspects of the principle at work. I argue that Russell's logical atomism is not refuted by skeptical arguments against it drawn from Plato and Wittgenstein. Finally, I discuss how both the apparently atomistic principle of compositionality and the apparently holistic principle of contextuality have often been found in the work of Gottlob Frege. In that section I draw on some recent work suggesting that the compositionality was not significant at any point in Frege's work. I also put forward Brandom's preferred view that Frege is a sentential holist, but I do not wish to defend it in detail.

Chapter 3 discusses Fodor's semantic atomism. The main contention is that Fodor sees compositionality as requiring an atomistic metaphysics of meaning, and he sees compositionality as required in order to explain productivity. He presents these considerations in two ways. In Concepts: Where Cognitive Science Went Wrong, compositionality is presented as a natural explanation

of mental phenomena when the computational and representational theories of mind (CTM and RTM) are adopted. The argumentative support comes indirectly from arguments against “pragmatist” approaches — what Fodor terms any non-computational, non-representational approach. In LOT 2: The Language of Thought Revisited, CTM and RTM are required because of compositionality, and compositionality is required in order to explain productivity. Here, Fodor sees compositionality as what is needed to explain productivity. That seems intuitive at first. The problem is that productivity does not require his particular, atomistic semantics.

Chapter 4 will present my reading of Brandom, which emphasizes sets of counter-factually robust inferences as the primary components of meaning. I explicate arguments from Between Saying and Doing that this sentential holism (also called molecularism) is recursive and projectable. As such, it can explain productivity. I also attempt an explication of the semantics presented in Making it Explicit, which founded on the notions of inference, substitution and anaphora. I believe that this is an instantiation of the logical framework from Between Saying and Doing, and it provides a way of filling in the story of how speakers could employ meanings as Brandom conceives them.

The end result is somewhat of a disappointed for the metaphysically ambitious. Unfortunately, I do not think that lin-

guistic productivity or the principle of compositionality can prove either side of the abstract debate between atomism and holism. At best it rules out only some kinds of holism. Working out the details of this failure, however, is illuminating. I think Fodor's theory takes many people's initial intuitions about linguistic productivity down a logical path, but the end result can be peculiar in detail. The argument left is still that atomism is the only choice because of the failure of alternatives. But that cannot be simply because productivity is true and compositionality is the only choice. We must also attend to Fodor's arguments that an ontological notion of meaning is prior to anything epistemological or pragmatic in nature. Brandom's holism is also presented as the best choice because of the failure of atomism — it fails because behind atomism's seemingly straightforward explanation of our intentional relationship between symbols and the world lie intractable conceptual difficulties.

1. Does Compositionality Demand Atomism?

My concern is the validity of the following claim: the principle of semantic compositionality demands semantic atomism and is incompatible with semantic holism. The preceding introduction presented the contents of these concepts without analysis. This first chapter will provide that analysis. The relevant context of this claim is Fodor's use of it against Brandom. But that, in turn, must be seen in light of the broader use Fodor makes of this claim in support of his representational theory of mind and against all those non-representational theories he calls "pragmatism". Fodor probably offers more arguments-per-word count than any other philosopher, and recently almost all of them have involved some part of the above claim as an assumption, premise or conclusion. I cannot survey them all. Instead, let us start with this:

> (1) the semantics of language (and/or thought) is productive
> (2) productivity requires that language (and/or thought) be compositional
> (3) compositionality entails atomism
> (4) inferential role semantics is not atomistic
> therefore (5) inferential role semantics is false

I claim that this argument is not sound because it fails at either (2) or (3), depending on how we define "compositionality". In this chapter I will show that this argument is unsound because

atomism does not follow from either compositionality or productivity in the “necessary” or “a priori” sense that Fodor often presumes. As mentioned, however, he does not just presume that productivity demands an atomistic principle of compositionality and that it refutes holism, he provides many other arguments against holism which appeal to compositionality and productivity in some way. The goal of this chapter is simply to pave the way for a defense of Brandom by showing that holism is not ruled out in some logical sense.

First, note that we can take (3) as an assumption and conceive of compositionality accordingly, such that it does entail atomism. I think Fodor often does take it this way. However, with this restricted definition of compositionality the argument fails at premise (2) if a non-atomistic semantics can be productive. The standard argument for (2) is that compositionality is the only, or, failing that, the best, explanation of the productivity, novelty, and learnability of language. I will, by and large, take the common denominator of all these to be the fact that speakers and listeners can evaluate the meanings of complex expressions they never before produced or encountered and will mostly refer to this as productivity — although novelty and learnability are somewhat separable from productivity.

It is commonly agreed that the “standard argument” to the best explanation offered for compositionality insufficient when

the principle is meant to carry the philosophical burden of entailing semantic atomism.[1] Which is to say that Fodor should know better. But it just as common to run compositionality and atomism together as a common sense view of language. Consider these lines by John Perry:

> Words stand for things of various kinds and for various kinds of things. Because words do this, the sentences made up of words mean what they do, and are capable of expressing our thoughts, our beliefs and conjectures, desires and wishes. This simple idea seems right to me, but it flies in the face of formidable authority.[2]

It is not only Perry who thinks this idea just seems right. Put that way, it seems right to me too. It relies on a plausible view of the relation between intentionality and language: we are able to employ language because we are first able to direct our thoughts towards the world. I believe that a detailed look at how compositionality does or does not support semantic atomism provides an explication of this intuitive appeal that Perry refers to. Now the formidable authority Perry argues against is Davidson, whose position I do not wish to defend. Interestingly, he employs arguments from Wittgenstein's Philosophical Investigations against Davidson's holism, a move which seems particularly

1 See, for example, Zoltan Szabó's unpublished Ph.D dissertation for a discussion of how the "standard argument" is insufficient and what might be done about it.
Szabó, Problems of Compositionality.

2 Perry, "Davidson's sentences and Wittgenstein's builders", 23.

daring since that work of Wittgenstein's has arguably done more than any other to inspire countless philosophers to turn away from the atomistic, referential view that Perry think "just seems right". Surely it is Wittgenstein's considerations against the "Augustinian picture" of language found at the beginning of that work which are the real formidable authority against atomism.

In fact, I think that the skeptical arguments against atomism which one might find in the Philosophical Investigations provide the best chance at lending holism any kind of intuitive support. Without the tools given us by Wittgenstein, holists could very well be out trying to convince people that white horses are not horses. Brandom frames his own project in Between Saying and Doing as providing a non-quietist interpretation of the Investigations which allows for the possibility of semantic analysis. Sentences can be used to make various claims about the world, and function as premises and conclusions in inferences. As such they can be offered as reasons justifying our beliefs about the world and our actions within it. It is this use of sentences which makes our relation to the world conceptual and it is what allows us to express what non-linguistic animals just do. This places linguistic meaning in the social domain, in this case the social domain of reasons. Many inspired by the Investigations have held that meaning can only be understood through its employment amongst discursive creatures, and Brandom's theory is one

way of doing so. This reverses the order between intentionality and language: we are able to direct our thoughts toward the world because we are able to employ language. Now, I do not think anyone could say all that “just seems right”, not unless they are first convinced that there are deep and inextricable problems with the atomistic view.

1.1 What is Semantic Compositionality?

Put informally, the principle of semantic compositionality tells us that the meaning of a complex expression (typically a sentence) is determined by the meaning of its parts and the way they are combined. The “standard argument” is that it provides an inference to the best explanation of the productivity and systematicity of language. This point is aptly illustrated by supplying a sentence which the reader has likely never encountered before, yet still understands. For example: “Fred Rogers was killed during a police standoff today.” It is unlikely that anyone else has ever claimed that the gentle star of Mr. Roger's neighborhood was involved in that kind of event. Yet I can use it to convey a meaning which you can grasp. Many, perhaps a very large percentage, of the sentences we produce and encounter are novel in this way. Not only can we give them a semantic evaluation, apparent instances of successful communication suggest that we often give them the correct evaluation, all with little to no effort.

It is a natural assumption that we can do this precisely because we first grasp the meanings of each of the words that they are composed of, and the rules of their composition. From there, it is also natural to assume that the meanings of words must exist prior to the meanings of sentences since we can use familiar words to express novel thoughts. The principle of compositionality captures the former assumption, semantic atomism the latter. The individual words of language must already have a meaning in order to be used, but sentences do not have a meaning until they are actually produced by someone, and maybe also understood by someone else. The rules of composition allow us to combine the simple meanings into complex ones, and even to identify meaningful wholes from meaningless aggregates of words.

I take the “standard argument” for compositionality to be centrally concerned with explaining productivity, although a number of related features of natural language are also relevant. Language is productive because an arbitrarily large number of expressions can be produced, as well as understood, from a finite, and much smaller set, of symbols and rules for their composition. Some proponents of compositionality stress this point by saying that the number of novel expressions we can produce and understand is infinite. However, as Richard Grandy points out, it is not necessary for proponents of compositionality to commit themselves to this perhaps hyperbolic claim. It is enough that novel

sentences are understood and that there is an exponential growth between words known and sentences understood.[3]

Productivity is naturally run together with arguments about language learnability. How could finite human beings learn languages if they were not compositional? Formal properties of productivity might be abstracted from psychological concerns of learnability, but I will not be too concerned with doing so here. I take a central concern of learnability arguments to be an explanation of productivity: compositionality is part of what allows someone to learn how to evaluate the meanings of complex, novel expressions of natural language.

Another aspect of the standard argument for compositionality is the systematicity of language. Language is systematic because if we understand the form of an expression, we can understand and produce another meaningful expression of the same form. Explanations of this feature often rely on examples along these lines: if you understand “John loves Billy”, you can easily produce “Billy loves John”. It is difficult to identify precisely just how consideration of systematicity is meant to support compositionality in a way <u>different</u> from productivity. In §3.2 I employ Kent Johnson’s arguments that language is not systematic against Fodor’s appeal to systematicity. Nevertheless, it might be held either that the standard argument for compositionality

3 Grandy, “Understanding and compositionality”, 558.

can stand on considerations of productivity alone, or that the important aspects of systematicity for the standard argument are its connections with, and similarities to, a more general argument from productivity.

Systematicity might be seen as a special case of productivity — one of the ways in which language is productive and novel. Systematicity might also be considered an important feature of language because it draws a link between syntax and semantics. Understanding the form might presuppose a certain understanding of the meanings involved — what meanings can be the object of X's love, for instance. Pagin and Westerståhl suggest that one way systematicity offers new considerations for compositionality is by offering predictions as to what novel utterances will be meaningful. However, they also point out that this can be problematic for proponents of compositionality when the predictions fail.[4]

The "standard argument" might refer to compositionality as the best explanation of language productivity, systematicity or novelty. Pagin and Westerståhl draw conceptual distinctions between these variations. They also add "inductions on synonymy" as well as "intersubjectivity and communication" to the list of things compositionality is meant to explain in "the standard argument".[5] My concern is more with what ontology of meanings follows from compositionality, rather than what compositionality

4 Pagin and Westerståhl, "Compositionality II: problems and arguments".
5 Pagin and Westerståhl, "Compositionality II".

follows from, so I will not draw the distinctions between these variations of the standard argument too sharply. The common denominator to all, I take it, is explaining how the meanings of complex, novel expressions can be evaluated. This explanation should be logically possible, given the indefinitely large set of complex expressions it is meant to explain. It should also be, at the very least, psychologically plausible.

This draws out two important components to the intuitions behind compositionality: psychological and metaphysical. The psychological component arises from the fact that compositionality is meant to provide an explanation of how creatures like us are able to understand and learn language. Compositionality is “in the head”. It is a basic principle that is presupposed in our daily use of language. The metaphysical component arises from the expectation that if compositionality is philosophically interesting, it will tell us something about the ontology of meaning. Simple and complex meanings, whatever they are, must reflect this compositionality or we could not communicate successfully.

There are plenty of reasons to be dissatisfied with the “standard argument”. One reason is precisely that these metaphysical and psychological intuitions tend to be blurred. Another is that inferences to the best explanation do not seem decisive. Yet many have been apt to take these suggestions about language at face value. As Pelletier says, there is a “warm fuzzy

feeling", many get when a theory claims to be compositional.[6] Furthermore, atomistic views of meaning tend to "just seem right". But it is not easy to see how the argumentative support for compositionality can be bolstered beyond inference to the best explanation. I think Fodor's arguments add a lot to understanding compositionality, and support for it. But I still do not think he can draw the metaphysical conclusions about meaning he wants from the fact of linguistic productivity.

Alongside much recent work at clarifying arguments for compositionality are attempts to formulate more precise definitions of the principle. This requires resolving some ambiguities in informal variations of the principle. A problem that arises here is that the principle is used in different disciplines, especially psychology, linguistics, and philosophy, with different purposes in mind. Zoltan Szabó, for example, has given three variations of the principle for each of these three disciplines.[7] As the main concern of this thesis is Fodor's use of the principle to support semantic atomism, I am interested in which do entail such a commitment. Szabó first analyzes this formulation:

> [Φ] The meaning of a complex expression is a function of the meanings of its constituents and the way they are combined.

6 Pelletier, "The principle of semantic compositionality. 11.
7 Szabó, "The case for compositionality".

Fodor is committed to such a principle: a thought has the meaning it does because it is composed of concepts which have the meanings they do. It seems that a theory which defines meanings recursively (as Brandom does) might also adhere to [Φ]. However, according to Szabó, this formulation of the principle entails the stronger, problematic, principle of compositionality [C]:

> [C] The meaning of a complex expression is determined by the meanings its constituents have individually and the way those constituents are combined.[8]

Analysis of two other possible definitions of compositionality will be given in §1.3, but [C] suits the purposes of both explaining productivity and requiring the ontological dependence between simple and complex meanings which entails semantic atomism. [C] is the result of resolving 3 ambiguities in [Φ].

The first ambiguity is between “function of” as a “merely mathematical” function as opposed to a dependence relation. [C] resolves this ambiguity in favour of a dependence relation. Philosophers are generally not interested in what might be called “trivial” accounts of compositionality which rely on a “merely mathematical” functions between words and sentences. A function might be “merely mathematical” if it does not show a dependency relation between the semantic parts and wholes. This requirement of a dependency relation bars functions along the lines of Zad-

8 Ibid.

rozny's demonstration that a compositional function could be trivially given for any semantics.[9] A function might also be considered "merely mathematical" if it is so complex that it is not a plausible principle for language users to employ. This disambiguation, then, gives the ontological dependence relation which Fodor suggests, but it also opens the principle up to supposed counter-examples from natural language which will be seen below.

The second ambiguity is between an interpretation of "the meanings of its constituents" as being read individually or collectively. [C] resolves this in favour of an individual reading. This requires that every part which is a component is assigned a meaning. Not all parts are components. For example, "the cat ate" is part of "the cat ate the rat", but it is not generally considered a component according to the rules of composition. If we suppose a final analysis which will identify a finite number of components, there will be some finite, simple components which are assigned meanings. If it is this aspect of compositionality which productivity requires, it can be met in a non-atomistic way. It is just that the meanings which analysis arrives at will be seen as defeasible and theoretic, rather than reductions of the complex meaning.

The third ambiguity is between "they" in "the way they are combined" as referring to the constituents or the meanings of the

9 Zadrozny, "From a compositional to a systematic semantics"

constituents. [C] resolves this in favour of “they” referring to the constituents. The result is that the rules of combination operate on the grammatical constituents rather than their meanings. This suggests an identification between the rules of semantic composition and the rules of syntactic composition. I think that here we see how the principle might be supported by arguments similar to the “poverty of the stimulus” which many, following Chomsky, have seen as establishing the existence of an innate grammar. It is fairly uncontroversial to claim that syntax must be compositional. Whether a complex expression is well-formed surely depends only on the syntactic properties of its components and the way they are combined. If semantic compositionality could piggy-back on this claim, it would be a boon to atomism.

All three of these resolutions work to ensure that [C] is the strongest possible interpretation of the principle — one that draws a clear dividing line between theories which accept it and theories which do not. Fodor does, Brandom does not.

There have been plenty of alleged counterexamples to the principle of compositionality, especially formulations like [C]. Szabó provides three examples which I think correspond nicely to three classes of counterexample to compositionality:

> [CE-AMB] Ambiguity: [C] does not allow for context to determine the correct interpretation of semantically ambiguous constructions. Szabó mentions examples where the scope of a quanti-

fier is undetermined. For example, “every student drank a beer” could mean that each student drank his or her own beer, or that all of them drank one beer. Other examples include sentences with only one syntactic parsing but multiple semantic interpretations. For example “the philosophers lifted the piano”.[10]

[CE-MOD] Modifier-head constructions: The same modifier can seem to apply to different heads in different ways. AN-phrases provide clear illustrations. For example, “Red” modifies the following nouns differently: red apple (is applied correctly if the apple is red outside), “red grapefruit” (is correctly applied if the grapefruit is red inside), “red wine” (is correctly applied if the wine is red throughout but can also be more of a purple, really) and “red ideology” (is correctly applied if the ideology is communist).[11] Since in every case it is the same word applied in the same way, it seems it ought to contribute the same meaning in each case.[12]

[CE-SYN] synonymy/Frege Cases: accepting [C] leads to direct reference theories, where

10 “The philosophers lifted a piano” is an example due to Pelletier. It is different from “every student drank a beer”, because in the latter, the semantic ambiguity can be resolved by choosing a particular syntactic parsing. The former has only one syntactic parsing, it might mean that the philosophers collectively lifted a single piano, or each individually lifted a piano.
Pelletier, “The principle of semantic compositionality”, 20.

11 These different variations are discussed at length in:
Lahav, “Against Compositionality”.

12 Szabó, “The case for compositionality”.

> phrases like “Cicero is Cicero” and “Cicero is Tully” must be synonymous. This seems counter-intuitive since I can trivially know that Cicero is Cicero without knowing, or even believing, that Cicero is Tully.

In chapter three I will construct responses to these counterexamples on Fodor's behalf. [CE-AMB] and [CE-MOD] might be considered pragmatic rather than semantic problems. As such, they might be simply pushed aside by atomistic theories as problems to be addressed by psychologists or linguists on the basis of semantics. In fact, Fodor often uses examples of the [CE-MOD] type to argue against holistic theories. Fodor takes counterexamples of the [CE-SYN] type seriously, and offers a direct, albeit partial response which will be discussed in chapter three. These counterexamples are centrally difficult, because they suggest that information which identifies the reference be included along with the reference as the meaning of a name. But allowing that introduces pragmatic aspects which Fodor wants to bar from a notion of meaning.

Another much discussed issue in recent literature is whether the principle is in fact true. Aside from counterexamples of the 3 sorts listed above, there are <u>prima facie</u> reasons to think that the principle is not <u>universally</u> true of every complex expression, at least. Any theory which thinks productivity requires compositionality in the strong sense of [C] will need to account

for this sort of exception. However, as long as the set of idiomatic expressions are finite, they do not pose a challenge to the view that an indefinite (maybe infinite) set of linguistic utterances are compositional. Second, even if compositionality is true, it seems we can sometimes get along fine without it. One strategy of training young people to be good readers, for example, is to teach them to infer the meaning of an unknown word from its context. It seems we can grasp a lot of the meaning of a complex sentence even when we do not grasp all the component meanings. We can know a lot about what "the gavagai is hungry" means, for example, without knowing exactly what a gavagai is. Fodor's own position is an acceptance that natural language is not strictly compositional, but that thought must be. He even employs this in an argument that intentionality must inhere in thought prior to language.[13] Thus, non-compositional ambiguities in natural language are resolved in the strictly compositional and atomistic language of thought.

No one denies that language is productive — if not infinitely productive then at least productive in Grandy's sense of exponential growth. But productivity and compositionality are often equivocated. Even if productivity is a readily apparent fact of language, compositionality is not. The word "compositionality" has been sometimes used to mean whatever it is that explains

13 Fodor, "Language, thought and compositionality".

productivity. Fodor often talks this way. But it can also be used to mean a commitment to semantic atomism as an explanation of productivity. Fodor often talks this way too. In the former sense, a philosopher could use “compositionality” to mean a recursive theory which assigns meanings to component expressions on the basis of meanings first assigned to the complex expressions which use them. In that case, the meaning of constituent expressions depends on the meaning of the complex expressions in which they appear — a reversal of [C]. Conversely, “compositionality” might refer to a theory which explains language productivity specifically by appeal to semantic atomism. This is the case for anyone who accepts something like [C]. But [C] goes beyond how we can produce and understand novel complex meanings, and entails a certain ontology of those meanings.

Brandom has wavered in whether his theory ought to be called “compositional”. In a response to Fodor, he agrees that “productivity demands compositionality”, but denies that “compositionality implies the priority of sub-sentential semantics to sentential semantics.”[14] In that case, it seems that he takes compositionality to be whatever explains productivity. However, in Between Saying and Doing, he describes his position as a rejection of compositionality in favour of a holistic explanation of productivity. In that case, he is taking compositionality to

14 Brandom, “Inferentialism and some of its challenges”, 677.

refer to a metaphysical thesis about the priority of sub-sentential semantics to sentential semantics.[15] In this thesis I will take the latter nomenclature as the one he settles on, and will therefore refer to his position as “non-compositional”.

It can seem like everyone since Frege believes in the principle of compositionality. It can also seem like Frege talked about compositionality, but that is mistaken. In fact, Frege himself probably did not adhere to the principle, at least not in the sense of [C].[16] In fact, direct discussion of the principle was extremely rare until a recent explosion of debate across the disciplines of philosophy, psychology, and linguistics. Only a tiny fraction of the philosophical literature has been surveyed here. I suggest the results of much of this recent work on compositionality show that Fodor is making a mistake by drawing an ontological conclusion about meaning from facts about linguistic productivity. But if he is making that mistake, he is not alone. Chapter 2 will show how analytic philosophers have been making similar claims about meaning and productivity ever since Frege.

15 Brandom, Between Saying and Doing, 133-6.
16 Janssen, “Frege, contextuality and compositionality”.
Pelletier, “Did Frege believe Frege’s principle?”.

1.2. What is Semantic Atomism?

Semantic atomism is the thesis that the meanings of constituent parts are ontologically prior to the meanings of the sentences of which they are a part. Atomism is entailed by the definition of compositionality [C], introduced above. This is because [C] holds that every constituent part of an expression is given a meaning, and that the meanings of these constituent parts determine the meaning of the whole. Read in a way which only allows the determination to happen one way (it is never the meaning of the whole which determines the meaning of the parts), this entails that the meanings of the constituent parts be ontologically prior to the meanings of sentences.

All atomistic theories share central theoretical commitments. For one thing, atomism requires that there be meanings which correspond to the simple parts. The meanings of any given simple part, furthermore, must be independent of any other meaning. For example, whatever the meaning of “horse” is, it does not depend on the meaning of “animal” or any other word. So you could grasp the meaning of “horse” even if you have not grasped any other meanings. Fodor often argues that if the meanings of the component parts did not have this ontologically independent meaning, then we could not say how having the meaning of “white” and the meaning of horse “horse” is sufficient to grasp the meaning

of “white horse”.[17] You have the meaning of a word or not, and any descriptive explanation of that meaning may well be insufficient to capture it. In atomism, simple meanings are not definitions. When they are grasped, it is perhaps similar to Russell's direct acquaintance, which will be discuss in §2.3.

Holism, on the other hand, is united more by a denial of atomism than adherence to any specific, common, positive ontological commitments. Holists deny that the meanings of simple parts can be independent in the way described above. Beyond that, some holists, like Brandom, are committed to the existence of complex meanings like propositions. Others, like Quine, hold that meaning only exists at the level of a speaker's entire theory about the world. Holists might might even profess a skepticism toward meanings of any kind, as Davidson suggests. In a holistic theory, for example, part of the meaning of “horse” might be that horses are animals but not cows. As such, this direction of meaning determination is better suited to epistemological rather than metaphysical inquiry: if you want the meaning of a word, look to its use.

The fact that holists and atomists approach meaning with epistemological and metaphysical approaches, respectively, creates difficulties for fruitful comparison. One is tempted to say that they are simply talking about different, but not necessarily incompatible notions of meaning. We might say that the atomists

17 This argument is discussed in §3.2.

are concerned with literal meaning, the holists with figurative; the atomists with connotation, the holists with denotation etc. But this overlooks the very goals of semantics. These theories are not just giving a story about way symbols of a natural language are combined, but also a story about how things in the world are combined. Theories about language are perhaps in a unique position of offering us insight into metaphysical questions which lie beyond our sensory perception. If natural languages necessarily frame our conceptual representation of the world, then insight into the principles of natural languages work just is insight into how the world must be (for us). Entwined with this is the wonder at how quickly two year olds are able to master this infinite resource.

I think that the hope that we can gain this insight through the principle of compositionality rises from a deeply held belief that the rules of semantic composition parallel the rules of syntactic composition. Jaakko Hintikka makes a similar assessment: he sees an entire semantic tradition resulting from Frege and Russell which draws atomistic conclusions from a close connection between syntax and semantics but argues that there is no a priori reason to make this connection.[18]

In his early work, Chomsky argued that syntax is independent of semantics, and illustrated this with “colourless green

18 Hintikka, “A hundred years later”, 40-1.

ideas sleep furiously".[19] The idea is that this sentence is syntactically well-formed, but meaningless. Of course, examples like this as not decisive. Fiona Cowie points out that some current theories view such sentences as ungrammatical.[20] Nevertheless, the early work of Chomsky arguably advanced linguistics a great deal by showing how syntactic components can be defined without any appeal to semantic values. The conceptual separation between the two also seems to work the other way; some expressions seem to carry a meaning even if they are syntactically ill-formed. For example, most English speakers can understand Ralph Wiggum's ill-formed exclamation, "Me fail English? That's unpossible".

If it could be demonstrated that the rules of semantic combination were parallel to the rules of syntactic combination, that would count as a definitive reason to prefer some semantic theories over others. Take, for example, Barbara Partee's preferred interpretation of compositionality: it is not an empirical claim but rather a good working principle.[21] This is not just a good working principle for linguists, mind you, but for every language user qua folk theorist about semantics: "we hold the principle of compositionality constant in working out (unconsciously) what shifts [in lexical meaning] our interlocutors may

19 Chomsky, Syntactic Structures, 15.
20 Cowie, "The logical problem of language acquisition", 37.
21 Partee, "Formal semantics, lexical semantics and compositionality", 7.

be signaling".[22] From there, it might be argued not just that everyday language users unconsciously employ compositionality as a good working principle, but that they must do so. Thus, compositionality might be an innate linguistic principle, argued for on the basis of poverty-of-stimulus arguments like Chomsky's Universal Grammar. It can thus been seen as part of what is called the "logical problem" of language learning: the evidence that a learner is exposed to is insufficient for her to restrict her hypotheses about her language's syntax. Now, if mastery of a language requires mastery of an infinite semantics then the case might be made that there is a problem in explaining how a learner, qua folk theorist, obtains that mastery on the basis of the evidence she is presented. Fiona Cowie discusses this logical problem of language acquisition in support of Universal Grammar. She argues that at most the nativist is entitled to the hypothetical claim that if some plausible learning theory ascribes prior knowledge of Universal Grammar, then that knowledge must be innate.[23] Now, if knowledge of compositionality is, in some way innate, then maybe concepts really must be atomistic, and if concepts are atomistic, maybe the world has to be too.

I have no good argument that that is true. The best I can do is to offer suggestions that language learning must also require a mastery of semantics which is similar to our mastery of

22 Ibid., 16.
23 Cowie, "The logical problem of language acquisition", 33.

syntax. Language allows us to refer to the world. We often do this successfully and there are lots of such meanings that probably could not be communicated without a syntax. Consider the following:

> (A): The proposition (assume there is a unique one) expressed by “Please pass the salt you keep under your pillow to me next Thursday after work.”
> (B): The proposition expressed by “Put that salt here.”

Assume, just for this paragraph, that demonstrating the desired behaviour is sufficient to demonstrate that the meaning was grasped. The meaning of (B) could be grasped not only by someone who has never heard such a meaning expressed, but possibly even by someone who has never heard its component parts expressed. Appropriate pointing could determine the reference of “that salt” and “here”; miming can indicate the action of putting; and acting can communicate the desire of an imperative. But in grasping the meaning of (A), it sure seems salient how the syntactic relation takes mere signs and forms them into grammatical terms. And the syntactic relation seems compositional. That amounts to a whole bunch of seeming. If, for example, Gérard Depardieu can express (A) in a game of charades, that is surely because his acting is syntacticly productive and it points us to the correct combination of meanings we already possess.

The next section will present to formal principles of meaning combination. One views complex meanings as a function of the meanings of the constituent parts. The other views constituent parts as equivalence classes defined by substitutions amongst complex expressions. The functional variation suggests atomism, the substitutional variation suggests sentential holism. Yet, I believe, both are able to account for semantic productivity.

1.3 Does Productivity Demand Atomism?

One of Fodor's central claims is "the principle of semantic compositionality demands semantic atomism". I do not think it does, and the goal of this section is to show that we should not be surprised if Fodor's claims for compositionality are wrong (or maybe just hyperbolic).

First I need to dispense with a way of reading this claim which makes it uninformative. In discussing compositionality above, I identified [C] as my working definition, and claimed that [C] entails atomism.

> [C] The meaning of a complex expression is determined by the meanings its constituents have individually and the way those constituents are combined.[24]

With this definitions strict one-way determination of meaning from simple to complex it might seem like a supposed analytic

24 Szabó, "The case for compositionality".

truth to say that compositionality demands atomism, but that is not my target. Instead, we should disambiguate Fodor's claim into this inference: productivity demands compositionality and compositionality demands atomism. Given a definition of compositional such that it entails atomism, the question is whether productivity requires atomism. I do not think it does.

First, it might be expedient to note that perhaps very few beyond Fodor mean anything like [C] when they use the word "compositionality". There are a number of ways that the meaning relation could be called compositional. Pagin and Westerståhl identify functional and substitutional variants of the principle.[25] Functional versions take it to mean that there is some function which takes component meanings as inputs and yields combined meanings as output. Substitutional variants take the principle to mean that if you replace a component part with a synonym, you will not change the meaning of the whole.

Pagin and Westerståhl view compositionality as a formal syntactic point — it requires that languages be semantically interpreted, but does not require a certain interpretation. As such, both are equivalent under the domain principle (DP): subterms of meaningful terms are themselves meaningful terms. The suggestion is that adherence to (DP) is what is required to explain productivity. Pagin and Westerståhl thus argue that the functional

25 Pagin and Westerståhl, "Compositionality I".

and the substitutional variants are both able to account for the semantic evaluation of complex utterances on the basis of the meanings of their components.

However, I think that when we apply these two principles in a semantic theory, we see an important difference between the two: they differ in the direction of determination. Substitutional versions allow the meaning of complex expressions to partially determine the meaning of components, and functional versions do not. I think that when we interpret the functional definition as a semantic theory, as what language users do when they evaluate sentences for example, it requires that there be atomic meanings which are independent of any other meanings. The substitutional variation, on the other hand, does not require this atomism.

I present these definitions in a recursive fashion: first by presenting my version of them, then providing an informal explication of their components. My version comes from choosing amongst interpretations which Pagin and Westerståhl offer in their own explication and substituting these choices back into the definition. My goal is to interpret these not just as capturing a “formal syntactic point” about productivity, but as applicable semantic principles. I do not offer much in the way of justification of the interpretative choices I opt for, other than that they seem natural. Without further ado:

> Funct(μ) For every syntactic rule $\alpha \in \Sigma$ there is a meaning operation Υ_α such that if $\alpha(u_1,...,u_n)$ has meaning, $\mu(\alpha(u_1,...,u_n)) = \Upsilon_\alpha(\mu(u_1),...,\mu(u_n))$
>
> Subst($\cong_\mu$) If $s[u_1, ..., u_n]$ and $s[t_1, ..., t_n]$ are both meaningful terms, and if $u_i \cong_\mu t_i$ for $1 \leq i \leq n$, then $s[u_1, ..., u_n] \cong_\mu s[t_1, ..., t_n]$[26]

Here μ is the meaning relation. It relates grammatical terms to semantic values. Grammatical terms may be either complex expressions or atoms. Σ is a set that has a subset of partial functions which acquire grammatical expressions from grammatical atoms. Thus the functional variety says that if a syntactic rule α yields a complex grammatical term, some operation, Υ_α, will apply the meaning relation to each of that expressions grammatical components. The meaning given by μ on the complex grammatical term, then, is equivalent to the meaning given by the meaning operation Υ_α to the meaning of each component.

So, if we consider how this definition assigns meaning to a novel expression, we see that Υ_α can act as a function whose domain is the meaning of each component and whose range is the meaning of the complex expression. This formulation is trivial if we allow μ to assign the same meaning to every term.[27] Thus I suggest reading into the definition the semantically recursive con-

26 Ibid., 254.
27 Ibid., 254.

dition, which requires the function to assign more than one meaning.[28] The functional variety of compositionality then requires that μ be applied to a expression's grammatical components to yield possibly distinct meanings.

Pagin and Westerståhl point out that the functional definition presupposes the domain principle (DP) because it presupposes that μ can be applied to each component part to give a meaning. Read in a way which makes it semantically (rather than only syntactically) recursive, this entails a metaphysical commitment to the prior existence of atomic meanings which can be mapped onto atomic grammatical terms by μ, because μ will be applied until there are no more component parts, i.e. until it assigns meaning to the atomic terms, and it will yield at least some distinct meanings. This variation thus draws the close connection between syntax and semantics which I believe is behind atomistic intuitions about compositionality.

Pagin and Westerståhl also point out however, that the determination of meaning from grammatical terms does not need to specify anything about the semantic values. Using a definition of synonymy as grammatical equivalence, they are able to define the substitutional variety of compositionality. Under this definition, terms are synonymous only if they are grammatically identical. For example “bachelor” is synonymous with “bachelor”

28 Ibid., 255.

but not with “unmarried man”. Defining synonymy grammatically yields a version of compositionality which is far too weak to explain semantic productivity so I defined synonymy in the definition as requiring a congruency of meaning, here symbolized $\approx_{\mu}$.[29] This requires that there be a least some expressions which are assigned the same meaning, and allows both grammatically equivalent and non-equivalent terms to be assigned the same meaning. Equivalence classes of meaning can include words which are not type-identical, so “bachelor” might be assigned the same meaning as “unmarried man”.

Reading synonymy as congruency also requires that we postulate meanings at some level of complexity. However, unlike the functional version of compositionality, it does not require that there be independent semantic atoms. It is committed only to the existence of something like complex propositions which have component parts. Furthermore, since it assigns meanings only to the “derivation histories” or “analysis trees” of complex terms, it cannot allow for the existence of the most general proposition (the monistic “absolute”, perhaps).[30] Therefore, it is committed to what we might call a “mid-level” existence of meanings.

Now, if (DP) is assumed, both varieties show how the meaning of complex expressions can be determined by the meaning of

29 Ibid., 257.
30 This point also applies to the functional variety. Ibid., 252.

component parts. Referring back to Subst($\cong_\mu$): if every u_i and t_i is a meaningful term, then we can substitute just one u with a synonymous t without changing the meaning of the complex expression. The meaning of the complex expression depends on the meanings of its component parts in the sense that if non-synonymous component parts were substituted in, the meaning would change or disappear. Because of this equivalence under (DP), both Funct(μ) and Subst($\cong_\mu$) account for formal determination of the meanings of complex expressions by the meanings of their components. If productivity only requires (DP), then Subst($\cong_\mu$) accounts for productivity just as well as Funct(μ). Furthermore, since the Subst ($\cong_\mu$) does not presuppose (DP), it is more general. This has the result that Subst($\cong_\mu$) can more readily deal with apparent counterexamples where the subterms of meaningful terms are apparently not meaningful.

Even though both the functional and the substitutional varieties are formally equivalent in explaining how the meaning of complex terms can be constructed from the component meanings, reading them as requiring at least some different meanings and at least some synonymous meanings results in incompatible metaphysical entailments. In requiring atomic semantic values to map onto grammatical primitives, the functional variety entails atomism. In only requiring the existence of semantic values to map onto

grammatical terms at some mid-level point of complexity, the substitutional variety suggests holism.

Part of the difference between the functional and the substitutional ways of assigning meaning is in the distinction between analysis and reduction. Both the function and the substitutional varieties analyze the meaning of components from the meaning of a complex expression. The functional variety, however, also “reduces” meaning in that it is committed to the independent existence of atomic meanings which are the result of that analysis. Such atomic meanings cannot change with the addition of new expressions. Under the substitutional view, on the other hand, the results of the analysis might be viewed as “merely” theoretical constructs. They are theoretical in that the meanings assigned to component expressions may change with the addition of new expressions to the set of evaluated expressions. Note that, as a explanation of novelty, both varieties will need to allow us to predict the meaning of new expressions based on previously evaluated expressions. Chapter 4 will explicate how Brandom's substitutional account of productivity can allow for such novel evaluations. If successful, this shows that productivity requires analysis of meaning, but not a reduction of it.

As long as a meaningful expression has syntactic components we can analyze the meaning of the component parts. A commitment to (DP) seems to be just this commitment to the possibility of

analysis. Proponents of substitutionality are committed to some complex expressions being analyzable but functionality entails a commitment to all complex expressions being so analyzable. A further difference is found in whether analysis (when possible) yields defeasible evaluations. The reductionism of functionality does not seem to allow that it can. As an explanation of language learning and novelty, what is interesting is not that the functional variety is reductive, but that it allows us to induce over results and successfully evaluate novel expressions. Atomists could hold that we could only perform these inductive applications of μ under a reductionist view.[31] The induction of novel meanings over component meanings, in such a view, is more like mathematical induction than the sort which leads to probable conclusions.

Pelletier points out the similarity between the principle of compositionality and inductive definitions but argues that non-inductive definitions can achieve the same results as long as they are "semantically grounded".[32] To see how Funct(μ) could be instantiated inductively, let us take another look at the final clause which defines the meaning of the complex expression in terms of the meanings of its parts:

31 We could also employ an analysis/synthesis distinction, in lieu of analysis/induction, here. Roughly, analysis "works backward" to assign meanings to component parts, and synthesis "works forward" to assign meaning to a complex expression. Then, synthesis is what we are really interested in as an explanation of productivity, and the proponents of functionalist compositionality can be seen as claiming that synthesis is not possible unless analysis is reductive.

32 Pelletier, "The principle of semantic compositionality", 21-23.

$$\mu(\alpha(u_1,...,u_n)) = \Upsilon_\alpha(\mu(u_1),...,\mu(u_n))$$

Now, let us assume that $\alpha(u_1,...,u_n)$ is a novel expression, for which we do not know the value of μ. Assume that we already have the values of $\mu(u_1)$ to $\mu(u_n)$, i.e. that we grasp the meaning of the components. What we want is a definition of the operation Υ_α over those component parts, since this is equivalent of the operation μ over the novel expression. First we define a basis clause:

$$\Upsilon_\alpha(\mu(u_i)) = \mu(u_i)$$

and an inductive clause:

$$\Upsilon_\alpha(\mu(u_{i+1})) = \mu(\alpha(u_i)(u_{i+1}))$$

And finally an extremal clause:

$$\Upsilon_\alpha(\mu(u_1),...,\mu(u_n)) = \mu(\alpha(u_1)) \ ... \ \alpha(u_n))$$

This definition might formally capture how we can parse a novel expression. It assumes we know α, which is just the expressions syntax. We want to apply that same operation to ($\mu(u_1)$,..., $\mu(u_n)$), the meanings of the components of the novel expression. The base clause ensures that applying Υ_α to any simple component parts yields just the meaning of that component part. Υ_α combines meanings in a way congruent to how α combines grammatical terms. The inductive clause applies Υ_α to grammatical terms of increasing complexity, and the extremal clause ensures that the opera-

tion stops when it reaches the final grammatical term of the expression.

What is important is not that this mode of evaluating novel expressions is inductive, but that it is semantically grounded. It is semantically grounded because it is non-circular. It would be circular, for example, if we defined:

$$\Upsilon_{\alpha}(\mu(u_{i+1})) = \alpha(\mu(u_i))\mu(\mu(u_{i+1}))$$

Or if we defined:

$$\Upsilon_{\alpha}(\mu(u_i)) = \alpha(\mu(u_i)).$$

Thus, this functionality version of compositionality avoids circularity because Υ_{α} and μ are distinct operations, even though both assign meanings to grammatical terms. μ is a “first level” operation that assigns meaning, and Υ_{α} is an “n-level” operation which inductively assigns meaning from μ.

Because Funct(μ) is non-circular, it is semantically grounded. Pelletier suggests that it is <u>this</u> feature which proponents of compositionality are after in order to explain productivity, novelty and learnability. Using, as an example, a definition of an operator in sentential logic non-inductively but still validly, Pelletier shows that this groundedness can be achieved without requiring that <u>only</u> the meaning evaluations of the components be used in evaluating the meaning of the complex expression. He suggests that a “semantic evaluation can, in general,

bring into play all kinds of facts."[33] This includes facts that are not part of the meanings of the components, such as facts about context, inferences and knowledge of the world, as long as it does so in a manner which is semantically grounded.

Fodor's arguments against holism, however, often turn on alleged issues with the holistic theory of concept individuation. Fodor argues that inferential role semantics cannot explain how we can possess complex concepts corresponding to AN-phrases, like "brown cow" simply by possessing the concepts corresponding to the adjective and the noun. One argument to this effect is that if I infer from "Betsy is a brown cow" to "Betsy is dangerous" but you do not, then we cannot be talking about the same things by "brown cow" or "dangerous".

However, Pagin argues that a many-one inferential holism can explain how these complex concepts are composed of their component concepts.[34] We take the assignment of meanings as a two step process: assigning a global role and assigning a meaning. The global role is a set of inferential pairs between concepts. Meaning is assigned by a function which operates on the domain of global roles and yields the range of semantic values. If that function is one-one, whereby it cannot assign the same meaning to any two distinct global roles, then it has the absurd consequence that changing any of your inferential pairs changes the semantic

33 Pelletier, "The principle of semantic compositionality", 22.
34 Pagin, "Is compositionality compatible with holism?".

value assigned. However, allowing the function to be many-one allows the meaning function to assign the same meaning to distinct global roles. In that way, it allows us to both assign the same meaning, for example, to “dangerous”, even if you inferentially pair it with “brown cow” and I do not. This draws a distinction between meaning and concept possession, which deflates this argument against the inferential-role theory of concept possession. For Fodor, the meaning is identified with the concept possessed. Inferential role semantics must have a different notion of meaning which differentiates it from the concept in order to explain meaning productivity. But it can be productive.

In summary, one way of arguing that Brandom does not violate the requirements of productivity is to show that his semantics is “compositional” in the substitutional sense, which is equivalent to the functional sense as an explanation of productivity because it can also adhere to (DP). Another way of arguing this point is to show that his semantic evaluations are “grounded”, which is to say non-circular. An even further way is to show that his theory is semantically productive because semantic values are not identified with the inferential roles.

Many similar points are argued for by Henry Jackman in a paper which defends Davidson from Fodor’s criticism.[35] Jackman argues that inferential role is only refuted if we take it to mean

35 Jackman, “Compositionality and semantic properties”.

that a term's meaning is strictly identified with its inferential role.[36] But that would be something of a straw-man, since the claim of inferentialists is simply that inferential role plays some constitutive role in a term's meaning, not that it plays the only role. Jackman argues that there is no inconsistency in being an "atomist" in describing what the semantic values of our words are, while providing "holist" theory about how those semantic values came to be.

The key is different notions of semantic value. Brandom might be "atomist" about semantic values in that semantic evaluation can proceed by analysis until it evaluates some kind of "smallest part". That is what we need to explain productivity. However, he cannot be an atomist about semantic values in the ontological sense whereby those atomic parts at the end of analysis correspond in a one-one relationship to things in "the world". Fodor is atomist about semantic values in this ontological sense. Another difference is that, for Brandom semantic values are defeasible and can be altered under new evaluations. For Fodor, they cannot. Defeasible semantic values, however, can still be semantically productive.

Since Fodor and Brandom are both free to conceive of semantic values in any way they want, it should come as no surprise that productivity alone does not demand ontological atomism.

36 Ibid.

Thus, in evaluating Fodor's claim against Brandom, we need to consider more than just his conception of compositionality, but also his related arguments against pragmatism.

Finally, it should be conceded that not every holism can account for productivity. I have already mentioned that a holism which is committed to the existence of "the most general proposition", from which all other meanings are defined, could not be compositional in either the functional or the substitutional senses. Pelletier also draws a useful distinction between holism and wholism. Wholism cannot account for productivity because it is incompatibile with (DP). But non-wholist holisms are and can.

Wholism holds that some "properties can only be attributed to entities that are not individuals."[37] Holism holds that "some properties of an object are defined in terms of the same type of property of some other object(s), and these properties of the other object(s) are in turn defined by means of the first property."[38] Here we are interested in the property "being a meaning of", and the entities are grammatical terms. Contrary to wholism is the atomistic view that a "whole in a compositional system is built up from materials in the parts."[39] Contrary to holism is the compositional view that "The μ of a whole is a function of the μ's of its parts and the ways those parts are combined."[40] Put in

37 Pelletier, "Holism and compositionality", 5.
38 Ibid., 6.
39 Ibid., 2.
40 Ibid., 3.

terms of meaning, the atomistic view denies that anything other than the meaning of the components can contribute to the meaning of the complex expression. The compositional claim denies that the meaning of a grammatical term can be defined by the meaning of some other grammatical term. Fodor is committed to both the compositional claim and the atomist claim.

Brandom, however, is committed to the holist claim but not the wholist claim. He is not committed to the view that meaning can only be attributed to complex grammatical terms. Meaning can be attributed to simple grammatical terms, it is just that this meaning attribution may need revision on the basis of new evidence. Fodor takes this to mean that Brandom's semantics amounts to confirmation holism. As a semantic theory, confirmation holism holds that the hypothetical assignment of meaning to a grammatical term cannot be confirmed or denied in isolation of all other meaning assignments. The confirmation-holistic view of meaning is incompatible with the atomist claim that the meaning of a component part cannot be defined in terms of the meanings of other component parts. Confirmation-holism is holism (about the properties of things), but it is not wholism (which is an ontological thesis about the things themselves). As such, it is compatible with, but not committed to, the atomistic view that the meaning of an expression is built-up from the meaning of its parts. Brandom's holism, then, does not commit him to wholism. He does, of course,

deny atomism. However, he is able to explain productivity by making the property of “being a meaning of” to be be an epistemic property. Although he is committed to an ontology of propositions, he does not claim that only propositions fulfill the property of being a meaning. Therefore he does not make the wholist claim that meaning can only be attributed to complex expressions.

Summary

The goal of this chapter was to provide an analysis of the central concepts employed in this thesis: compositionality, atomism, holism, and linguistic productivity.

§1.1 suggested the use of this definition of compositionality from Szabó: [C] The meaning of a complex expression is determined by the meanings its constituents have individually and the way those constituents are combined. I argued that this definition entails a commitment to semantic atomism, and as such, can serve as a formal definition of what Fodor is committed to. We also saw how this strong definition is susceptible to counterexamples from natural language which do not seem strictly compositional: ambiguous constructions, adjective-noun phrases, and synonymous phrases.

§1.2 identified semantic atomism as the thesis that there are atomic meanings which are independent of all other meanings. Holism is the denial that there are such independent meanings —

you cannot grasp just one meaning. I suggested that the argument for semantic atomism from compositionality might be seen as congruent to “poverty of the stimulus” arguments for an innate Universal Grammar. The suggestion is that we could not learn language without employing the principle of compositionality, and we could not successfully employ the principle of compositionality unless our concepts were atomic and independent of other concepts.

§1.3 introduced two approaches to productivity: functional and substitutional. I argued that both could capture the productivity of natural language, but that the substitutional version does not require atomism. Holism is able to explain how the meaning of complex expressions can be constructed from component meanings, given certain constraints. For one, the assignment of simple meanings from complex ones must not be circular. The meaning of “dog”, for example, cannot depend on the meaning of “animal” which in turn depends on the meaning of “dog”. Secondly, the meanings of component expressions cannot be identified with their inferential role. Rather, we can only say that the meaning supervenes on the inferential role, or is partially dependent on the inferential role, or something similar. Last, holism cannot be committed to the “wholist” claim that meaning is only attributable to complex grammatical terms. Such a theory would clearly be

unable to explain how it is we parse complex meanings on the basis of their components.

Chapter 4 will explicate my reading of Brandom's inferential role semantics. Admittedly, very little has been said about what this "inferential role semantics" is, and how language users might actually employ it. So far, my claim is hypothetical: If Brandom's semantics able to obey the constraints set forth in §1.3, then, it is able to account for productivity and is not refuted by Fodor's arguments.

2. A Brief History of a Humble Property

The notion that we can we can grasp novel meanings because they are built from their component parts is oft-mentioned in philosophy of language. Bertrand Russell hinted at compositionality as support for his logical atomism, calling it an "apparently humble property" of language which allows us to understand novel propositions.[1] Wittgenstein, on the other hand, begins his Philosophical Investigations with a long and detailed critique of exactly this atomistic view which so many philosophers at the time (his younger self included) often took for granted. This chapter will show how notions of compositionality appeared in early analytic philosophy before any precise notions of it were worked out. Compositionality was often treated as a fact of language as readily apparent as the productivity of language itself. It is often mentioned and employed as support by philosophers without much in way of explication. This tends to give the illusion of agreement — apparently everyone believes in compositionality. As we will see, however, the relatively uncontroversial thesis that language is productive ought to be separated from the more problematic thesis of semantic atomism.

Along with tracing a relatively unquestioned identification of productivity with atomism, this chapter will also trace skeptical criticisms of semantic atomism. In Plato's Theaetatus we find an epistemological dilemma posed for the atomistic mode of

1 Russell, The Philosophy of Logical Atomism, 54.

composition. This is the problem of the one and the many. Should we call a complex whole “one” insofar as it is a unified whole, or should we call it “many” insofar it is an amalgamation of parts. Either answer is problematic, insofar as the atomic parts are meant to provide a foundation for explaining the whole.

In Wittgenstein’s Philosophical Investigations we can find two broad sets of criticisms against semantic atomism. First, in his opening critique of the Augustinian picture of language Wittgenstein expresses a skepticism that atomic parts of language can themselves be meaningfully identified and talked about. The suggestion, then, is that what we call meaning of any given word be at least partially conceived in terms of its use. A positive argument for this non-atomistic notion of meaning can be seen in Kripke’s interpretation of Wittgenstein’s private language argument. When we consider that the meaning of a word ought to determine its future use, we see that the atomistic theory of meaning is incoherent.

Bertrand Russell worked out his philosophy of Logical Atomism unaware of Wittgenstein’s new found skepticism towards such projects. However, he did struggle with a problem strikingly similar to Plato’s dilemma. Russell was concerned with explaining what it is that holds a proposition together as a unified, meaningful whole. Russell felt that he failed to do so, but this problem is at least partially addressed by the multiple relation

theory of judgment, whereby a propositional attitude like “x believes a is b” is seen as a relation between three terms, x, a and b, rather than a relation between a subject and a proposition. I believe this feature allows Russell to avoid Plato’s problem though Fodor cannot. In response to Wittgenstein, however, both Fodor and Russell could only reply that epistemological doubts cannot entirely discredit their metaphysical approaches to meaning.

Finally, the roots of both Russell’s atomism and Wittgenstein’s holism might be found in the work of Gottlob Frege. The tension is aptly captured in two principles, compositionality and contextuality, both of which have been credited to Frege. While compositionality suggests that the meanings of words are prior to the meaning of sentences, contextuality suggests the opposite. This has caused interpretative difficulties for Frege scholarship. However, recent work has shown that Frege was not committed to an atomistic notion of compositionality, and only maybe committed to a sentential holism.

Brandom sees his own work in semantics as following from his interpretation of early analytic philosophy. Specifically, he does not read the Philosophical Investigations as offering a philosophical quietism in place of analytic philosophy. Rather, he sees Wittgenstein as explicating and defending an alternative view of analysis, one which starts with sentences and is not re-

ductive. It is, perhaps, compatible with many of the goals of Russell's atomism, but without his problematic metaphysics. It is this interpretation of early analytic philosophy that I wish to put forward, in defense of the overall argument that sentential holism is a plausible approach to linguistic theory.

2.1 Socrates' Dream

True to Whitehead's adage that the safest general characterization of Western philosophy is as a series of footnotes to Plato, the historical roots of semantic atomism and holism can be found in the Theaetetus. Indeed, Gilbert Ryle goes so far as to read the movement of logical atomism back into that work.[2] Specifically, he claims that the logical atomists took up a position which Socrates relates to us as a dream:

> I seem to have heard some people say that what might be called the first elements of which we and all other things consist are such that no account can be given of them. Each of them just by itself can only be named; we cannot attribute to it anything further or say that it exists or does not exist, for we should at once be attaching to it existence or nonexistence, whereas we ought to add nothing if we are to express just it alone. [...] But when we come to things composed of these elements, then, just as these things are complex, so the names

2 Ryle, "Logical Atomism in the Theaetetus".

> are combined to [give an explanation][3] Accordingly, elements are inexplicable and unknowable, but they can be perceived, while complexes [...] are knowable and explicable, and you can have a true notion of them.[4]

Before looking at how Ryle sees logical atomism as a version of Socrates' dream, I would like to put it in a broader context. Verity Harte's Plato on Parts and Wholes discusses this passage in the context of Plato's concern with "the problem of the one and the many".[5] The problem manifests in a number of different paradoxes which arise from a model of composition whereby a whole is identified with its parts. The paradoxes arises because a whole seems to be both one thing and many things. It is many things by way of being composed of parts. It is one thing by way of those parts being unified.

It is this model of composition which Plato is presenting in the Theaetetus, but it is a model which, Harte argues, he ultimately rejects. Harte identifies the central thesis claimed by Socrates' dream as the "Asymmetry Thesis" which is "the thesis that there is an epistemological asymmetry between elements and complexes: elements are unknowable, complexes are knowable.[6] She reads Socrates discussion of the dream as providing two refuta-

3 Plato here uses the word "logos". I render it as "an explanation" to match discussion below, and follow Harte in holding that alternate renderings of it do not affect this discussion.
4 Plato, Theaetetus. (202c-202c).
5 Harte, Plato on Parts and Wholes, 3.
6 Ibid., 33.

tions of the Asymmetry Thesis: a dilemma which shows that the primitive elements are just as knowable or unknowable as the complexes and an argument from experience that the primitive elements are actually more knowable than the complexes.[7]

What counts as knowable stems from the notion that if something is known, we should be able to give an account or explanation of it. An explanation, according to Socrates, can mean one of three things: (1) giving your opinion on the object of explanation; (2) enumerating the parts which make up the object of explanation; or (3) describing the mark of difference between the object of explanation and all other things.[8] According to the Asymmetry Thesis of the dream, the primitive elements are only recognizable as such. This is similar to Russell's idea of direct acquaintance. Now it might be that since they are recognizable as such, we could describe the mark of difference between one element and other elements, along the lines of (3). However, the complex wholes cannot only be explained according to (3), but also according to (2), since we can enumerate their parts. The complex wholes also seem more susceptible to explanations along the lines of (1) than elements. So, even if we say that the elements are "knowable" because "recognizable", it seems reasonable to hold that the complex wholes are "more knowable" because we can give more of an explanation of them.

7 Ibid., 33-4.
8 Plato, Theaetetus. (206c-208c).

Harte's explication allows for a concise formulation of the dilemma: is a syllable the same as its letters or is 'some single form' resulting from the combination of the letters? Assume the syllable is the same as the letters. Then if we know the syllable we must also know the letters, since the syllable just is the letters. Assume a syllable is a single form which is different, but results solely from, its letters. Then the syllable, as a single form, does not have parts. So, if the letters are not knowable, as the dream suggests, the syllable cannot be either. Both horns of the dilemma depend on identifying a whole with its parts. Contrary to the Asymmetry Thesis, both the whole and the parts are equally knowable (first horn) or equally unknowable (second horn).[9]

Ryle argues that the logical atomists, by which he means Meinong, Frege, Moore and Russell, conceive of particulars and facts in a way congruent to the model of syllables and letters related in Socrates dream.[10] Particulars are held as knowable only in the sense of direct acquaintance but facts are complexes which can be explained through their parts. However, if the same parts of a given fact are combined in a different way they yield either different facts or no facts at all. This suggests that facts are also a single form arising from its parts, and this single form

9 Full discussion of the dilemma, which places it in context to avoid prima facie fallacies, is found here:
Harte, Plato on Parts and Wholes, 35-39.
10 Ryle, "Logical Atomism in the Theaetetus", 33.

is itself a particular. The logical atomists therefore find themselves susceptible to the dilemma Socrates presents.

According to Ryle, logical atomism is concerned with the composition of truths and falsehoods.[11] Truths have an internal complexity and knowledge of them can be stated. The simple, nameable objects (particulars) of which they are composed, however, can only be known through direct acquaintance. We can state our knowledge of facts because there are proposition which, when true, correspond to facts. These propositions must also themselves be complex wholes that can be analyzed into names which correspond to particulars. Because propositions can be analyzed this way, they must be identified with their parts in a way similar to the identification of syllables with their parts in the dream. This gives the Logical Atomist position all three aspects of Socrates' dream: facts are nothing more than the particulars from which they are composed; facts can be stated by propositions and can be explained through the particulars from which they are composed; but particulars are only knowable in the sense of direct acquaintance.

Now consider propositions, which are composed from names and (when true) correspond to facts. Faced with Socrates' dilemma, logical atomists must ask whether a sentence just is the names of which it is composed, or whether it is a higher-order

11 Ryle, "Logical Atomism in the Theaetetus", 30.

name of a complex object. Ryle claims that although propositions are complex wholes, the atomists also see them as higher-order names which correspond to objects called “objectives”, “propositions-in-themselves”, or “facts”.[12] This follows from considering the proposition corresponding to “Brutus killed Caesar”. Ryle sees the atomists as viewing this, on the one hand, as a complex consisting of “Brutus”, “Caesar” and “murder”, but also, on the other hand, a higher-order name which combines them in a particular way.[13] This is because combining them in another way, namely “Caesar killed Brutus”, corresponds to a different proposition —— as can be seen by the fact that it is no longer true. This amounts to accepting the second horn of the dilemma, that propositions are a single form and which entails that propositions are only knowable in the sense of direct acquaintance, like names.

This, by Ryle's reckoning, leads the atomists to be concerned with a question which also puzzles Socrates but should not normally puzzle English speakers, namely “how can I think falsely?” This puzzles the atomists because thinking falsely is thinking about an object, namely a proposition, which is not.[14] This is puzzling because the proposition is not only a complex but also a name knowable through direct acquaintance. Thinking “Brutus killed Caesar” is thinking about what <u>is</u>, because it cor-

12 Ibid., 33.
13 Ibid., 33.
14 Ibid., 33.

responds to a proposition which corresponds to a fact. Thinking "Caesar killed Brutus" is thinking about what is not because it corresponds to a proposition which does not correspond to a fact (but must somehow still exist if one is able to think about it). Ryle thinks that English, unlike Greek, should not normally compel us to think about this since does not normally use "is" and "exists" interchangeably.[15]

Ryle mentions that he finds Russell's knowledge by acquaintance and knowledge by description to be largely unsuccessful at capturing two distinct senses of knowing.[16] Perhaps this puzzlement about "thinking what is not" shows why Ryle thinks this. In both cases, knowledge requires representation, and even thinking falsely requires representation. Ryle himself is often praised (or blamed) for instead drawing attention to the distinction between knowing-how and and knowing-that. He argues that knowledge-how cannot be defined in terms of knowledge-that, and that knowledge how is logically prior to knowledge-that.[17] In this way, thinking, for Ryle, does not need representation, so we need not be concerned with how we can "think what is not". Part of the wide influence of this doctrine can be seen in semantic theories which define the meaning in terms of use, as Brandom does. Fodor is particularly bothered by this thesis, and flatly denies that

15 Ibid., 25.
16 Ibid., 27.
17 Ryle, "Knowing how and knowing that".

possessing a concept can be defined in terms of knowing its use, or that we can know how to do something, or even plan to do something, without knowing that something is the true, or at least representing the world in some way.

Whether or not Ryle's interpretation of atomism is fair (I argue that it is not in §2.3), it seems plausible to view atomism as something like Socrates Dream. To hold that the analysis of propositions is possible can be seen as congruent to holding a view of composition whereby the whole is identified with its parts. Socrates' asks Theaetetus what alternative there is, other than to hold that the syllable is not the letters, and the letters are not parts of the syllable. But that is absurd. If we apply such an analogy to linguistic meaning, for example, it is akin to saying that the meaning of a sentence is independent the meaning of the words which compose it and that is contrary to productivity.

However, Socrates does not endorse this view either, at least, not in this dialogue. Harte suggests that he endorses the other refutation of the dream: the reversal of the Asymmetry Thesis. This is an argument from evidence that it is actually the elements which are knowable and the complex which is unknowable. Harte sees this refutation of the dream as partial work toward an alternative model of composition which is presented in other dia-

logues, notably the Parmenides. What is to be rejected as problematic is the identification of a whole with its parts.

The rejection of the Asymmetry Thesis is foreshadowed earlier in the dialogue, when Socrates reminds the audience that the great philosophers tell us that the objects of nature, which are constantly being created and destroyed, cannot be fixed, or brought to a standstill by names, and that "anyone who talks so as to bring things to a standstill is easily refuted"[18] Identifying the whole with its elements amounts to trying to fix it with a name. However, having elements, the whole is subject to change with a change of any of its parts. So anyone who does try to fix it with a name is easily refuted and the identification of a whole with its parts remains a dream.

2.2 The Limits of Language Games

Ryle's employment of the Theaetetus to criticize logical atomism is not unique. Wittgenstein had already made such a connection in his Philosophical Investigations. The Investigations open with an extended critique of atomistic notions of meaning, in preparation for the suggestion that meaning should be conceived in terms of use.

One theme that the later Wittgenstein carries over from the earlier Wittgenstein of the Tractatus Logico-Philosophicus is the

18 Plato, Theaetetus. (157c-157c) p. 862.

attempt to draw the boundaries of meaning. A central difference lies in how each work draws this boundary. To make a comparison, it is enough to be reminded of the broad outlines of the Tractatus. Wittgenstein begins that work by delimiting the world as the domain of meaning: the totality of the world defines what is or is not the case; the world divides into facts; the world just is the totality of facts and their being all the facts; the facts are in logical space.[19] The world is thus defined so as to be compositional in a way which it can correspond to a compositional and fully extensional language. This world, in turn, is what we can speak meaningfully about. He ends the Tractatus with the dictum that "what we cannot speak about we must pass over in silence."[20] He begins his Investigations by criticizing this notion of meaning.

The Investigations start with a picture of language he attributes to Augustine where "the individual words in a language name objects", "sentences are combinations of such names" and "every word has a meaning" which is just the object for which it stands.[21] This is similar to the Tractatus where propositions project a possible situation.[22] That possible situation is actual if it composes a fact, or state of affairs, out of objects (things).[23] Wittgenstein does not completely reject this picture

19 Wittgenstein, Tractatus Logico-Philosophicus, §1.
20 Ibid., §7.
21 Wittgenstein, Philosophical Investigations, §1.
22 Wittgenstein, Tractatus Logico-Philosophicus, §3.11.
23 Ibid., §2.01.

at the beginning of the Investigations. Rather, he suggests it is only partially correct. It "does describe a system of communication; only not everything we call language is this system."[24] The Tractatus drew the boundary of meaning too narrowly. The Investigations reveal just how much must be "passed over in silence" according to the former view, along with the observation that we do not pass these things over in silence. We do speak about them and this talk serves various uses. Many systems we legitimately call language are not captured by the narrow definition of the Tractatus.

So language can be used in a strictly denotational way where the meanings of words just are things in the world. This use of language can be viewed as compositional in an atomistic way. Wittgenstein is lead from considerations of naming toward considerations of composition in §39, where he discusses whether a name really ought to signify a simple. The suggestion is that the reference of a true (or fully analyzed) name cannot itself be composed of parts.

He takes this suggestion to follow naturally for those who hold that the meaning of a name is its reference, when faced with this conundrum: if a name means an object then that name would become meaningless were the object to be broken into parts; hence sentences using that name would be nonsensical; however such sen-

24 Wittgenstein, Philosophical Investigations, §2.

tences often do make sense. For example, the sentence “Excalibur has a sharp blade” continues to make sense even if Excalibur is broken into parts. So “Excalibur” itself is not a real name, but needs to be analyzed into words which name atomic parts.

A lesson drawn in §40 tells us not to confound the meaning of the name with the bearer of the name. The fact that “X is dead” is meaningful even if that person dies shows that the meaning of “X” exists even if the person does not. If the atomic parts, then, could cease to exist, then the meaning of the “real names” which result from the analysis of ordinary names is not exhausted by their reference. To the suggestion that a name ought to signify a simple, then, is added the suggestion that the simple have a kind of perpetual existence.

But if we deny that words signify simples, how can their meanings be specified? Wittgenstein suggests that for many words, but not all, it can be done through their use in a language game.[25] He contrasts this approach with one that attempts to specify the meanings of names from within a language, likening such attempts to what Socrates says in the Theaetetus: that the essence of speech is the composition of names so the bare elements which are named cannot be given an account. Examples of such bare elements are given as Russell's “individuals” and the “objects” of the Tractatus.[26] These simple parts of speech, in order to have

25 Wittgenstein, Philosophical Investigations, §41-43.
26 Ibid., §46

meaning, must refer to simple parts which do not cease to exist. This can be done, as it is in the Tractatus. But in the Investigations Wittgenstein is skeptical about whether these simples, in turn, correspond to reality.

Wittgenstein's skepticism suggests that those who hold that a name can be used to signify a simple, but perpetually existing atom of meaning find themselves in a chicken/egg sort of problem:

> "Asking 'Is this object composite?' outside a particular language-game is like what a boy once did, who had to say whether the verbs in certain sentences were in the active or passive voice, and who racked his brains over the question whether the verb "to sleep" meant something active or passive."[27]

The boy in question found himself ensnared in an unanswerable riddle because he misunderstood which language game was being played. When he wondered whether to sleep was to do something or not, he made the mistake of not moving from a language game played within his language (where the words "active" and "passive" denoted properties of things in the world) to a language game about his language (in which "active" and "passive" denoted grammatical aspects of his language). Speaking of names as properly only signifying simples can be seen as a language game philosophers play. Even Wittgenstein took part in the Tractatus. However, it is one thing to construct abstract, atomic "meanings"

27 Ibid., §47.

which are composable. It is quite another to then describe the relation between these atomic meanings and the actual world. The Investigations suggests that doing the former is tantamount to asking “is this object (the world) composite?” outside of any particular language game.

To the question “Is X composite?” Wittgenstein tells us that the correct answer is “that depends on what you understand by “composite”.”[28] What you understand by “composite”, in turn, depends on what you understand by “part”. He illustrates the large variety of ways such questions can be answered with the chessboard example. A chessboard seems to be a composite object, but what are its parts? We could say it is composed of 64 chess-piece spaces, or 204 squares, or 18 lines. But outside of any context these answers do not communicate any information about the chessboard. If we are inquiring about how the board relates to the game, then it is informative to say that it is composed of 64 chess-piece spaces; if we inquire as to how many regular quadrilaterals are to be found on the board, it is informative to answer that is is composed of 204 squares; if we inquire as to number of lines which are needed to create it then 18 lines is an informative answer.

Furthermore, there is a meaningful, and necessary, dependence relation to be found amongst these different ways of compos-

28 Ibid.

ing the chessboard. The game of chess requires that the board be composed of 64 spaces for the pieces to be placed in. But this is still only to single out that one mode of composition as important by placing in a context of use: learning the game might require knowing that the board is composed of 64 spaces. Other ways of composing might be necessary for other uses: solving a math problem might require knowing that it is composed of 204 squares; programming a lathe to etch boards might require knowing that it is composed of 18 lines. If we further ask “what is it composed of?” apart from any particular use of this question, there is no meaningful answer because we need a context of “use” to first define the whole.

But of course those who see the possibility of analysis as a clear fact of language need not view themselves not playing any language game whatsoever. Dummett explains one motivation for talking about atomic parts and composites of meaning: we might want to explain why some propositions seem knowable a priori, such as “all boys and girls are children”. Tarski drew attention to a distinction between meta-language and object language, which is useful here.[29] Given the language whose set of expressions contains the expression “every boy or girl is child”, we can then construct a meta-language where all individuals who satisfy the predicate “is a boy” also satisfy the predicate “is a child”.

29 Tarski, “The Semantic Conception of Truth”, 349-350.

This meta-language surely captures (at least) part of the meaning of the natural language predicates “x is a boy” and “x is a girl”. Furthermore, it is the relevant part for explaining why “every boy is a child” can be known a priori.

A difficulty arises if we are asked to provide an explanation for our meta-language. We can take “child” as a simple concept, and both “boy” and “girl” as composite concepts (of “is a male child” and “is a female child” respectively) or we can take “boy” and “girl” as simple concepts and child as a complex concept (“is a boy or girl”). However, there is no clear and obvious way to choose between them except by mere stipulation.[30] We simply choose one concept as basic and defined the others accordingly, putting an end to any further explanation.

Wittgenstein suggests that explanation rightfully comes to an end when it is no longer needed to prevent a misunderstanding.[31] In discussing how it is that we can explain what we mean by the name “Moses” by giving a description of him, Wittgenstein shows how further and further explanation might be demanded by an interlocutor. If we describe Moses as “the man, if there was such a man, who lead the Israelites out of Egypt” it can be further demanded that we explain “Egypt”. An interlocutor, however, not unlike a child who has recently learned the use of “why?”, might demand some aspect of this description of Egypt be explained, but

30 Dummett, Thought and Reality, 7-8.
31 Wittgenstein, Philosophical Investigations, §87.

such explanations come to an end when there is no more misunderstanding.

Thus, Wittgenstein's notion of a language game allows us to play language games whereby we appeal to the meanings of component parts to explain the truth value of certain sentences. But only to prevent misunderstanding. In order to prevent the misunderstanding that a boy is not a child, for example, we could provide a Tarskian truth-schema which shows that the expression "a boy is not a child" is always false and argue that this captures the relevant meanings of "x is a boy" and "x is a child". Now, as Dummett points out, this goal can often be achieved in more than one way. What Wittgenstein disallows is that there can be a further explanation which prioritizes one way of dividing the concepts over the other, because there is no apparent further misunderstanding that needs to be explained.

Although Wittgenstein's language-games do not rule out <u>all</u> uses of analysis, they do diminish the significance of the functional, atomistic principle of compositionality. The principle of compositionality can be seen as a way of stating the general form of a proposition: The meaning of an expression just is the meaning of its parts and the way those meanings are combined. Wittgenstein muses that if one repeats such a thing to oneself "one thinks that one is tracing the outline of the things nature over and over again, and one is merely tracing round the frame through

which we look at it".[32] Like Plato in the Theaetetus, what leads Wittgenstein to this conclusion is consideration of apparent epistemological difficulties in concluding that it really is the general form.

Now, it is one thing to accept that the form cannot be given further explanation. It is another to accept that because it cannot be given further explanation it is not the full picture. That it is not the full picture itself seems to be a metaphysical conclusion. Why should proponents of a functional version of compositionality, with its seemingly necessary entailments for the metaphysics of meaning, accept the conclusion, metaphysical in itself, that it is not the full picture of meaning because we cannot know that it is? In the next section, I present a reading of Russell which presents his logical atomism as largely unaffected by Wittgenstein's skepticism. In the next chapter, we will see that Fodor completely ignores Wittgenstein and instead charges that those who are inspired to define meaning in terms of use have things backwards.

We can find a defense of the alternative view of meaning as use later in the Investigations, where Wittgenstein considers rule-following and private language. These passages provide further motivation for the view that an atomistic semantics will be unable to explain important features of language. Wittgenstein

32 Ibid., §114.

here shows that the atomistic view is not the full picture of linguistic meaning by offering an alternative picture which captures aspects of meaning left out in atomism. The frame is redrawn around use rather than reference. A look at these passages is worthwhile because they are often employed as reasons for having to accept a holistic view of meaning. The argument inspired by Wittgenstein's considerations of rule-following is that any theory of language must account for both meaning and use but atomistic views which prioritize the meaning of subsentential components will be unable to adequately explain use.

Kripke's reading of these passages have come to overshadow alternatives. Following Kripke, it is common to see the argument against private language as starting around §201 of the Investigations.[33] Here, Wittgenstein presents a skeptical paradox that arises from considerations of whether rule-following can be construed as in introspective affair. Consider that knowing the use of a word is to know a rule governing its future application. What those who want to define use in terms of meaning want to say is that when someone grasps the meaning of a word, they are able to say "now I understand", "now I can do it", and "now I can go on" and use the word in future situations — just as a pupil who understands the rule of a mathematical function can go on correctly applying it beyond the example cases.[34]

33 Kripke, Wittgenstein on Rules and Private Language.
34 Wittgenstein, Philosophical Investigations, §114.

By Kripke's reckoning, a skeptical argument can be read into Wittgenstein's passages based on inductive considerations. Wittgenstein shows that indefinitely many courses of action can be found to be in accord to the rule. Any explanation of the rule will always be a finite set of instructions or based on a finite set of examples. But understanding the rule requires that we repeat our use the rule correctly in the future for indefinitely many cases. So any explanation of the rule will always underdetermine its future use. It seems as though understanding the rule instructs me on how to act in all future cases. But, as Kripke explicates, "when I concentrate on what is now in my mind, what instructions can be found there? [...] the infinitely many cases [of correct use] are not in my mind for my future self to consult."[35] This notion of an epistemic gap between our understanding of an explanation of the rule and our understanding of the rule itself suggests that understanding the rule must transcend understanding the explanation.

The next step is to argue that this is not so. The epistemic gap cannot be bridged by introspection. That is to say, understanding a rule cannot be understanding the publicly available instructions plus a private understanding achieved through intro-

35 It should be pointed out that Kripke (and Wittgenstein) are talking about rules for the correct use of a mathematical rule. This allows Kripke to say "infinitely", whereas holding that there are "infinitely" many uses of non-mathematical words might not be so clear. There are at least enough uses of any word to make its future use underdetermined by the instructions one has introspective access to.
Kripke, Wittgenstein on Rules and Private Language, 69.

spection. Cyrus Panjvani offers a reading of these passages that goes beyond this point about the under-determination of induction. He reads Wittgenstein as making a conceptual point against the "rule-realist" who accepts that understanding a rule must transcend understanding an explanation of it and holds that this epistemic gap can be bridged through introspection.[36] This is a conceptual point about what constitutes understanding the rule. It is not just that indefinitely many courses of action can be determined by the rule, which would be an epistemic point, but that any course of action can be determined by the rule if the normative criterion is only privately grasped. Of course, if any course of action can be interpreted in accord with the rule, then there is nothing that constitutes following the rule whatsoever.[37]

First, note that the rule-realist must deny that there is any difference between having an intuition and knowing what it is to follow it correctly.[38] Otherwise, even though one cannot be wrong about having an intuition one could still be wrong about having an intuition to follow the rule correctly — making the rule-realist position vulnerable to verificationist objections. Wittgenstein points out that an important aspect of following a rule is being able to repeat it.[39] So an intuition cannot be an intuition for following a rule unless the same intuition can be

36 Panjvani, "Rule-following, explanation transcendence, and private language", 312-313.
37 Ibid., 319-320.
38 Ibid., 322.
39 Wittgenstein, Philosophical Investigations, §199.

repeated. So there must be an external condition for identifying repeated instances of the intuition. But, “if a determination of correctness is separated from the having of an intuition, then anything can be interpreted to be correct (for this correctness is not given by the intuition)”.[40] So “no course of action could be determined by a rule, because any course of action can be made out to accord with the rule.”[41] Invoking a type/token distinction between the rule and different applications of it will only push the difficulty back a step. There must be a normative criterion to tell us whether a given performance is indeed a token application of this rule-type rather than that one.

Now we can apply this skeptical argument to the case of understanding the correct use of a word by first grasping its meaning. Grasping the meaning of the word must come from a finite set of examples (if not examples, then at least some other sort of publicly available explanation). But any explanation will underdetermine future uses of the word. So it must be held that grasping the meaning must transcend any explanation of it. It might be that grasping the meaning amounts to an intuition which will determine future use but this intuition must be repeatable. In order to identify that intuition across instances an external criterion of identity is needed but that separates the criterion

40 Panjvani, “Rule-following, explanation transcendence, and private language”, 323-4.
41 Wittgenstein, Philosophical Investigations, §201.

of correctness from having the intuition, so anything can count as a correct use of the word, and so there ends up being nothing which constitutes having a rule which governs that word's use.

In the Philosophical Investigations, then, I think two skeptical arguments against compositionality and atomism can be found.[42] The first argument is derived from the Excalibur and chessboard examples. It is a direct negative argument against the compositionality as a metaphysical principle. However the epistemic doubts it castes on that view might be ignored as insufficient to refute the metaphysical claims of atomism. As we will see in Chapter 3, Fodor claims he is as little worried by such arguments as physicists are by brains-in-vats. The second argument is derived from considerations of private language and rule-following. This is an indirect argument to show that atomistic views of language will, in principle, be unable to account for the use of linguistic expressions. In particular, atomism will be unable to account for their repeatability. As we will see in Chapter 4, Brandom takes the repeatability of expressions to be a central mode of conferring content.

42 It seems orthodox to note, at places like this, that Wittgenstein was not really presenting "arguments" in the sense that word is typically understood in philosophy.

2.3 Russell's Logical Atomism

In the Philosophy of Logical Atomism Russell orients his position against the monistic logic of "people who more or less follow Hegel."[43] He characterizes atomistic logic as holding, in opposition, that there are many separate things. Justifying this claim requires justifying analysis.[44] Finally, justifying analysis, in turn, requires answering two questions: "Are things that look like logically complex entities really complex?" and "Are they really entities?".[45] With previous discussions of Socrates' dream and Wittgenstein's language games in mind, we can anticipate some of the difficulties that will arise in providing answers to both these questions. He takes the claim that there are many things to be an empirical matter (as opposed to the monistic claim of there being only one thing, which uses an a priori proof).[46] Russell's answers to the two questions above, then, will be empirical. He sees atomism, in turn, as entailed by this empirically justified standpoint.

In answering the first question he starts by ruling out things which look like complex objects but are not logically complex objects. For example, tables, chairs, Picadilly and Romania are all ruled out because analysis of them will not reveal any logical constituents.[47] He instead focuses on facts as prima facie

43 Russell, The Philosophy of Logical Atomism, 36.
44 Ibid.
45 Ibid., 49.
46 Ibid., 47-48.
47 Ibid., 50.

logically complex entities. To answer whether they really are logically complex, however, he turns to the propositions that express them. His discussion here takes for granted a close connection between facts and propositions — a connection which will need to be further explored in his answer to the second question. He admits that "we really should start from the complexity of the fact" which cannot be "merely psychological" but he doubts whether "complexity, in that fundamental objective sense in which one starts from the complexity of a fact, is definable at all."[48]

The fear that the complexity of propositions might be "merely psychological" stems from his suggested answer to the first question. Propositions really are complex because they are composed of words, and you can understand novel propositions if you understand the words.[49] This "humble property" is what marks the proposition "as complex and distinguishes it from words whose meaning is simple."[50] This is where the atomisms of Russell and Fodor are most similar: both see the composition of "complex" propositions from "simple parts" as a fact straightforwardly read off the novelty of language.

The relation between propositions and facts give a connection between two different complexes, seemingly by stipulation.

48 Ibid., 57.
49 Ibid., 53.
50 Ibid., 54.

Symbols (words) combine to form propositions. Propositions correspond to facts, which are in turn complexes that are composed of the meanings of symbols. Russell is also committed to something like the Asymmetry Thesis discussed above, whereby the meanings of symbols and facts are known in two distinct ways. He tells us that "all analysis is only possible in regard to what is complex, and it always depends, in the last analysis, upon direct acquaintance with the objects which are meanings of certain symbols".[51] Russell states to an interlocutor that he does not think it "necessarily implied" that there are simple parts, simply because we asserted that propositions are complex.[52] Rather, evidence for the simple parts is based on the results of analysis and is empirical in nature.

Thus far, Ryle's reading of logical atomism seems consistent with Russell's presentation in the Philosophy of Logical Atomism. Insofar as propositions correspond to facts on a one-one basis, they seem to be single entities which function much like names. However, Russell explicitly denies that propositions are names for facts.[53] He says he was led to this by Wittgenstein, who pointed out to him that propositions have a true/false dichotomy which facts do not. Propositions can have two relations to a fact: true to the fact and false to the fact. Names, on the other

51 Ibid., 54.
52 Ibid., 64.
53 Ibid., 46.

hand, only have one relation. If a name does not mean some particular, then it is not a name (but rather a description, as would be shown by final analysis).

It is easy to see, however, how Ryle might feel justified in ignoring this distinction. Russell goes on to make another distinction between atomic propositions and molecular propositions, which correspond to atomic and molecular facts, respectively.[54] Atomic facts are only atomic in the sense that they can be combined with other atomic facts (expressed by combining atomic propositions using logical operators). Atomic facts are complexes made up of a relation and terms of that relation. Atomic propositions are not names for atomic facts because by names Russell means proper names which are words for particulars, and particulars are terms of relation in atomic facts. Most symbols we use everyday as names fail to be proper names in Russell's strict sense. They do not name particulars but rather are abbreviated descriptions.[55] A proper name can only only name something you are directly acquainted with. Particulars are indeed slippery things in a world where swords can be broken apart and still referred to. They are, for Russell, only present in the here and now, and perhaps in this way only the demonstrative "this" comes close to being a logical name.

54 Ibid., 59-60.
55 Ibid., 62.

I think these pieces give us just enough to take a look at Russell's answer to the second question: are these logically complex entities really entities? This could be asked about facts, propositions, or both. If it is asked only about propositions, Socrates' dilemma can be easily avoided as there would be no need to postulate an epistemological asymmetry between symbols and propositions. Speaking about propositions alone does not make any difference if we undertake to provide a definition (/explanation) of the symbols, or if we view the propositions as grasped through recognition. Explaining how we understand novel propositions requires a kind of asymmetry only because understanding the proposition requires grasping the fact it expresses.

Furthermore, if it is only propositions that really are entities then it could be charged that this is an empty formalism. To co-opt the jargon of the later Wittgenstein, it is not a frame with which to view the world at all, narrow or otherwise. The interesting questions, then, start with whether facts really are complex entities. Since facts are only known through propositions this question brings up epistemological difficulties which cannot be so easily avoided. I think that Russell takes the complexity of facts for granted, given that they must correspond to apparently complex propositions. I think we could fill out the suggestion of an argument this way: given that at least some propositions are true, and some false, there must be some extra-

propositional thing which makes them true or false. Lets call them facts. Since propositions are complexes their parts can be rearranged in a different way to form a different proposition. But if that's the case, it can be made from true to false, or false to true. Likewise, facts must also be complexes, such that if their components were were rearranged they would be a different facts.

However, although Russell holds that facts really are complexes formed from particulars, he denies that they are entities. That is to say, facts are a kind of object which make up the world but they do not have the "kind of reality" which belongs to simple parts from which they are composed.[56] Russell identifies three kinds of simple parts: particulars, qualities, and relations.[57] Relations hold between particulars and can be roughly seen as corresponding to verbs in propositions. They can come in various orders of hierarchy, depending on how many particulars they relate. It is these simple parts which are the real entities, because they can be known through direct acquaintance.

In this way we can see that Russell does not submit to Socrates' dilemma. He does hold that propositions and facts are both complexes, composed of words and simple parts. He does hold an epistemological asymmetry: facts can be given an account, the simple parts are known only through direct acquaintance. However,

56 Ibid., 142.
57 Ibid., 142.

since facts are not entities, propositions cannot be names for them. So he is not forced to hold that they can also be known through direct acquaintance. Facts, like the propositions which correspond to them, really are complex but they are not really entities. It is the simple parts which are entities. Furthermore, Russell holds that in justifying analysis it is enough that "in theory you can get down to the ultimate simples".[58] In theory we can arrive at them because analysis presupposes them and analysis can be done. For one thing, perhaps analysis explains how novel propositions can communicate newly understood facts. If that is so, then Russell's logical atomism is also immune to considerations along the lines of Wittgenstein's chessboard example. We do not need to justify that these particulars and those relations really are the ultimate simples outside of any language game. Indeed, Russell accepts that you cannot do so — giving a justification would be providing an explanation, and explanations go beyond direct acquaintance. Rather, we can give an explanation of the facts, which Russell repeatedly calls "logical fictions", created out of series of particulars.

Russell thus avoids Socrates' dilemma by holding that propositions are complexes built of names for particulars but not themselves names for particulars. He avoids the latter Wittgenstein's puzzle of composition by maintaining that it is enough to

58 Ibid., 142.

hold that the ultimate simples which compose reality exist but unnecessary to make claims about what they are. Only the "logical fictions" which arise from those simples can be given an account (consistent with the epistemological asymmetry between knowledge through acquaintance and knowledge through description). Russell avoids the difficulties of Socrates and the later Wittgenstein by trying to have a number of things both ways, for example: atomic simples are the only things which really exist but we cannot say precisely what they are; propositions/facts really are complex but are not really entities.

A number of aspects of the metaphysics of logical atomism are important to explaining how Russell might have these things both ways. Bernard Linsky offers a broad overview in "The Metaphysics of Logical Atomism" which is useful for seeing how the pieces fit together. Of particular interest for my purposes is the multiple relation theory as a way of explaining judgment without relying on the notion of propositions as unified entities that can be named. This view differs from the language of thought found in both Fodor and the Wittgenstein of the Tractatus. Whereas Fodor and the earlier Wittgenstein view judgments such as "X believes that a is b" as a relation between X and the proposition "a is b", Russell views it as a complex relation between X, a and b.[59] No fact can be included in another — all facts are atomic in

59 Linsky, Bernard. "The metaphysics of logical atomism", 373.

the sense of having just one relation and one or more individuals it relates.[60]

I think the multiple relation theory of judgment takes the wind out of the sails of a number of pragmatist gripes with atomism. But there are problems in working out the details. One such problem is put by Linsky this way: "How can Russell both deny the existence of propositions and then use variables 'p' and 'q' for them?"[61] Russell must allow for quantification over propositions because, for one, a complete description could not be given of the world without general facts. Particular facts might describe all the relations in the world, but for a complete description we would at least also need a general fact which states that those particular facts are all the facts. One way of answering this is to apply a nominalist interpretation of Russell's way of quantifying over propositions. The variables which range over propositions are not to be seen as ranging over objects, "but really are schematic letters holding the grammatical places of sentences" so that "propositions are just sentences, and propositional functions nothing more than predicates."[62] Russell gives an interpretation of the extensionality of propositional functions which allows functions to be retained in logic while eliminating talk about classes.[63] This is done by combining two

60 Ibid., 381.
61 Ibid., 373.
62 Ibid., 374.
63 Ibid., 378.

slogans about extensionality found in the second edition of Principia Mathematica: coextensive functions are identical and a function can only enter a proposition through its values. So if two propositions are "coextensive and also only contribute the truth values of the propositions that are their values to higher-order contexts, their intensional aspects will make no difference and as indistinguishable they will be considered identical."[64] In this way, if John believes that the cat is on the mat and Bill also believes the cat is on the mat, we can say that John and Bill believe the same thing without commiting ourselves to the view that there is some proposition which they both believe. Rather, we say they share the same relation toward the cat and the mat.

This pushes the problems of intepretting Russell onto his ontology. To go on with the same example, even though the proposition "the cat is on the mat" is made true or false by relating (or not relating) to a fact, the fact of the cat being on the mat is not something with which we can be acquainted. Indeed, not even the cat nor the mat are things with which we can be acquainted. Rather they are to be further analyzed into series of particulars with which we are acquainted. There must be some sort of final analysis, even though we might never come to it, or it would not be possible for us to analyze facts into

64 Ibid., 377.

constituents at all. Unlike with Fodor, representation does not go beyond present particulars. To answer Wittgenstein, there must be ultimate simple parts of reality even though we might not be able to specify them outside of a particular language game.

One complaint I have not discussed yet can still be launched at Russell. It is the insistence that you cannot really be directly acquainted with particulars in a meaningful way. Brandom develops this argument by starting with the Quinean consideration that deixis alone cannot confer semantic content — pointing while you make a noise is insufficient to determine the thing pointed at as the meaning of that noise. A logical name for a particular which we are directly acquainted with would not be repeatable. As such, it is not meaningful. What is needed is the ability to repeat that name; to refer to the same particular even though it is no longer present.

The multiple relation theory of judgment also seems to go against the intuition that propositions do form unities. Leonard Linsky describes aspects of Russell's progress toward logical atomism in the context of his struggle to account for the unity of the proposition. This problem is, in a sense, prior to the problems raised by Socrates' dilemma because that dilemma only arises when we consider the dual epistemic relations from propositions to facts and words to atomic meanings. Here difficulties already arise from attempts to account for what holds

a proposition together as unified whole in the first place — as a proposition rather than a meaningless aggregate of words.

Leonard Linsky notes that Russell declared himself completely defeated by the problem.[65] For Russell, a proposition is a unified complex rather than a mere aggregate because it is held together by a verb, which denotes a relation between the rest of the words in the proposition. The problem is that verbs can play dual roles: as a verb they can relate terms, or as a noun they can denote the relation itself. For example, the verb "to run" can be used to say "John runs to work" and "Running is how John comes to work". In the former, "runs" denotes a concept which relates John to his work. In the latter, running is treated as a term, not a relation. The fact that English can mark the nominalized usages with a gerund formulation of the verb and a placement in subject position is of no help. The fact that relations can sometimes themselves become terms means we cannot appeal to a dichotomy between relations and terms to explain why "running red blue" does not form a unified whole. "Russell's problem is that he has no explanation whatever for what it is for the verb to be used as a verb rather than as inert verbal noun and logical subject. How do relations relate?"[66]

Russell seems to have given up on accounting for the unity of a proposition. He does not return to it after presenting it in

65 Linsky, Leonard. "The unity of the proposition", 249.
66 Ibid., 245.

The Principle of Mathematics.[67] His ultimate denial of the unity of the proposition and adoption of the multiple-relation theory of judgment, however, provides at least a psychological solution.[68] Judgments, like verbs, can play a dual role — as relations and as objects. For example, Bill can believe that John believes that the cat is on the mat. Here, the second occurrence of "believes" is not itself a relation. We can understand a difference in the state of affairs whereby the cat is on the mat and the state of affairs whereby Bill merely believes that the cat is on the mat. In the latter case, the cat and the mat are only united in Bill's imagination. This is analagous to the iterated belief Bill holds about John. In that case John does not (necessarily) hold a relation of belief, it is only Bill who does. Although there may not be a unity to the proposition expressing a thought, there is a unity to the thought itself.

By contrast, Frege accounts for the unity of the proposition by accepting a strict dichotomy and holding that, after analysis, we can only use concepts but not talk about them. Frege's notion of concepts is different from Russell's — they are functions which take the objects of propositions as their domain and yield truth values. Using a metaphor of complete and incomplete, Frege strictly delineates their roles. Concepts are not just incomplete but essentially incomplete because in order to have a sense they

67 Ibid., 259.
68 Ibid., 258.

require objects. Objects, on the other hand, are complete in that they may have a sense individually. In a way consistent with the typical charge that Frege is dismissive about the features of natural language, this simply denies concepts the dual role which plagued Russell. This is not done without "embarassment" however: Frege notes the seemingly absurd conclusion that, for him, the concept horse is not a concept because saying that it is (or is not) requires treating it as an object in the position of logical subject. Sometimes, Frege says, our expressions miss the thought — we mention an object when what we mean is a concept. He asks that he be granted this with a "grain of salt".[69] This embarassing conclusion arises from a defect in our ordinary language, which allows us to speak of concepts as though they were objects by placing them in the grammatical position of subject. However, "the syntax of our language is not a logical syntax" and "[i]n a correct logical notation, Frege's embarassing sentence 'The concept horse is not a concept' cannot be written."[70]

According to Leonard Linsky, Wittgenstein followed Frege in this by "taking seriously" the contextuality principle, which means starting with the unity of the proposition rather than the individual constituents, as suggested in proposition 3.3 of the Tractatus: "only propositions have sense". Using Frege's metaphor of completeness, it is only the proposition which is complete.

69 Ibid., 248.
70 Ibid., 270.

The individual constituents are incomplete in that their logical form consists in their powers of combination with other constituents. The problem of the unity of the proposition, then, arises only from an inversion of priorities, from the misguided effort to begin with the constituents of the proposition conceived as independent building blocks and then to seek for the cement that will hold them together in the proposition.[71]

2.4 Two Principles in Frege

The crux of the debate between atomism and holism can be found in competing interpretations of Frege. There has been a long tradition of philosophers from both persuasions finding support in Frege. This has sometimes been taken as a central difficulty in Frege interpretation. The key is the relation between Frege's early work, notably the Foundations of Arithmetic, and his later work, notably "On Sense and Reference". The former contains a dictum which has been called the principle of contextuality: "never to ask for the meaning of a word in isolation, but only in the context of a proposition".[72] In his later work, after drawing the distinction between sense and reference, Frege is can be seen as suggesting the principle of compositionality: "[t]he possibility of our understanding propositions which we have never heard before rests evidently on this, that we construct the sense

71 Linsky, Leonard. "The unity of the proposition", 255.
72 Frege, The Foundations of Arithmetic, X.

of a proposition out of parts that correspond to the words."[73] The former principle suggests a sentential holism, the latter suggests atomism. The question arises as to how much Frege's philosophy of language changed with his introduction of the sense/reference distinction. Michael Dummett, for example, argued for a continuity in Frege, and struggled with attempting to bring these two principle together in a coherent fashion.

But some recent scholarship has argued that there is no real tension, because Frege was not really committed to anything like what we now call the principle of compositionality. Pelletier gives a broad survey of the the variety of ways either one or both of these principles have been attributed to Frege, and the different theses which this attribution has been meant to support.[74] He concludes that there is little to no textual evidence that Frege considered himself committed to anything like the contemporary principle of compositionality. Furthermore, although contextuality is more clearly found in Frege's work, it is not clear that Frege took it to commit him to holism. Theo Janssen argues that Frege always adhered to the principle of contextuality, but was not conflicted with a contrary commitment to compositionality.[75] Contrary to Dummett's germinal scholarship, these scholars find no conflict between the earlier and later Frege.

73 Frege, "Letter to Jourdain", 43.
74 Pelletier, "Did Frege believe Frege's principle?".
75 Janssen, "Frege, contextuality and compositionality".

This is, perhaps, more of a blow to those who want to read Frege as an atomist than those who want to read him as a holist. The case can still be made, I think, that Frege was a sentential holist throughout his career. Danielle Macbeth, for instance, argues that Frege offers a non-reductive view of analysis which is a viable alternative to Russell's atomism.[76] James Levine offers a similar appraisal.[77] For Frege, signs only express a sense prior to their use, and can express a thought only when those signs are put together in a sentence. That thought, in turn, is what can be analyzed into function and argument in a number of different, but equally valid ways.[78] This is crucially different from Russell's notion of analysis, which presumably only allows for one "final analysis". For Frege, on the other hand, while senses are composed into thoughts, the concepts designated by those thoughts are only understood relative to an analysis of a given thought.[79] In Frege's mature conception of a sentence, it displays three levels of structure: at the highest level the sentence expresses a thought and designates a truth value; at the lowest level it is a collection of primitive signs of the language arranged in a certain way; and in between are the object names and concept words that are revealed by analysis.[80]

76 Macbeth, "Logical analysis, reduction, and philosophical understanding".
77 Levine, "Analysis and decomposition in Frege and Russell".
78 Ibid., 479.
79 Ibid., 480.
80 Ibid., 483.

Macbeth's reading dissolves Dummett's criticism that Frege puts two incompatible demands on the sense of propositions. In Macbeth's view, there is a division of labour: "it is the senses of primitive expressions that are relevant to compositionality, and the senses of those concept words that are the result of analysis that contain modes of presentation of functions from objects to truth-values."[81] So the senses of words are composed into thoughts, those thoughts are then analyzed into function and argument, which reveals the concepts and objects referred to by the subsentential expression. It is then the sense of those concepts and objects (revealed by analysis) which denote the truth value. I think this position can be called sentential holism, because the referents of the simple parts are not given prior to a use in a sentence.

Although both Macbeth and Levine disagree with Dummett that Frege has incompatible commitments stemming from compositionality and contextuality, both follow him in trying to make sense of Frege's apparent holism. Dummett warns against attributing to Frege the simple slogan that it is only sentences which have meaning. This slogan is, he says, either truistic or nonsensical.[82] If we take it to mean that the words contribute nothing to the meaning of the sentence, it is nonsensical because it violates the obvious fact of linguistic productivity. On the other

81 Ibid., 480 fn.
82 Dummett, Frege: Philosophy of Language, 3.

hand, if the slogan simply means that we cannot say anything meaningful with a sequence of words that do not make up a sentence (and where there is no additional information supplied from context which completes a sentence), then it is merely truistic. Neither of these ways capture, according to Dummett, the real contributing insight Frege made with the context principle. Dummett says that if the consequences of Frege's approach be reduced to a slogan, it should be this: "that in the order of explanation the sense of a sentence is primary, but in the order of recognition the sense of a word is primary."[83] Macbeth's reading of Frege is compatible with this this revised slogan. In explaining (or discovering) the truth of a thought, it is the sense of the sentence which is primary because that is what we must analyze into concepts and objects. However, in grasping the thought which is expressed by that sentence, it is the senses of the words which is primary, since we compose the thought from them.

Under this interpretation, Frege never gave up the holistic view suggested by passages like this in the Foundations of Arithemetic:

> "If I give someone a stone with the words: Find the weight of this, I have given him precisely the object he is to investigate. But if I place a pile of playing cards in his hands with the words: Find the number of these, this does not

83 Ibid., 4.

> tell him whether I wish to know the number of cards, or of a complete packs of cards, or even say of points in the game of skat. To have given him the pile in his hands is not yet to have given him completely the object he is to investigate; I must add some further word — cards, or packs or points."[84]

Brandom reads this passage as support for his argument that objects (as the meaning of singular terms) cannot be given to us without a sortal concept.[85] A sortal concept is necessary to provide a criterion of identity for the object being sought. To assert that there is a unique analysis of "the number of this bunch of cards" is to assert that some way of sorting is already given with the presentation of the object (perhaps the number of individual cards) and others derived (perhaps points in a game of skat). This is holistic because it denies that words can denote themselves in and of themselves, which is a denial of the atomist thesis that atomic meanings are independent of each other and their use in complex meanings.

The final issue I want to address is why Frege introduced the notion of sense. Frege realized that the notion of reference was no longer adequate to fully capture what meaning is when he considered the relation of equality, what we might also call synonymy. If the meaning of a symbol is just the object it refers

84 Ibid., The Foundations of Arithmetic, 28-9.
85 Brandom, Making it Explicit, 438.

to, then we ignore the fact that “statements of the form a = b often contain very valuable extensions of our knowledge”.[86] If what we call the meaning of a statement is only its reference, we are leaving out the fact that the statement may contain actual knowledge. So, according to Frege, “It is natural ... to think of there being connected with a sign ... besides that to which the sign refers ... also what I should like to call the sense of the sign, wherein the mode of presentation is contained.”[87] The sense of a word is the knowledge which allows us to pick out its reference. It was important to Frege that this not be the same as an idea, which is subjective — it must be objective.[88]

It is the sense of words which are composed into thoughts, and so, if Frege was talking about compositionality at all, it was perhaps the compositionality of senses. Interestingly, however, it is the notion of sense which points to difficulties in atomistic semantics. Grasping the sense of a sign must give us knowledge required to pick out the sign’s reference. But since senses include some information which allows us to identify an object, they cannot really be independent of one another. Including information as an aspect of meaning makes meanings interdependent on one another. If you do not yet know that Hesperus is Phosphorous, for example, your grasp of the senses of those words

86 Frege, “On sense and reference”, 23.
87 Ibid., 54.
88 Ibid., 26.

is incomplete. Fodor, as we will see, wants to restrict meaning strictly to reference, to avoid any slide into holism. As such, however, he is required to explain the considerations of equality which lead Frege to introduce sense in addition to reference.

Summary

The goal of this chapter was to show some of the tensions between semantic atomism and holism in early analytic philosophy. In §2.3 I argued that Russell's logical atomism could avoid the epistemological difficulties found in Plato and Wittgenstein.

These epistemological difficulties were introduced earlier in the chapter. §2.1 introduced a dilemma which Plato posed for the atomist position. The atomist wished to explain a complex unity on the basis of primitive parts, but there is a danger of that complex unity itself being seen as much like a part. In §2.2 I explicated two lines of criticism from Wittgenstein's Investigations. The first line stems from Wittgenstein's considerations against the "Augustinian Picture" of language. It suggested that atomism captures only a limited part of what we call linguistic meaning. The second line stems from the so-called "private language argument", and suggests that the atomistic notion of meaning will be insufficient in principle to capture those aspects of language which it leaves out.

Such skeptical arguments cannot disprove strictly metaphysical semantics like Russell's atomism, but they might strain their plausibility. Russell felt that atomism was justified because it was presupposed by the fruitful method of analysis. However, in §2.4, I presented a reading of Frege which views him as offering an alternative, non-atomistic, notion of analysis.

3. Fodor's Conceptual Atomism

Fodor's theory of meaning takes seriously the common intuition that the basis of meaning is the fact that words stand for things. Fodor's theory holds in draws a much simpler, straightforward connection between words and the world. I want emphasize the intuitive aspects of Fodor's criticisms against Brandom, even thought they ultimately fail to refute Brandom's semantics.

§3.1 sets out to explicate how compositionality fits in with Jerry Fodor's representational and computational theory of mind. This is meant to serve in lieu of offering arguments directly supporting compositionality, which Fodor tends to take for granted. §3.2 will focus attention on the role of compositionality in Fodor's arguments against inferential role theories of concept possession. This section will set the ground for Chapter 4, where I discuss Brandom's inferential role semantics and consider Fodor's objections. It will also fill in some of the argumentative support for the representational and computational theory of mind approach which was missing from section one. §3.3 will explain how compositionality fits in with Fodor's notion of concept acquisition. This is meant to set the ground for Chapter 4 where I compare Fodor and Brandom on the issues of language acquisition and intentionality. In the last section I consider objections to Fodor's metaphysical approach to semantics from those of the pragmatist persuasion.

Note that I follow Fodor's preferred notation, whereby words writ in capital letters denote concepts, underlining refers to semantic values (whether they be referents, senses, or meaning generally), and single quotes around expressions which are mentioned rather than used. For example 'dog', for Fodor, refers to the concept DOG, which in turn refers to dogs. It should be clear from context wherever underlining or capital letters are used for emphasis instead.

3.1 Compositionality as a Consequence of RTM

No other contemporary philosopher has matched Jerry Fodor in both the number of appeals to and consequences derived from the principle of compositionality. Compositionality plays a central role many of his arguments against pragmatist (in general) and inferential role (in particular) approaches to semantics. He claims that compositionality demands that intentional content be purely referential, that the ontology of concept individuation be atomistic and that thoughts have a constituent structure.[1]

Unfortunately for those not inclined to grant the truth of compositionality, he offers little support for the principle itself. He is even loath to attempt any precise definitions of it: "[s]o not-negotiable [sic] is compositionality that I'm not even going to tell you what it is."[2] Although he takes the standard

1 Fodor, LOT 2, 20.
2 Fodor. "Language, thought and compositionality", 6.

arguments from productivity and systematicity to be decisive, he rarely ever rehearses them, and never in rigorous detail. I propose to follow Fodor's lead on this and not try to construct an argument for the truth of compositionality.

Instead, I think that Fodor's reasons for taking compositionality for granted can be better understood by considering his general approach to philosophy of language and the mind. Specifically, Fodor sees compositionality as both implying and implied by his commitment to the computational theory of mind (CTM) and representational theory of mind (RTM). Computationalism requires that meaningful wholes be made up of meaningful constituents which can serve as inputs and outputs in computations. Representationalism require that the world be divided up into atomic parts (corresponding to those constituents) which can be represented in a human mind. Drawing out the details of this fit will be the topic of this first section. The reader should be forewarned, however, that Fodor also often avoids direct arguments for RTM and CTM. Rather, they are most often presented as the best theories we have left once we see that pragmatist approaches fail — and one such argument Fodor frequently wields against pragmatism will be explicated in §3.3.

One possible approach to compositionality and atomism is to establish compositionality as true, then arrive at an ontology of concepts which follows from it. In Concepts, Fodor's approach is

the opposite: he presumes an atomistic ontology of concepts and compositionality comes out as a natural and elegant way of explaining mental phenomena. For example, concepts are mental particulars. Some of these, in turn, may be simple, but others are complex: “white horse” is a particular concept which is at the intersection of the concepts “white” and “horse”. Compositionality is a natural explanation of these processes.

Fodor sets out 5 key theses of RTM: T1 Psychological explanation is nomic and intentional;[3] T2 Mental representations are the primitive bearers of intentional content;[4] T3 Thinking is computation;[5] T4 Meaning is information;[6] and T5 Whatever it is that distinguishes coextensive concepts is “in the head”.[7] He believes that if one accepts these, then one is further lead to accept 5 non-negotiable conditions on a theory of concepts: C1 Concepts are mental particulars — they satisfy whatever ontological commitments are necessary to function as causes and effects;[8] C2 Concepts are categories and routinely employed as such;[9] C3 Compositionality — concepts are the constituents of thoughts as well as constituents of complex concepts and mental representations inherit their contents from the contents of their constitu-

3 Fodor, Concepts, 7.
4 Ibid.
5 Ibid., 9.
6 Ibid., 12.
7 Ibid., 15.
8 Ibid., 23.
9 Ibid., 24.

ents;[10] C4 A lot of concepts are learned;[11] and C5 Concepts are public.[12]

I propose to discuss these theses and conditions in turn, trying to explicate their connections with compositionality. I intend to discuss T1 – T3 in detail, and give only a brief mention of the significance of T4 and T5 at the end. Likewise, C1 – C3 will be discussed in detail but C4 and C5 briefly mentioned.

In some cases, the reasoning that leads one from accepting these theses to accepting the conditions is relatively clear. Accepting T1 leads to accepting C1 because it is supposed that only mental particulars will be able to serve in nomological explanations. Nomological explanations, in turn, are desirable for explaining intentionality scientifically.[13]

An obvious counter to this thesis and the condition which follows from it is a skepticism about any such nomological explanations of mental events in the first place. An alternative is to model the mind as being disposed to certain behaviours. For example, you have the concept of ‘and’ if you are disposed to use conjunction introduction and elimination appropriately (Fodor points out that starting with an explanation of logical concepts is a favoured strategy amongst pragmatists and that is true of Brandom in Between Saying and Doing).

10 Ibid., 25.
11 Ibid., 27.
12 Ibid., 28.
13 Ibid., 7.

Fodor argues against this in LOT 2, targeting Boghossian in particular who interprets being disposed to a behaviour as being disposed to following a rule.[14] First, you can only be said to be following a rule if that rule, R, is an intentional object of one of your mental states. Then, ‘X follows R’ can be read transparently or opaquely. The transparent reading is took weak for pragmatism's needs and also cannot accommodate compositionality, but the opaque reading is really just a representationalism in disguise. Under the transparent reading, we say you are disposed to follow R if you are disposed to follow any rule that's equivalent to R, but the goal is to define R not something equivalent to R (for example, the goal is to define the concept AND, not merely equivalence to AND). However, on the opaque reading of ‘X follows R’, X must, in some way, grasp R. This means that X is representing R or knowing-that R prior to knowing how to use R —— surely the opposite of what the pragmatist explanation set out to do. This is a recapitulation of the same argument against Boghossian that Fodor makes in “Brandom beleaguered”. He believes that Brandom's approach must be something similar.

Rejecting pragmatism as circular or ineffectual by reasoning along these lines leads one to accept T2, that mental representations are the primitive bearers of intentional content. Fodor characterizes this stance in LOT 2 as being “hyper-realist” about

14 Fodor, LOT 2, 38.

propositional attitudes and propositional attitude explanations: “tokens of cognitive mental states are tokens of relations between creatures and their mental representations”.[15] Fodor also explains this with the “belief box” metaphor — to believe P is to have P in one's “belief box”.[16] Clearly, such mental representations must be finite to fit in a “belief box”. Indeed T1 tells us that they are particulars. This yields an atomistic theory about concepts, specifically, about concept possession: one can have any given representation, mental particular, or belief, without the need to have any other. You could, according to this view, have the concept CAT even if you do not have the concept ANIMAL insofar as it is possible to think about cats without thinking of them as animals. Note that this could not be the case in inferential role theories since inferences most often involve more than one concept (invoking the law of identity being an exception). On the other hand, you could not think of BROWN COW without thinking of both BROWN and COW, but this is readily explained with a primitive/complex concept distinction and the principle of compositionality. There is still no need, on Fodor's account, to subsume the concept of BROWN COW under the concept of ANIMAL in order to have the former.

This condition contrasts with Russell's multiple-relation theory of judgment. It might be that representations lie at the

15 Ibid., 5.
16 Fodor, Concepts, 8.

bottom of Russell's semantics (if direct acquaintance of particulars is a kind of representation). However, under the multiple-relation theory of judgment, one does not judge a proposition. Fodor's theory suggests that a judgment is a representation of a representation. For now, take Fodor's protests that this does not lead to homuncular regress at face value.[17] Unlike Russell, however, Fodor seems clearly faced with Socrates' dilemma. Being "hyper-realist" about propositions means that propositions are there to be referred to even though they must be known a way different from simple representations like "this cat".

C2 is the condition that concepts be thought of as categories. This is just to say that concepts apply to things in the world. Applying concepts requires semantic evaluation.[18] The evaluation conditions of concepts are essential to the contents of the concepts.[19] So, to use Fodor's example, to judge that Greycat is a cat, is to say that the individual, Greycat, falls under the extension of CAT. Since the concept of CAT is atomistic then whether or not Greycat is a cat must not depend on its being (or not being) anything else. That his being a cat is incompatible with his being a tree but entails his being an animal is irrelevant. The meaning of "Greycat is a cat" is that Greycat falls un-

17 Fodor, LOT 2, 216.
18 Fodor, Concepts, 24.
19 Ibid., 25.

der the extension of CAT and that fact depends on nothing other than Greycat and the set of cats.

Now, Fodor does recognize that confirming the truth of "Greycat is a cat" is holistic: "Confirmation is an epistemic relation, not a semantic relation, and it is generally theory mediated, hence holistic."[20] Actually evaluating whether Greycat is truly a cat depends on being in the right place (the same place as Greycat) at the right time (in well-lit conditions perhaps). It might also depend on knowing that if he is a cat then he cannot be a tree. However, being unable to confirm or disconfirm the sentence does not show that one does not grasp its meaning.

This is another point which Fodor makes frequently: semantic evaluation must be prior to empirical confirmation. If one wants to explain how it is we could know whether Greycat is a cat, one might start with an analysis of the proposition "Greycat is a cat"; however, it still remains that the meaning of that sentence is ontologically dependent on the meaning of the words contained in it. In this way, Fodor's atomism extends not just to objects of the mental world, but also to objects of the physical world, even though the relation between minds and the physical world must be explained in holistic terms.

The idea that concepts are categories seems simple enough. Since it is the meaning of "x is a cat" which allows us to decide

20 Ibid.

whether or not Greycat is a cat, then the meaning of “x is a cat” must not be determined (even in part) on whether he is or is not a cat. This is a rule-realism about meaning determining the use of a word. Drawing from Wittgenstein's skeptical doubts about this rule-realism (see §2.2), however, we might ask where the representation of the categorical concept CAT comes from. It cannot come from epistemic considerations of cats, because then it could not serve as the rule for determining whether Greycat is a cat in a non-circular way. Although atomism about concepts seems to intuitively follow from their use in categorization, there is a definite tension between a concept being both categorical and atomistic. A representation of Greycat comes from Greycat. But if a representation of CAT cannot have anything to do with epistemic conditions of being a cat, then what is it a representation of? This forces Fodor to make the source of concepts something of a mystery. Fodor's answer to these difficulties will be elaborated in §3.3.

Fodor reasons from an assumption of RTM and an atomism about concepts to C3, the demand that concepts be compositional, by considering systematicity. Systematicity of thought is the ability of the mind to think thoughts which are of a similar logical or syntactic form. Demonstrations of systematicity often rely on supposedly clear examples rather than clear definitions: if a mind can think that John killed Mary, it could also think

that Mary killed John. It is the mind's capacity to think that is systematic — it is not necessary to actually think Mary killed John when presented with the former thought. The claim is that the mind must be systematic before acquiring language. Fodor asks how an unsystematic mind could learn a systematic language "given that the [language] is ipso facto able to express propositions that the [unsystematic mind] is unable to entertain"?[21]

The argument suggested to me is this: there is systematicity somewhere (by empirical fact); it is not in the world, but it is in the mind or language (known empirically); it cannot be that minds become systematic through learning a language, because they could not learn a language if they were not systematic; therefore minds are systematic. Lastly, the best explanation of systematicity is compositionality, so the constituents of thoughts (i.e. concepts), must be compositional. Furthermore, since systematicity itself inheres in intentionality — one cannot think or produce sentences systematically without thinking or talking about something — this is tied in with the issue of whether intentionality inheres first in thought or first in the mastery of a public language. Of course it bears mentioning that he does not take this as as apodictic: "If a serious alternative proposal should

21 Ibid., 26.
Yet surely being able to entertain a thought requires more than systematicity; it also requires that I have the concepts contained therein. There are many propositions in English, about rocket science for example, that my mind cannot entertain, although I learned English.

surface, I guess I'm prepared to reconsider what's negotiable. But the prospect hasn't been losing me sleep."[22]

One could certainly question whether minds or language really are systematic, despite the supposed obviousness. Kent Johnson points out that what often goes unmentioned is that systematicity always implicitly depends on linguistic natural kinds (substitute mental natural kinds, if you prefer to discuss the systematicity of thought) and argues that no matter how you divide up the kinds, they do not actually display systematicity.[23] One might think that the productivity of language better demonstrates the need for compositionality, and so serves Fodor's point better. I will focus on productivity in §3.2. That being said, a closer look at Johnson's arguments will reveal more difficulties in drawing the close relation between syntax and semantics which is seen in atomistic theories.

First, the substitutions which allegedly prove systematicity are always supposed to be of the same kind. "John killed Mary" can be systematically become "John killed Fred", but the fact that it cannot systematically become "John killed redly" is not supposed to disprove systematicity. So, at the very least, systematicity must be restricted to involve only substitutions amongst similar grammatical kinds.

22 Ibid., 27.
23 Johnson, "On the systematicity of language and thought".

Johnson identifies the following grammatical categories within which acceptable substitutions are meant to demonstrate systematicity: verbs, quantifiers, connectives, adjectives, nouns, and singular terms. For each of these categories he provides counter examples which show that the grammatical behaviour of words depends on syntactic and phonological properties of the words, as well as pragmatic implicature.[24] To use my own example, we might wish to disallow "John killed the desk" as a systematic substitution of "John killed Mary" insofar as the desk cannot be the victim of a killing. Yet both "Mary" and the "the desk" are singular terms and singular terms are the paragons of systematicity. Johnson provides a more decisive counter example from Attic Greek, which marks nouns for case, and a similar counter example in English where "He killed her" is acceptable but "Her killed he" is not.[25]

Johnson goes on to generalize this argument.[26] Depending on how we define our natural kinds, either language is trivially systematic or not systematic at all. An example he gives which demonstrates the former starts with dividing language into two kinds: words that Sally likes and words that she does not like. Obviously, if we substitute one Sally-kind for another Sally-kind

24 Ibid., 116-120.
25 This is a clearer example because one might argue that the meaning of 'to kill' could be metaphorically extended such that John could kill the desk, similar to how he might also love the desk, though hopefully not in precisely the same way that he loves Mary. Still, I think my counterexample is applicable to Fodor, for whom meaning is reference.
26 Ibid., 121-6.

in any given sentence we will yield a new sentence which still contains a Sally-kind. But this tells us nothing interesting about language. Now, we might divide our kinds up in ways that are theoretically useful, by making finer grained distinctions in the grammatical categories which reflect broad semantic distinctions (say verbs which are +CAUSE and verbs which are not). However, in such cases we will find that although these finer-grained kinds overlap, the linguistic evidence does not support calling the intersection of them a third kind. This is because, in principle, natural kinds do not distribute across Boolean combinations. According to Johnson, language is not systematic in any theoretically useful sense.

However, it might be believed that thought is able to escape the problems that language faces in terms of systematicity. Fodor could accept that language is not strictly systematic as long as thought is. Thoughts, as Johnson points out, are not overtly recognized in phonological or orthographic forms. This avoids some of the above mentioned problems.[27] Yet most suggestions that thought is systematic rely on an assumption that thought is similar to natural language in its systematicity and this includes instances where Fodor invokes systematicity. Furthermore, there is still the problem of barring substitutions

27 Ibid., 130.

which are unsystematic on the basis of some semantic feature in a non-arbitrary way.

Another way that someone might be led to accept that thinking is systematic is by first accepting T3, the thesis that thinking is computation. As Fodor puts it, this just to explain mental processes causally, in a manner which, unlike associationism, explains the ability of thought to preserve truth.[28] Computation is surely a better candidate than either thought or language for being obviously systematic. In Chapter 4 of LOT 2, Fodor distinguishes his version of computationalism about the mind from others. Computations, Fodor argues, must be defined over local properties. The argument goes like this: the local properties of X are independent of anything except X;[29] The part/whole relation is paradigmatic — being a part of a cow does not depend on being or not being a part of anything else; since computation processes are defined over the syntactic structure of mental representations; and since syntactic structure is constituent structure; and since constituent structure is a part/whole relation then computational process are local.[30] This suggests a way that Fodor can avoid Johnson's arguments against systematicity. If systematicity holds only amongst the parts of a given representation then it seems much clearer and more plausible than saying that

28 Fodor, Concepts, 10.
29 Fodor, LOT 2, 107-8.
30 Ibid., 108.

all thought is systematic. To use a simple example, given a complex representation like “the cat is on the mat”, computational processes could yield the systematic “the mat is on the cat” but not the unsystematic “the mat on the cat is” because “the cat is” is not a constituent in the first place, let alone one which is substitutable for “the mat”. Since, as Fodor points out, some mental processes (like fixation of empirical belief) are not local, and hence not computational in Fodor's sense, the concession would be that some (perhaps much) of thought is not systematic after all. It is just that none of the non-systematic thought is relevant to intentionality or semantics.

I will finish this section with only brief remarks about the remaining theses and conditions, which tend to be less directly related to compositionality. T4 says that meaning is information and is meant to capture Fodor's belief that the semantics of thought must be purely referential. Fodor believes that only a purely referential ontology of conceptual content will satisfy compositionality. Roughly, it is supposed that if there is anything that is intermediary between concepts and referents the content of the concepts will not be atomic. Relatedly, T5 is a denial of Fregean senses acting as intermediaries between concepts and referents. He addresses Frege's problem of substitution of co-referential terms at length in Chapter 3 of LOT 2. T4 ad T5 together recall the difficulties discussed with concepts as

categories above. If the meaning of concepts is determined strictly by direct reference then how do they function as categories and what are they representations of?

C4 seems to be an utter contradiction of Fodor's infamous Language of Thought doctrine that no primitive concepts can be learned. I will discuss this argument in detail in Section 3 of this chapter. The confusion rests on the fact that here Fodor is using "learn" to refer to concept acquisition broadly (i.e. he is saying that concepts must be acquired after a causal interaction with their reference) but his argument in LOT focuses on "learning" as an inductive method. He accepts that concepts must be acquired in some way from experience with the relevant objects; he denies that they can be acquired inductively. C5 is the condition that we must be able to share concepts. This is usually appealed to on intuitive grounds: how else would we be able to communicate successfully? This point is often stressed against holistic theories. If concepts are determined by inferential role, for example, and inferential roles vary between individuals, then how can people be talking about the same things? However, it is also difficult to see how Fodor is able to justify his holding of C5. If concepts are mental particulars whose content is determined solely by reference then how can they be shared?

There are a number of places in Fodor's theory of concepts with the potential for circularity. That representation can be

seen as a kind of ability is not one of them. He characterizes his version of RTM as being Cartesian about concepts: having the concept X is being able to think of Xs as such. Thinking about dogs as dogs is enough to have the concept DOG; it is not necessary to be able to think that dogs are not cats. In LOT 2, he considers the objection that this too must be a kind of ability, that is, the ability to think about Xs as such.[31] He replies that you can consider it an ability if you like, however, unlike other abilities which pragmatists appeal to (like the ability to sort Xs) it gets the order of analysis right because it supervenes on representation. Furthermore, we know that it gets the order of analysis right because the ability to think about Xs as such, unlike to ability to sort Xs, composes.

A real potential circularity, however, arises in this note of warning against Brandom: “It is a bad idea to confuse semantics with the epistemology of interpretation/translation; and it is a bad idea to confuse semantics with the psychology of learning; and it is a bad idea to confuse semantics with the psychology of language production/comprehension.”[32] Semantics is a metaphysical pursuit. One of the many cardinal sins that pragmatists commit is to use epistemological arguments to make claims about semantics. But Fodor himself takes compositionality to impose serious constraints on semantic theory even though arguments

31 Ibid., 47-8.
32 Fodor and Lepore, “Brandom Beleaguered”, 679.

for compositionality are inextricably tied up with the psychology of language production and comprehension, if not also epistemological concerns.

Returning to Fodor's definition of having a concept: it is hard to understand what thinking about dogs (or anything) as such amounts to if one is not thinking of dogs that they have some kind of property. But if you are thinking that, then you need two concepts, one that denotes dogs as such and one that denotes the property you think they have. But that is not atomistic, so it cannot be what Fodor means by "thinking of dogs as such".

3.2 Compositionality and Concept Possession

Fodor has used compositionality as an objection to pragmatism, inferential role semantics and even Brandom's approach specifically. All these objections are variations on the same theme: epistemic conditions on concept possession do not compose, so having a concept cannot be construed as knowing how to do something. A favoured approach to making this argument is using adjective-noun phrase examples such as (my example) 'revolutionary song', or better yet, the complex concepts denoted by such phrases, such as REVOLUTIONARY SONG. I will restrict my explication to the complex concepts rather than to the natural language

phrases which denote them since Fodor thinks that thought must be compositional even if natural language is not.[33]

First of all, it should be noted that these adjective-noun phrase examples might seem odd when put to Fodor's ends. After all, as we saw in §1.2, they are tokens of a type of example against compositionality [CE-MOD]. In fact, Fodor's arguments are also arguments against taking these cases to be counter-examples. The argument that epistemic conditions should not be a part of semantic content also rules out any aspect of context constituting any part of content, since adjusting semantic contents on the basis of context would be an epistemological affair.

Fodor thinks that epistemic conditions fall under two kinds: following a rule or sorting things. He prefers to use the ability to sort things as an example to show why such conditions do not sort, so I will start with that. He presumes an ability like sorting will ultimately have to be appealed to in order to explain the content of empirical concepts. If having REVOLUTIONARY SONG means to know how to sort revolutionary songs from other things, then that know-how cannot be the sum of knowing how to sort SONGS from NON-SONGS and REVOLUTIONARY things from NON-REVOLUTIONARY things. Clearly you cannot be expected to sort SONGS from NON-SONGS under any conditions — for example, if the song is being quietly played amongst many other loud noises — in order to

33 Fodor, "Language, thought and compositionality", 7.

be said to have that concept. So some conditions must be better than others. There must be normalcy conditions, whereby under such conditions you need to be able to tell songs from non-songs in order to possess the concept SONG. However, normalcy conditions do not compose. For example, the normalcy conditions for spotting things with the property of being REVOLUTIONARY might include being in the midst of a revolution, but the normalcy conditions for spotting SONGS might include being in an otherwise quiet environment.

Fodor generalizes the argument: any epistemic constraints on concept possession will violate compositionality since any epistemic constraint requires associated normalcy conditions;[34] either language or thought must be compositional; concept possession must not violate compositionality. So pragmatic theories of concepts fail, inferential role theories, fail, and Brandom, if indeed concept possession for him is epistemic, fails too. This is all unsurprising since knowing something necessarily requires believing it and for Fodor believing something requires representing it.

It is not obvious to me, however, that the argument does generalize. To start with, consider a simplistic pragmatism where having a concept means knowing when one's linguistic peers will agree that you used the word correctly. It seems to me, then, you

34 Fodor, LOT 2, 46.

could say that having REVOLUTIONARY SONGS is the sum of the two: knowing when one's linguistic peers will agree that the two words were put together correctly. At least, there is no obvious reason why the details of this theory could not be worked out so that it does compose in this way. Let us say that there are rules (whether the speaker knows them or not) for the correct use of both words, then why would not the rules for the correct use of their combination be the union of those of each word?

The response from Fodor could be that actually this theory is representational. My simplistic pragmatism described above is ambiguous between two kinds of representing which determine whether or not I have the concepts. Either I mean that I have the concept if I can represent some fact about my linguistic peers or I mean I have the concept if I can represent some rule about its use. This is a species of Fodor's point that representation is at the bottom of any reasoning about the world because being able to think thoughts which are susceptible to truth values must be prior to the ability to plan a course of action.[35] The pragmatist approach, according to Fodor, is to get everything backwards and if they are getting things right (if they are able to explain compositionality) then it must be because they are starting to get it the right way around.

35 Ibid., 13.

Fodor says loudly, “ADOPTING A DISPOSITIONAL ACCOUNT OF RULE-FOLLOWING WON’T SAVE AN INFERENAL ROLE SEMANTICS FROM THE CHARGE THAT IT IMPLIES A CIRCULAR ACCOUNT OF CONCEPT POSSESSION.”[36] I wonder, however, at the risk of being yelled at, if it could save an inferential role semantics from the charge that it violates compositionality. Remember Fodor’s criticisms of Boghossian mentioned in the previous section: Where ‘X follows rule R’ is read opaquely, it is circular (or really representationalism). But if it is read transparently, where X is following R if X is following anything equivalent to R, then the pragmatists are not really providing a definition of R (merely equivalence to R). Let us take the transparent horn of this dilemma. I guess that inferential role theorists are not supposed to take this reading because they are meant to be providing definitions of concepts via their inferential roles. Fodor says that the transparent reading is too weak for their purposes.[37] Just for now, let us assume that defining equivalence classes of concepts serves their purpose fine: as long as one follows any set of rules which are equivalent to the set of rules which defines a concept, one has that concept.[38] Well then, if one is disposed to follow any set of rules which is equivalent to the set of rules defining REVOLU-

36 Ibid., 39.
37 Ibid., 38.
38 Think of it this way: Fodor has his own concept PROTON and James Watson has another, but they are equivalent and hence Fodor and Watson could, in principle, talk about the same thing. Fodor argues that there’s no viable notion of equivalence (or similarity) that could sustain this.

TIONARY then one has that concept (or an equivalent one), and mutatis mutandis for SONGS. Then, if one is disposed to follow any set of rules which is equivalent to the union of those two concepts (or any two equivalent concepts), one also possesses REVOLUTIONARY SONG. Perhaps this could be the start of a compositional, pragmatist account of complex concepts.

Note that epistemic skepticism does not apply: it does not matter how unwieldy these sets of equivalent rules are since the pragmatist is precisely denying that one must represent the rules in order to count as "knowing how" to follow them. It is like riding a bike, I suppose. There are consequences that may seem strange to some. For example, you might have some particular concept, yet be the last person in your linguistic community to know that you have it: you have a concept because you are disposed to a particular behaviour, but you have no word for it.

However, a challenge which does apply is a demand for a principled way of identifying equivalent concepts. Fodor charges that this would require a tenable analytic/synthetic distinction, which would, in turn, require answering to Quine's arguments against such a distinction. Fodor claims no one has been able to do this. Brandom thinks that one could follow Quine and not draw any such distinction, or that one could follow Sellars and draw the distinction from amongst inferences which are counter-factually robust and those that are not. I think the latter is a vi-

able option, but I am not sure the former is. At any rate, it still requires a long story about how such inferences sort meanings into equivalence classes, something I attempt in Chapter 4.

I do not think that it has been demonstrably shown that pragmatist theories of concept possession simply cannot explain the arrangement of concepts into novel thoughts. We might still say, however, that it is much more difficult for them to do so. I will stop trying to talk of pragmatist theories generally and focus on Fodor's criticisms of Brandom's inferentialism specifically. Fodor admits that Cartesian representationalists such as himself do not exactly have the details of an account of compositionality worked out (formal semanticists are still at it).[39] However, it is easy to see how well compositionality fits with their approach. Cartesians take representations to be the primary bearers of intentional content. Singular terms are the cornerstone: 'Fodor' means FODOR, which means Fodor, and so on for every referring concept I can name. Since these primitive representations are metaphysically prior to complex representations, Cartesians can accept the demands of compositionality right from the start, namely that the semantics of complex representations are determined entirely by the semantics of their constituents and the way those constituents are put together. Indeed, Cartesians can even assume such things.

39 Fodor and Lepore, "Brandom's Burdens", 472.

We should not expect it to be so straightforward for inferentialists. Inferentialists prioritize judging, and since the smallest unit of judgment is a complex whole they prioritize propositions, thoughts, sentences and such. Brandom describes his approach as "top-down": semantic explanations start with a sentence, and work towards explaining the sub-sentential components.[40] Any account of compositionality from this explanatory starting point is going to require some acrobatics but it is not impossible in principle. Fodor and Lepore agree that if Brandom's attempt to define singular terms in Articulating Reasons were successful it would amount to an account of compositionality.[41] They deny that it is successful.[42] Brandom agrees that productivity demands compositionality. He denies that compositionality in turn demands atomism — at most it demands molecularism.[43] Fodor and Lepore respond by agreeing that molecularism would solve many problems of holism but argue that such molecularism requires an analytic/synthetic distinction in order to make sense of identical content.[44] Alternatively, there needs to be at least some way of identify similar content, if one is just to give up on the notion of identical content.

40 Brandom, "Inferentialism and some of its challenges".
41 The details of the account are first worked out in Making it Explicit, then reported in Articulating Reasons. In the article referred to, Fodor and Lepore are reviewing Articulating Reasons.
42 Fodor and Lepore, "Brandom's Burdens", 472.
43 Brandom, "Inferentialism and Some of its Challenges", 671.
44 Fodor and Lepore, "Brandom Beleaguered", 689.

Fodor has claimed that productivity demands compositionality; that compositionality demands representation be prior to inferring; and that representation demand semantic atomism. I think that these claims are hyperbolic. He himself tends to back off from them, at times, and agrees that pragmatic/molecular/inferential accounts could explain productivity without semantic atomism, but not without a further explanation of sameness or similarity of content. However, I still think it can be claimed that Fodor's explanation of the productivity of language/thought will be intuitively more elegant than one which is not committed to atomism, which prioritizes inferring over representing and works top-down from sentences to words. I would like to end this section with a quote which ends Fodor and Lepore's critical attack on Articulating Reasons:

> "You can't rely on the notion of inference to abstract the world from the theory of meaning. [... If] the language that you talk/think in happens to be productive, then you can't subtract the notion of reference from the notion of truth, since productivity demands compositionality, and compositionality implies the priority of subsentential semantics to sentential semantics, and referring is the typical job that subsentential expressions perform. And, finally, you can't subtract the notion of an object from the notion of reference because,

> in the paradigm cases, objects are what subsentential expressions refer to."[45]

I think I have shown that the features of language which arise from its composability, such as productivity and systematicity, do not quite make inferential role semantics impossible. At most, we can say that a semantic theory which makes the ontology of subsentential, referring expressions primitive has an easier job of explaining compositionality than one that does not. This may still yet count as a good reason to have a referential semantics. Notice that Fodor says that compositionality implies the priority of subsentential semantics to sentential semantics. "Priority" can be understood in a number of ways — ontological priority, explanatory priority, epistemic priority, etc. In the next chapter, I will argue that Brandom endorses the explanatory and epistemic priority of sentential semantics over subsentential semantics, remains neutral with regards to ontological priority and still offers an explanation of how sentential semantics are constructed from subsentential semantics. If "priority" here is to mean what one starts with, then, as Fodor himself says, it depends on what you are explaining and to whom.[46]

45 Fodor and Lepore, "Brandom's Burdens", 480.
46 Fodor and Lepore, "Brandom Beleaguered", 677.

3.3 Compositionality and Concept Acquisition

The fourth condition on concepts is that lots of concepts must turn out to be learned, yet Fodor is infamous for arguing that concepts cannot be learned. In LOT 2 he introduces a clarifying distinction: concept acquisition is any mind/world interaction which can alter a conceptual repertoire; concept learning is a kind of concept acquisition which happens through inductive inference.[47] Under this terminology, C4 should perhaps be reworded to say that a lot of concepts should turn out to be acquired from mind/world interactions but not turn out to be learned through inductive inference.

Explaining how concepts are acquired but not learned is what Fodor terms the doorknob/DOORKNOB problem. Any theory of concepts needs to explain why we acquire empirical concepts from certain kinds of interactions but not others — why we acquire DOORKNOB from interactions with doorknobs and not from interactions with cows. It seems as though any such process must involve learning the fact that DOORKNOB applies to only and all doorknobs. However, learning that fact in turn seems to require inductive reasoning on a hypothesis about DOORKNOB, and Fodor argues that inductive inference cannot account for concept acquisition.

In The Language of Thought, he argues that primitive concepts cannot be learned. The argument is that any concept you

47 Fodor, LOT 2, 131-2.

learn inductively requires a tokening of that very same concept in order to form an hypothesis about it in the first place. For example, you might figure that confirming the hypothesis "GREEN applies only and to all green things" would result in acquiring green but clearly that hypothesis already requires thinking about green.[48] In LOT 2 he generalizes this to include complex concepts as well. In order for you to formulate a hypothesis about complex concepts, you first need to have all the concepts which are its constituents.[49] So reasoning about hypotheses cannot be the way you acquire even complex concepts. On the other hand if you do have all the constituent concepts (and the rule(s) that govern their composition), then compositionality already explains how it is you could think about, and thus have, that complex concept. The moral of the story: "[hypothesis formation] is most naturally construed as a theory about what goes on in acquiring beliefs. It applies to learning concepts only via the (very tendentious) assumption that concept learning is itself a species of belief acquisition. But beliefs are constructs out of concepts, not the other way around".[50]

So we have two problems relevant for any theory of concepts: the DOORKNOB/doorknob problem, and the circularity of concept learning. I think that one or the other, not necessarily

48 Ibid., 137.
49 Ibid., 139.
50 Ibid..

both, must be answered by a theory of concepts.[51] If a theory of concepts does insist that acquisition of empirical concepts is a kind of inductive inference, then it does not have to answer to the DOORKNOB/doorknob problem because surely inductive inference is a very natural and direct way of explaining how it is that experience precisely with doorknobs results in acquiring the concept DOORKNOB. On the other hand, such a theory will have trouble with the circularity problem. For a theory which chooses to deny that empirical concepts are acquired from inductive inference, the circularity problem disappears but the problem of providing a plausible explanation for how it is that certain experiences are required for acquiring certain concepts becomes more difficult.

For example, it would seem that inferential role theorists would have little problem explaining how DOORKNOB is acquired from doorknobs — presumably interaction with doorknobs is a primary source of the inferential roles of DOORKNOB. However, inferential role theories will have far more trouble with the circularity problem. After all, a concept only has an inferential role insofar as it is the constituent of beliefs. This seems to make the circularity of concept learning both glaringly explicit and inescapable for them. If a concept is constituted by its inferential role it seems having at least some beliefs about a concept is ne-

51 Fodor presents these two problems as forming a kind of dilemma in Chapter 6 of Concepts.

cessary in order to have the concept. However, in order to have beliefs about a concept it must be somehow accessible as an object of intentionality which suggests one must already have that concept. So how could coming to believe certain things result in the acquiring of a previously unknown concept? Furthermore, unlike representationalists, inferential role theorists must confront this circularity. Since atomistic representationalists think that having a concept does not depend on having any beliefs about it they are able to sidestep the circularity by positing non-inferential processes of concept acquisition.

This tactic, however, makes the DOORKNOB/doorknob problem more pronounced for the atomists/representationalists. Even those who have no problem proclaiming some concepts innate might think it ridiculous to include concepts like DOORKNOB among them. Furthermore, it is unlikely than any posited mechanism for acquiring such concepts will draw the connection between doorknobs and DOORKNOB as neatly as inductive inference does.

Fodor's own answers to these problems are more suggestive than substantive. At bottom, he believes that the connection between empirical concepts and the properties they refer to must be causal rather than inductive. In Concepts, he argues that whether something is a doorknob partially depends on whether we take it to be a doorknob, because DOORKNOB has no conceptual analysis and doorknobs have no hidden essence which make them

doorknobs.[52] Specifically, "being a doorknob is having that property that minds like ours come to resonate to in consequence of relevant experience with stereotypic doorknobs."[53] This provides a non-inductive DOORKNOB/doorknob connection. It is meant to alleviate nervousness about radical nativism since it is not the empirical concept which is innate but rather some kind of a propensity of our minds. Lastly, it is not circular since "[t]he only theoretically interesting connection between being a doorknob and satisfying the doorknob stereotype is that, contingently, things that do either often do both."[54] Still, for many, saying that we are prone to take something as a doorknob prior to any experience with doorknobs might be just as bewildering as saying that we already have the DOORKNOB concept regardless of experience.

He revisits the issue in LOT 2. There, in what seems like a change of tactic, he suggests that a nativist might as well just say "concepts are, as it were, there from the beginning. We have the concepts we do because we have the neurology we do [...]. We have [DOORKNOB] for the same reason we have ten fingers".[55] This, of course, just raises the DOORKNOB/doorknob problem all over again: what role does experience play in such an account? In considering this reiterated objection, Fodor suggests that experi-

52 Fodor, Concepts, 136.
53 Ibid., 137.
54 Ibid., 138.
55 Fodor, LOT 2, 146.

ence might play a triggering role. This is consistent with what he said in Concepts and indeed he further suggests that learning a concept's stereotype might be a stage in concept acquisition, even though it is not sufficient for concept acquisition.[56]

Fodor leaves out many details which would be needed for an account of concept acquisition which is compatible with his ontology of concepts, but the rough outline he provides is sufficient to make a comparison with what sort of theory an inferential role theorist would need. The suggestion is that the acquisition of empirical concepts must ultimately rely on a causal rather than epistemological connection with the world. This leads to a commitment to some kind of innateness, and just what it is which must be innate is one of the details which needs to be worked out. Fodor delights in claiming that he is not at all scared of a theory which posits all empirical concepts as innate, but it is not clear what his answer to the DOORKNOB/doorknob problem would be in such a case. At any rate, he suggests a possible alternative in positing that things in the world trigger concepts. In such a case, it might not be that the concept itself is innate, but rather the potential for that concept. Finally, epistemological considerations cannot be completely circumvented. In Fodor's account, learning a concept's stereotype is a necessary

56 Ibid., 150.

step in triggering concept acquisition, though it is not sufficient for acquiring that concept.[57]

Paul Griffiths argues that proponents of innateness tend to conflate a number of empirically distinct concepts under the banner of innateness.[58] I think Fodor can be accused of this. On the one hand, he seems to view an innate concept simply as one that is not-learned by induction. However, he does not always stick to this conception. When faced with the hard-to-swallow conclusion that all concepts are innate, he switches to suggestions that innateness means some propensity to acquire concepts non-inductively. However, in that sense, innateness means more than just "not learned", it means "present at birth", "shared by all members of the species" and perhaps "unchanging through development". In short, invoking innateness in this way tends to just make things more mysterious then they need be.

Fiona Cowie argues that innateness arguments from what are called the "logical problem" tend to discount evidence which can help learners form concepts. Fodor is not arguing from the "logical problem", which essentially says that the evidence language learners are exposed to is insufficient for them to form the correct hypotheses. However, I think some of the same considerations are implicit in his arguments. Using Cowie's example, although no

57 This is because if concepts are prototypes, then they aren't compositional. Sere chapter 5 of Concepts for detailed discussion.
58 Griffiths, What is Innateness?

one ever teaches us what a curry is or is not, we tend to find out what it is on what seems like very little evidence.[59] Nevertheless, there is more negative evidence than first appears — we do not call hamburgers curries.[60] Also, a single piece of evidence is not enough to verify or falsify our hypothesis of what a curry is. Finally, arguments for innateness tend to assume that we need to explain how we all arrive at the same theory (same concept), when actually we do not.[61] We might well disagree about whether a mulligatawny falls under the curry concept, just as we might disagree about whether some particular thing is a doorknob or not.

At any rate, some of the problems Fodor finds here are surely the result of the peculiarity of his position. He is faced with explaining two difficulties in that concept possession is all-or-nothing and must be infallible. According to this view, then, acquiring a concept like GREEN, if it was learned, would mean confirming a belief like all and only green things are green. There is, however, no reason why an inferentialist would follow along. It is true that in order to have content you would need some beliefs about the concept GREEN. You might begin with the belief that if some patch of space-time is green then it cannot also be red. That is enough to begin to use that concept, and thus enough to begin to have (shades of) the concept. Even a

59 Cowie, “The logical problem of concept acquisition”, 34.
60 Ibid., 35.
61 Ibid., 37.

false belief about a concept's applicability is enough to allow the person to use that concept. Although they may use it mistakenly at first, the could, in time, be corrected. There is no reason to suppose, anyways, that my (very fuzzy) concept of "proton" is equivalent to a physicist's concept of "proton".

3.4 Skepticism about Concepts

Fodor's considerations about concept possession and concept acquisition yield a view of compositionality compatible with Funct(μ), which identifies the meaning relation with the syntactic relation. Concepts are (or at least can be) the meanings of grammatical terms. Complex concepts are made from combinations of these concepts. This is a commitment to a meaning operation which combines the meanings of simple grammatical terms. Although Fodor does not spell out his notion of compositionality, in order for it to have to consequences he claims it has, it must also be compatible with [C], which says that the meanings of component parts determine the meanings of complex expressions, and hence that the component meanings must be ontologically prior to the expressions. The result is that constituent grammatical terms have independent meanings which determine the meanings of complex grammatical terms. This combination, however, must be equivalent to the way that those grammatical terms are syntactically combined. The result is semantic atomism but one which is different

from Russell's logical atomism. Unlike with Russell, these combined, complex meanings must themselves be representations. In the previous chapter I suggested ways in which Russell could avoid the skeptical arguments of the later Wittgenstein. The question is whether Fodor can similarly avoid this skepticism.

At any rate, when Russell delivered the lectures which form The Philosophy of Logical Atomism he was unaware of Wittgenstein's new stance towards atomism. Fodor, on the other hand, is all too aware that he works amongst and against a tradition of philosophers inspired by the Philosophical Investigations, yet he chooses to ignore the problems Wittgenstein raises. At times he does explain why he thinks Wittgenstein should be ignored. Here is a comment he makes on Dummett, which could well be addressed to anyone who thinks Wittgenstein put an end to the search for meanings:

> "If, however, skepticism really is the skeleton in Dummett's closet, the worry seems to me to be doubly misplaced: first because the questions with which theories of meaning are primarily concerned are metaphysical rather than epistemic. This is as it should be; understanding what a thing is, is invariably prior to understanding how we know what it is. And, secondly, because there is no obvious reason why behaviourally grounded inferences to attributions of concepts, meanings, mental processes, communicative intentions, and the like

> should be freer from normal inductive risk than, as it might be, perceptually grounded attributions of tails to cats."[62]

I suppose the claim is that the post-Philosophical Investigations philosophers are mistakenly putting epistemology before metaphysics and they could not do so without inherently relying on false metaphysical theses. This may well simply baffle such philosophers: the very problem is that any claim of understanding of what a thing is is susceptible to epistemic doubt. Claiming that understanding that some thing must be prior to another does not do away with the fact that evaluating claims about what that thing is requires doing epistemology.

Fodor's second charge in this quote is perhaps more concerning. Pragmatists may not be any better off by directly addressing a skepticism about individuated concepts. In fact, whereas Fodor is not at all concerned with Wittgenstein's skepticism, Brandom absolutely is. In Between Saying and Doing, Brandom's project is to reconcile what he takes to be the methods of analytic philosophy with his holistic/pragmatic framework. In that work, he takes Wittgenstein's skepticism to be the most important objection to such a project that to which he must answer. For Fodor, pushing such objections to one side results in a clearer view of the relevant subject matter: the metaphysics of concepts.

62 Fodor, Concepts, 5.

Dummett argues that defining one concept as being composed of others often results in circularity.[63] The example already discussed in the previous chapter is that “child” might be defined as “boy or girl” or “boy” and “girl” might be defined as “male child” and “female child” respectively. However, Fodor happily accepts the impossibility of definitions. He claims, in fact, that it is the pragmatists who continue to seek definitions for concepts as definitions-in-use.[64] For Fodor, “boy”, “girl”, and “child” are all primitive concepts. The distinction between primitive and complex is not primarily a semantic distinction. Rather, it is a syntactic distinction. A discursive representation is semantically compositional, he says, if and only if “its semantic interpretation is exhaustively determined by its syntax together with the semantic interpretations of its lexical primitives”.[65] “Boy”, “girl” and “child” all function as primitive constituents in the English language because they don't exhibit constituent structure (like BROWN COW does). Fodor seems to take that at face value — presumably the corresponding mentalese concepts do not exhibit constituent structure either.[66] At any rate, he sets the problem of whether all monomorphemic words in a natural language correspond to primitive concepts to one side.[67]

63 Dummett, Thought and Reality, 7-8.
64 Fodor, LOT 2, chapter 2.
65 Ibid., 172.
66 Ibid., 58-61.
67 Ibid., 65.

As we will see in the next chapter, Fodor's application of his own all-or-nothing view of concepts to pragmatist views result in a misunderstanding of the positions he critisizes. Fodor believes that if concepts are determined by inferential role then inferential roles must give necessary and sufficient conditions for concepts along the lines of "x is a dog iff y". He rightly points out that no such definitions exist. But the whole point of the pragmatist view of concept possession is that there are shades of concept possession. Understanding the role a concept plays in at least some instances is enough to partially grasp the concept, even if it is not enough to provide an "if and only if" definition.

I promised I would try to provide an answer on Fodor's behalf to the three classes of counter-examples to compositionality: [CE-AMB] from ambiguity, [CE-MOD] from modifier-head constructions and [CE-SYN] from synonymy. These difficulties were: the principle of compositionality offers no way of determining the correct meaning in cases where there is an ambiguity of scope; the principle is unable to explain a shift in the meaning of the same modifier used across contexts; the principle raises Frege's puzzle of substitution for theories of direct reference.

I will start with Frege's puzzles as the central examples of [CE-SYN]. In LOT 2, Fodor sets out to "nibble away" at the puzzling cases rather than solve them tout court. The idea is to

identify the truly puzzling cases for his theory, and explain why they are indeed truly puzzling.

Cases involving complex concepts are not puzzling. Since complex concepts differ in their constituent structure they also differ in their possession conditions.[68] Because a person could possess the concept “the morning star” but not the concept “the evening star” it is not puzzling that we cannot substitute one for the other in propositional attitude descriptions.

This leaves primitive concepts as the truly puzzling cases. Standard examples include “Cicero = Tully” and “Paderewski = Paderewski”, where one might wonder if Paderewski the pianist is also Paderewski the politician.

Fodor's suggestion for handling these cases is to do away with the idea that proper names are pure referring expressions.[69] Since he defends a purely referential semantics, this entails that natural languages do not have a semantics.[70] Rather, it is only mentalese which has a semantics (i.e. refers), and Paderewski has two names in mentalese.

This suggestion, however, seems to require that Fodor make concessions to those who posit senses (or descriptions, or inferential roles), as playing an intermediary role between proper names and their references. He does this by making a place for

68 Ibid., 64.
69 Ibid., 72.
70 Ibid., 73.

such information in explaining mental processes. The content of a concept, he insists, is still independent, still compositional, and still purely referential.[71] If semantics is our domain of inquiry, we are asking about the content of a person's thought that Paderewski is Paderewski. The meaning of that thought is of the form x = x. If you wonder whether Paderewski is Padereski, then perhaps you do not know the meaning. If our domain of inquiry is psychology, we are asking about how a person represents things. In this sort of inquiry, additional information, in the form of inferential roles or descriptions, becomes relevant. One can wonder whether Paderewski is Paderewski because one represents the world as possibly having two Paderewskis.

Fodor does not address the problem of a modifier seemingly shifting meaning across contexts, but it seems his suggestion for Frege cases can also be applied here. The red of “red apple” and “red grapefruit” refers to the very same colour in each case. The semantics of the representations do not differ. However, representing an apple as red and representing a grapefruit as red requires also having, in addition to the reference, inferential or descriptive information about how each is red. A thought about “red ideology”, on the other hand, would seem to have no semantics. We cannot inquire about the meaning of such a thought.

71 Ibid., 88.

However, we could still inquire into how a thought is used in mental processes to represent the world.

Finally, it can also be suggested that restricting the domain of a quantifier like “every” in “everyone loves someone” is a psychological process. As such it is permissible for Fodor to say that this can depend on inferences without the meaning of the quantifier also depending on inferences. In that case, there cannot be just one meaning of any given quantifier across uses. In fact, it would seem that there would be potentially infinite number of meanings since we can quantify over a potentially infinite number of domains: all the beer in the world, all the beer in Canada, all the beer in my fridge and so on. Since a quantifier by itself, however, doesn't refer to anything in the world, there doesn't seem to be anything especially puzzling about this problem for Fodor's theory.

Summary

The goal of this chapter was to present just some of Fodor's arguments that the representational/computational theory of mind is the only viable approach to semantics. The aim was to show how compositionality is central to this theory.

§3.1 presented compositionality as a principle which fits naturally within the representational/computational theory of mind. Although Fodor does not typically argue directly for com-

positionality, it should be seen as an essential aspect of his atomism. Furthermore, his interpretation of compositionality is clearly opposed to Brandom's theory. For Fodor, compositionality explains how individual representations are combined into meaningful wholes which can serve as the objects of propositional attitudes according to computational processes of the mind. Part of what all that means is that it presupposes atomism of concepts.

§3.2 presented Fodor's argument that non-atomistic semantics violate compositionality. The idea is that any non-atomistic semantics requires individuating concepts according to epistemic conditions, and that epistemic conditions do not compose. I argued that this argument fails to general against any non-atomistic notion of semantic content. As long as a concept is not strictly identified with epistemic conditions, I think that it can be partially defined by such epistemic conditions.

§3.3 presented Fodor's argument against seeing concept acquisition as a kind of inductive learning. Once again, although I think that Fodor's position here follows naturally from his representational/computational approach, I do not think he is able to rule out other approaches. In particular, I think that inferential role theories are able to provide a plausible account of concept acquisition as a kind of theory testing. This is primarily because, according to inferential role semantics, concepts

might be only partially grasped. For Fodor, concept possession is all-or-nothing.

Finally, in §3.4 I offered some answers on Fodor's behalf to the kinds of skepticism which might be inspired by Wittgenstein. I also offered an explanation of how Fodor might deal with the supposed counter-examples to compositionality mentioned in §1.1.

4. Brandom's Conceptual Holism

Robert Brandom's semantic theory instantiates everything Fodor says will not work. In fact, Brandom's theory differs from Fodor's in such a myriad of ways it is difficult to choose one as the fulcrum for comparison. The ostensible goal of this chapter is to show that Brandom is also able to account for productivity. Doing this will also fill in some of the details as to why Fodor's other criticisms against inferential role semantics do not apply. Perhaps the greatest challenge is in making the connection between words and world intuitive; Brandom admittedly must tell a much longer and more complicated story about this than Fodor.

Brandom's theory does not stand strictly opposite Fodor's, however. Often, as I hope to show, his criticism of representational semantics are best read as a denial of certain theses without necessarily the endorsement of the contrary position. Fodor, for example, clearly demarcates semantics and pragmatics and prioritizes one over the other. Semantics is metaphysical, pragmatics is epistemological. Semantics is prior because in order to go approach epistemological concerns (such as whether Greycat is a cat) we need, first of all, to grasp the metaphysical concepts involved (be able to think of cats as cats). Brandom denies Fodor on both counts but he does not go on to endorse the priority of pragmatics over semantics. Instead, both are equally necessary aspects of the conceptual.

For my purposes, the central difference between Brandom (qua holist) and Fodor (qua atomist) is that Brandom gives sentences primary importance in semantic evaluations. In terms of an ontology of meanings, propositions form the central semantic objects to which Brandom must be committed. As with Fodor, propositions function both as linguistic objects (things which can be asserted) and mental objects (things which can be believed, denied or otherwise thought). Concepts, however, which may correspond to subsentential expressions and are the components of such thoughts have only a derivative existence in this theory. This, I believe, can be viewed much like a reversal of the Asymmetry Thesis of Socrates' dilemma. It is the meaning of complex linguistic expressions which are better grasped than the meaning of the components. The difficulty faced is that this notion of grasping a meaning seems to go against the common-grain of intuitions about meanings. Arguments from productivity to the principle of compositionality and then on to semantic atomism exploit these intuitions.

Starting with semantics, then, means starting in the middle of Brandom's usual story about linking words with the world. In Making it Explicit, semantics is presented as being derived from pragmatics in the sense that the the content of concepts are conferred by their norm-governed use in a linguistic community. The story presented in that work begins with justifying such theses

as "semantics must answer to pragmatics"[1] along with working out the details of a normative, pragmatic framework from which semantic content can be derived. Semantic content is worked out backwards, from the content of sentences down to the content of subsentential expressions and ending where Fodor begins — with the connection between words and the world.

That order of presentation, however, is not necessary in order to combat the claim that productivity/compositionality requires atomism. Making it Explicit sets out to argue, among other things, that semantics cannot be done in lieu of pragmatics, inference need not be defined in terms of a suitable notion of truth, and reference cannot replace expression. His claim that "semantics must answer to pragmatics" is the denial that semantics can be given an ontological status independent of pragmatics, it is not the further, mistaken claim that pragmatics could be given such an independent status.

In Between Saying and Doing, the order of presentation between pragmatic and semantic aspects matters less. There he presents an extended argument that:

> "every autonomous discursive practice [...] in order to count as deploying any vocabulary, must include performances that have the pragmatic significance of assertions, which on the syntactic side are utterances of declarative

1 Brandom, Making it Explicit, 83.

> sentences and whose semantic content consists of propositions."[2]

Utterances which function pragmatically as assertions, syntactically form declarative sentences and semantically have the content of propositions are all aspects of the "iron triangle of discursiveness." As the name suggests, this means that no one of these three are completely derivable from the others. The domains of syntax, pragmatics, and semantics are, we might say, metaphysical equals in explaining linguistic practice.

The project of Between Saying and Doing offers a conception of "meaning-use analysis" which does not require atomism or representationalism. As such, it goes beyond demonstrating Fodor's claims of compositionality wrong. It is also meant to show that, contrary to what Russell supposed, the possibility of analysis need not rest on a presumption of atomism and, contrary to quietist interpretations of Wittgenstein, analysis is possible even when considerations of use are included. These results, Brandom believes, do not depend on the normative pragmatics or inferential semantics of Making it Explicit.[3] Nevertheless, the discussion here will draw on both these works in order to demonstrate not just the possibility that a holistic semantics can answer compositionalist challenges, but that inferential role semantics in particular can.

2 Brandom, Between Saying and Doing, 117.
3 Ibid., 234.

In addition to the argument from compositionality is an entire constellation of related criticisms directed at Brandom which I will attempt to address. Following from holism is the alleged incommensurability of concepts and the impossibility of communication. Brandom himself calls this "the biggest challenge" that has always faced holistic accounts of semantics.[4] Following from the notion of content as inferential role is the demand for an analytic-synthetic distinction amongst inferences. Brandom considers this difficulty easily dispatched with, but Fodor puts relatively strong weight on it. One concern here is whether Brandom's answer to this problem leads him to one of the implausible holisms discussed in §1.3.

Finally, there are numerous challenges launched at Brandom's pragmatism. These start with considerations of compositionality and go further. Many of them can be grouped under the umbrella question "how does it all get started?" From considerations of compositionality: how do semantic evaluations of novel sentences get started? From considerations of knowing-how and knowing-that: how can knowledge-how to make assertions get started without first having a knowledge-that those assertions mean what they do? Additionally, one could question how language (or thought) become manifested in creatures like us if one is not taken to be "original" and the other derived. Following Fodor's

4 Brandom, "Inferentialism and some of its challenges", 663.

quip: how could one plan to be painted blue without first being able to representing oneself as blue? It is these last problems that I am most interested in, both because of their relation to the problem from compositionality and because of the interested differences I think they reveal between atomistic and holistic approaches to language.

§4.1 will present some considerations for taking the proposition, rather than possible constituents, as the primary bearers of meaning. Then §4.2 will show how inferential role determines propositional content. From that, constituent content can be derived using substitution and anaphora. §4.3 will look at the intentional relation between language users and the world which results from this semantic theory. The idea is to suggest how non-conceptualized pragmatic interactions with the world can be a basis for conceptualized interactions with the world. Finally §4.4 will put these pieces together to explain how language users can, under this theory, semantically evaluate novel locutions and address some of the related concerns with holism, such as incommensurability arguments.

4.1 The Priority of Propositions

Some features of semantic content are the same here as in Fodor. Corresponding to sentences are propositions, which are thoughts in that they can be the objects of propositional atti-

tudes. Corresponding to subsentential components are concepts. The fundamental disagreement is in demarcating the content of a concept. However, in order to explain the content of concepts in Brandom's theory we must first look at the content of propositions. For Brandom the content of concepts supervenes on the content of propositions. This provides a starting point with which to analyze the "iron triangle of discursiveness". On the semantic side, the evaluation of sentences into propositions is primary and allows for a key to evaluating the content of subsentential components and thus the content of concepts.

Giving propositions pride of place in terms of semantic content seems to be clearly favouring pragmatics over semantics. The atomist approach draws attention to the content of concepts first and foremost. Propositions are here given pride of place because they are the smallest semantic units which can serve as premises and conclusions in reasoning. It is their use in such reasoning which confers semantic content. The immediate challenge is justifying this move in the face of an apparent circularity. According to Fodor inferences require grasping meanings so they could not themselves confer meaning. I find four suggested lines of argumentation in Making it Explicit for the priority of the propositional. Three of them are positive arguments for the importance of the content of propositions. The fourth is a negative

argument against approaches which derive propositional content from conceptual content.

The first argument is inspired by Kant: grasping meanings is something we do; we bring things into a unity (i.e. subject concepts to rules); such an action is an aspect of judging; and judging requires a proposition.[5] Fodor's condition C2 on concepts, discussed above in §3.1, says that concepts are categories. As such, applying them is something we do in making judgments. However, Fodor does not follow this Kantian reasoning all the way down the line. Brandom's suggestion is that this activity on the part of the mind must also extend to the grasping of concepts themselves. For Fodor, concepts-as-representations must be first passively acquired before they can be used in judgments.

The second argument offers two considerations from Frege: an expression has content insofar as it makes a contribution to the truth value of a proposition in which it occurs, following Frege of "On Sense and Reference";[6] it is the sentence to which assertoic force attaches, following Frege of Foundations of Arithmetic.[7] In the first consideration it is the truth value of the proposition which we are ultimately interested in. A possible response is, of course, that that truth value can only be arrived at from the conceptual building blocks from which it is composed.

5 Brandom, Making it Explicit, 79-80.
6 Ibid., 81.
7 Ibid., 82.

A response to the second consideration could be to assert that assertoic force is itself derivative on literal meaning which is determined by conceptual content.

The third is an original argument which draws on a view of intentionality from Daniel Dennett:[8] intentional contents are attributed to non-linguistic creatures to explain their behaviour; the point of attributing such contents is to determine their practical significance; determining their practical significance typically relies on attributions of inferential reasoning; whatever can serve as a premise (for example, “the mouse believes that the cat is on the mat”) or a conclusion (for example, “therefore, the mouse avoids the mat”) must have a propositional content.[9] Here one might take objection with the idea that attribution of intentional content is meant to serve the purpose of explaining behaviour.

In all these arguments one might imagine the atomist nodding her head approvingly yet still insisting on the priority of conceptual contents. The fourth argument offers a multi-faceted attack on that alternative. Brandom argues that the representational model, which builds propositional content from conceptual content, faces two difficulties. First, “[i]t is not clear how to derive a notion of propositional contentfulness from the designa-

8 Brandom discusses Dennett’s view of intentionality in Making it Explicit, 55-62.
9 Ibid., 83.

tional representational model". This difficulty can manifest itself in the fact that the content of a proposition must be built from representations yet also be itself a representation, suggesting two different types of representation. In the Frege we have the idea that the references of sub-sentential components are things in the world, but the reference of propositions is a truth value. It is at least prima facie odd to consider that objects in the world compose into truth values. Even though the propositional is meant to be explained entirely in terms of its prior components, it often emerges as something different. Starting with propositions, on the other hand, we can identify a relevant property, and, using substitution, define sub-sentential expressions by invariance according to that very same property.

Second, "[c]onstruing content in representational terms still requires supplementation to explain proper use and consequences".[10] As we saw in the previous chapter, Fodor needs to make concessions allowing for epistemic information to be somehow attached to concepts even while he denies that information a constitutive role. His adherence to that denial becomes strained when faced with the doorknob/DOORKNOB problem, for example. Allowing pragmatic content to play a constitutive role in concepts, on the other hand, will yield a theory of concepts which does not need such supplementation in order to explain use.

10 Ibid., 84.

Brandom identifies two characteristic mistakes the representational model of content makes in addressing these difficulties: the first is assimilating sentences to complex names; the second is assimilating judging to predicating.[11]

This first mistake is precisely what Ryle accused the logical atomists of doing,[12] although Brandom offers a different explanation of why it is a mistake. Here, a mistake is revealed due to a resulting inability to distinguish between referring to a complex object and stating a fact about its components. For example, if we refer to a pictures as “the circle between two squares”. Are we simply representing that complex object, or saying what it is composed of? In the first instance, we are simply representing an object. In the second instance, we are representing an object as being presented in a particular way. Representationalism presumes that the latter form of representation is dependent on the first, but it is unable to draw a meaningful distinction between the two.

The second mistake results from representing something particular as something general (e.g. that apple is food). This makes an untenable distinction between two kinds of representing — representing as referring and representing as predicating. The result is a failure to distinguish and account for two important

11 Ibid., 84-5.
12 See §2.1 for a discussion of Ryle’s critique. In §2.3 I argue that it does not really apply to Russell.

aspects which reveal predication's use of concepts: regular differential response dispositions (for example, regularly treating tokens of “apple” as “food” by eating them) and normative constraints which differentiate between correct and incorrect predications. That this is a mistake is revealed by a resulting inability to employ such representations to account for use, which will be discussed in the next section below.

So it is because semantics, pragmatics (as well as syntax) are all equally necessary conditions for linguistic practice that semantics must answer to pragmatics. The overall strategy is to show that representational atomism is inadequate to account for the pragmatics of language, even when that semantic atomism is supplemented with work in linguistics and psychology. The alternative is to make the meaning of propositions central and show that this is able to account for all three aspects of the iron triangle of discursiveness. So far my goal has been only a clear explication of this strategy in order to prevent misunderstanding. A vindication of the position comes in the form of a demonstration that this approach avoids the difficulties and problems that representational atomism faces.

Brandom's position on original intentionality can also be difficult to parse. On the one hand, he endorses Davidson's claim that “to be a believer one must be an interpreter of the speech of others” but “neither language nor thinking can be fully ex-

plained in terms of the other, and neither has conceptual priority."[13] On the other hand, he describes the relational view of concepts and intentionality this way: "[c]oncept use is treated as an essentially linguistic affair".[14] That seems to suggest that thought can be fully explained in terms of language. In fact, this is meant to be a middle path. To see this we must note the distinctions which mark Brandom's approach to intentionality. For one, it is a view of intentionality in terms of linguistic practice as opposed to rational agency.[15] Viewing intentionality as inhering in rational agency, instead, suggests that rational agency can be separated from linguistic practice. That would grant thinking conceptual priority over speaking. Secondly, it views the contentfulness of intentional states as related to the contentfulness of speech acts, rather than as simply resembling to contentfulness of speech acts.[16] This is to say that it is essential, not merely accidental, to language and thought that, as Frege puts it, "even a thought grasped by a human being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new".[17]

The result is an approach to language which is, first of all, methodologically opposite to Fodor's. Brandom describes his

13 Brandom, Articulating Reasons, 5-6.
14 Ibid., 6.
see also Making it Explicit, 151.
15 Brandom, Making it Explicit, 149.
16 Ibid., 150-1.
17 Frege, "Compound thoughts", 1.

semantic inferentialism this way when he sets out three of its central theses: a) it is sententialist, or top-down, meaning that semantic evaluations are performed first on sentences and derivatively on subsentential expressions; b) it is expressive, or sense-based, meaning that semantic evaluations first assign senses to words, and only on that basis assign references and truth values; and c) it is rationalist in its choice of conceptual paradigm, meaning that semantic evaluation is modeled on logical vocabulary rather than non-inferential representation.[18] Fodor's approach, on the other hand, is: a) compositional; b) reference based; and c) empiricist. For Fodor, this approach is underwritten by his atomism about meanings. Brandom's own metaphysics is the result of his approach. The consideration that thoughts are essentially linguistic and that meaningful utterances essentially have pragmatic force result in ruling out particular views of the ontology of meanings. The holistic view of concepts comes out as a consequence of inferentialism, not something which is independently desirable.[19]

There plenty of room for error in comparing Brandom and Fodor's approaches. At the very least, noticing some of the ways a comparison can go astray should result in making it less likely that those particular mistakes will be made. Fodor opposes himself to Brandom in metaphysical terms about what the content of

18 Brandom, "Inferentialism and some of its challenges", 654.
19 Brandom, Making it Explicit, 89-1.

concepts must be. Brandom opposes himself to Fodor in methodological terms about where we must start showing how the content of concepts can be determined. Fodor, in turn, reserves the right not to understand talk about where one "starts".[20] For my part, I will try my best to focus on the metaphysical end result of Brandom's theory without worrying too much about whether he "starts" in the right pace.

That end result is a theory which is decidedly holistic in the sense that concepts are not independent of one another. As Brandom puts it, "one cannot have any concepts unless one has many concepts."[21] It is also a theory where the meaning relation is decidedly pragmatic — meanings are derived from their use. In the preface to Between Saying and Doing, Brandom says that "[o]ne of the reasons this kind of semantic relation has been overlooked is an overemphasis on semantic compositionality."[22]

4.2 From Inferences to Meanings

The distinction between what conceptual content is and how it is grasped is blurred in inferential semantics. Semantic content is, to some extent, what one believes it to be. This is a central point of contention. For Fodor, the result of this notion of content is an inadequate semantics. One way of showing that

20 Fodor and Lepore, "Brandom beleaguered", 689.
21 Brandom, Articulating Reasons, 15.
22 Brandom, Between Saying and Doing, xi.

the semantics is inadequate is by showing that it is non-compositional and cannot account for the productivity of language. But this is also a feature which Brandom uses against approaches like Fodor's. Only by including features of how meaning is grasped in conceptual content, Brandom alleges, can we have an adequate account of pragmatics. The challenge then, is to show that his semantics is adequate, at least that it can account for productivity.

Brandom identifies two necessary criteria for grasping the conceptual content of a proposition: knowing what entitles one to that proposition and knowing what that propositions in turn commits one to.[23] Commitment and entitlement constitute that proposition's conceptual content. If a sentence does not correspond to a proposition which entails or is entailed by other commitments, it is does not have a semantic content at all. "Propositionally contentful commitments are picked out in the first instance as those that can both serve as and stand in need of reasons."[24] Linguistic expressions are thus identified as meaningful by their ability to play a role in inference. The linguistic unit which typically plays such a role is a sentence. A sentence does not have a meaning unless it expresses a proposition which can be questioned and/or can be offered as a reason.

23 Brandom, Making it Explicit, 120-1.
24 Ibid., 275.

For a proposition to serve as or stand in need of reasons is for it to be a possible a premise or conclusion in an inference. This is true even for propositions whose semantic content is most apparently non-inferential, such as the proposition expressed by the sentence “this is red.” Brandom motivates this by claiming that the crucial difference between a genuine non-inferential report of “this is red” and a mere response exhibited by that of say, a red-reporting machine, is one of understanding.[25] The red-reporting machine is not employing a concept because it does not understand what “this is red” entails or what it could follow from. Now, this reasoning would surely not convince someone, such as Fodor, who insists that semantic content be sharply separated from pragmatic use. As Fodor puts it, the meaning of “red” is RED, not any set of conditions of when saying “red” is appropriate. A possible response from such an opponent might be that the question is not whether the red-reporting machine is employing concepts (it is not) but whether there are concepts expressed by its report (there are). The semantic content is expressed whether or not it is understood.

However, the notion of semantic content put forward by such an opponent would not include the normative dimension which defines correct and incorrect use of the concepts. The anti-pragmatist would likely say that such is as it should be: correct and

25 Ibid., 89.

incorrect use should not be considered part of the content of a concept. Now, if a notion of semantic content which does not include a normative dimension is in turn extended to explain its use, the only readily apparent evidence that can be pointed to is a reliable differential response to appropriate stimuli. The anti-pragmatist, then, must accept that, for all we know from the empirical evidence, the red-reporting machine is applying concepts. If the normative dimension of correct and incorrect use is not built into the content of concepts, then any criteria for correct or incorrect application must be external to the concept itself. In that case, it is possible to have a concept without any criteria of application. So it might even be further asked of the anti-pragmatist whether a machine which reported red completely randomly, not according to any rule whatsoever, could be seen as applying concepts. This argument takes its form from the considerations of private language and rule-following in Wittgenstein's Investigations (see §2.2). The anti-pragmatists explanation of the normative dimension of concept application must be based on empirical evidence, and any such empirical evidence will be insufficient to determine a rule.

So, even the content of non-inferential propositions is to be determined, at least partially, by inferential role. Inferential role, in turn, partially specifies correct and incorrect use because the salient properties inferential role encodes are com-

mitment and entitlement. Part of the meaning of any meaningful proposition P is what other propositions it commits or entitles one to, as well as what other propositions might commit or entitle one to it.

The next step is to explain what inferences are relevant and how they confer content. The challenge is to do this without a circular appeal to semantic notions. This requires conceiving of inference in a way very different from how it is often treated. For example, grasping the content of a proposition requires being able to make a distinction between good inferences, which are relevant to the content of the proposition, and bad inferences, which are not. The notion of a good inference is typically understood as one which preserves truth. But truth is a property of propositions and typically considered a semantic concept. As such, it cannot be appealed to in determining semantic content. Here, inference must be understood as the prior notion — truth (or the commitment involved in taken something as true) is what is preserved by good inference.[26]

Inference, then, is to be understood as something which can and is done without concepts and hence pre-linguistically. The ability to infer begins with non-conceptual classification as implicit in an organism's behaviour.[27] Brandom points to Hegel as one who has developed a naturalized, pragmatic version of impli-

26 Ibid., 113.
27 Ibid., 87.

cit classification by identifying animal desire as the source of classification.[28] When a mouse avoids the cat it is doing what is needed to show that it implicitly classifies the cat a certain way. What we are doing in explaining the mouses behaviour with a train of inferences is making that classification explicit through the use of concepts. Inference is a linguistic performance, but its roots are in implicit, non-linguistic performances: eating cheese but not rocks, avoiding mats that have cats on them, etc.[29]

A correct inference, then, is not defined as one which preserves truth. Rather, the act of making a correct inference is doing something appropriately. Since inference involves employing conceptual content, the relevant kind of doings include saying or believing something appropriately. Just as the roots of inference are to be found in non-linguistic doings, so are the roots of good inference: eating cheese but not rocks, perhaps. Concepts allow us to explain to each other why the inference from “that's cheese” to “that's edible” is a good one. Concepts allow us to verbalize inferences as the reasons behind what we do.

What is here called a “good inference” may not necessarily preserve truth. That is to say that good inferences on true premises might lead to falsity just so long as making that infer-

28 Ibid., 86.

29 Fodor's quip about planning to paint oneself blue is misleading. Planning to paint oneself blue might well require first representing oneself as blue (or, perhaps, having the concept blue) but simply rolling in blue paint requires neither a prior plan nor a prior representation.

ence is appropriate. Unless what counts as “appropriate” is brought in line with truth we cannot say that truth is what is preserved by good inferences when good inferences are understood as appropriate doings. It is difficult to see how “appropriate doings” could be defined in such a way that they are always truth preserving without first presupposing a grasp of truth and hence the danger of a circular appeal to semantic notions. The non-circular solution to this difficulty is to accept that appropriate doings at one time might turn out to have been inappropriate in the end. It might, for example, be appropriate, although ultimately a bad idea, to eat the rocks if they are somehow disguised as cheese. The central semantic notion becomes commitment to truth rather than truth itself. Good inference is what preserves commitment to truth. For example, if you are committed to the truth of A and the truth of A entails B then you are committed to the truth of B.

Content-conferring inferences need not be construed in such a way as to make them formally valid. A good inference may be as simple as inferring “this is edible” from “this is cheese”. Formal validity would require the insertion of a conditional to the effect that “all cheese is edible”. Brandom frequently cites three examples of inferences which are constitutive of the concepts which are involved in them: the inference from “Calgary is south of Edmonton” to “Edmonton is north of Calgary”; the infer-

ence from “today is Tuesday” to “tomorrow is Wednesday” and the inference from “lightning is seen now” to “thunder will be heard soon”. These are all material inferences, “[t]he kind of inference whose correctnesses essentially involve the conceptual contents of its premises and conclusions”.[30] A number of Fodor's criticisms against Brandom's semantics take hold here, as will be discussed in §4.4.

Brandom labels his form of inferentialism strong inferentialism and locates it between weak inferentialism and hyperinferentialism.[31] In contrast to weak inferentialism, Brandom holds that the inferential connections among sentences are sufficient, and not just necessary, to determine their contents. Unlike hyperinferentialism, Brandom's strong inferentialism allows for a broad set of inferences to determine content. This means that the content conferring relations extend beyond formally good inferences, including material incompatibilities between sentences as well as inferential relations between circumstances of appropriate circumstances and consequences of application.[32] This positions his theory in a middle ground: unlike weak inferentialism it is incompatible with referential semantics; unlike hyperinferentialism it is able to offers an account of how the content of even non-inferential reports can be given inferentially.

30 Ibid., 97.
31 Brandom, “Inferentialism and some of its challenges”, 656-7.
32 Ibid., 658.

Inference is the core of Brandom's semantics but it is not alone sufficient to explain all linguistic content. By itself, inference is only sufficient to pick out and confer content on grammatical terms complex enough to express propositions, typically sentences. Strictly speaking, what I should have said in the examples above is that the content of "Calgary is south of Edmonton" as a whole is partially determined by the fact that asserting that commits one (whether one knows it or not) to the proposition "Edmonton is north of Calgary". The semantic content of entire expressions is derived from the role they play as premises and conclusions in inferences. However, subsentential expressions cannot serve as premises and conclusions.

Brandom adopts a two-fold strategy for explaining the content of subsentential expressions. First, a strategy of substitution adopted from Frege in order to identify two kinds of contentful subsentential expressions.[33] Then, anaphora is invoked to explain how content is conferred on those subsentential expressions. Brandom construes anaphora as "a special mechanism for the inheritance of substitution-inferential commitments".[34] Anaphora is a fortunate aspect of language that allows us to refer to the same concepts even if we do not share the same beliefs about inferential roles. At the subsentential level inference and substitution help define the semantic significance of subsenten-

33 Brandom, Making it Explicit, 281.
34 Ibid., 283.

tial types but it is anaphora which defines the semantic significance of tokens.[35] This yields a three-part structure to the meaning relation: inference allows for the connection between sentences and propositional content; substitution allows for a connection between propositional and conceptual content; anaphora allows for the connection between conceptual content and non-linguistic reality. Fodor's arguments against Brandom ignore anaphora. For whatever reason, in Articulating Reasons, Brandom's summary of Making it Explicit and the source of Fodor's understanding of his project, substitution is the focus and anaphora is left to the side. This is unfortunate since anaphora is necessary to complete Brandom's story about the content of concepts.

Substitution does much semantic work in Making it Explicit. It is employed at every level of the explanation of content from propositions on down. My focus is on how it provides a means of sorting subsentential expressions by semantic content type. Prior to that, however, it is employed to define logical validity.[36] Substitution is used to generate multi-value assignments from a antecedently understood notion of designated value (i.e. a means of generating what is typically called truth value assignments with the reading of "commitment" and "entitlement" in place of truth).[37] The key principle in all cases is that substitution

35 Ibid., 466.
36 Ibid., 346.
37 Ibid., 341.

sorts linguistic components into equivalence classes according to invariance under some property.

Syntactically, the relevant property is grammatical well-formedness. Brandom defines grammatical categories thus: “Two subsentential expressions belong to the same syntactic or grammatical category just in case no well-formed sentence [...] in which the one occurs can be turned into something that is not a sentence merely by substituting the other for it.”[38] Semantically, the relevant property is pragmatic potential: “Two subsentential expressions share a semantic content just in case substituting one for the other preserves the pragmatic potential of the sentences in which they occur”.[39] “Pragmatic potential” means the set of claims that a given sentence commits or entitles one to make. It would be circular, after all, to define semantic content by appeal to invariance under a semantic property like truth preservation.

Only singular terms and predicates are identified and defined as subsentential components. Of these two kinds of subsentential expressions, Brandom, pays most attention to singular terms. Singular terms are, after all, often considered the paradigmatic way language connects us to the world. Brandom has to necessarily tell a longer story about how singular terms perform this function because their content is not conferred from an

38 Ibid., 368.
39 Ibid.

immediate intentionality between minds and the world but rather is derived from the pragmatic content of propositions.

Singular terms are picked out by two necessary conditions, one syntactic and one semantic. Jointly these conditions are sufficient for characterizing subsentential expression as singular terms. Syntactically, singular terms play the role of being substituted for in component expressions. What is left, the part of the expressions which remains unchanged before and after the substitution is a substitution frame.[40] As an example, in the sentence “Brandom admires Hegel”, “Brandom” can be substituted for “The philosopher who wrote Tales of the Mighty Dead” yielding “The philosopher who wrote Tales of the Might Dead admires Hegel.” which preserves grammatical well-formedness. Of course, well-formedness would also be preserved if we substituted “Fodor” for “Brandom”, yielding the novel sentence, “Fodor admires Hegel”. Since “Fodor” is synonymous with “The philosopher who wrote “Having Concepts: A Brief Refutation of the 20th Century”, we can assess the meaning of this novel sentence. Specifically, we can predict that it is false.

“Brandom admires Hegel” is a complex expression which can be substituted into. Brandom is a singular term which necessarily can be substituted for, which leaves the derived syntactic substitution role of being a sentence frame, “x admires Hegel”,

40 Ibid., 368-9.

where having the syntactic substitution role of being a sentence frame is necessary for being a predicate. First different values of x are identified by equivalence classes which give us different ways to pick out x. For any x, then, we can sort synonymous from non-synonymous x's. Then, we can predict novel uses of that sentence frame by substituting in non-synonymous x's by using other ways of identifying the reference of that singular term — typically a definite description.

Semantically, the mark of singular terms is that they are only involved in symmetric substitution inferences. Predicates, by comparison, are involved in some asymmetric substitution inferences. An inference is symmetrical if it is reversible.[41] For example, substituting "Brandom" for "The author of Tales of the Mighty Dead" will preserve appropriateness of inference in any substitutional frame. If we can infer from "Brandom admired Hegel" to "The author of Tales of the Mighty Dead admired Hegel" then we can infer from "The author of Tales of the Mighty Dead" admired Hegel" to "Brandom admired Hegel". Substituting for singular terms always produces inferences which are symmetric in this way. Replacing predicates however, will sometimes yield inferences which are only valid one way. For example, we can infer from "Brandom is person" to "Brandom is a mammal" but his being a mammal does not imply his being a person.

41 Ibid., 371.

This is another point where the critic might demand of Brandom something like an analytic/synthetic distinction. Substitution of singular terms only yields symmetrical inferences. Replacing predicates sometimes yields symmetric inferences, and sometimes not. For example, replacing “x is a bearded philosopher” with “x is a philosopher with a beard” yields a symmetric inference. Since singular terms are identified as only yielding symmetric inferences an epistemic problem might be raised. Consider the perspective of an alien linguist sorting our subsentential expressions into singular terms and predicates. Identifying predicates seems to be a relatively straightforward affair, if you find that replacement yields just one asymmetric inference, then it must be a predicate. However, since predicate replacements can sometimes, maybe often, yield symmetric inferences, how much evidence would be needed to conclude that the subsentential expression in question in fact has the semantic content of a singular term?

Another immediate objection is to question why playing the role of “substituted for” is a necessary condition of being a singular term. What is stopping us from thinking that predicates like “admires Hegel” can be substituted for, leaving the sentence frames like “Brandom x's”? Such replacements can be made, but Brandom must differentiate them from terms that are substituted for. Brandom's notes this objection after introducing the syn-

tactic conditions of singular terms but his resolution of the difficulty comes only after introducing the semantic distinction of singular terms from predicates in a section handily titled “The Argument”.[42]

This argument draws the connection between syntax and semantics which the principle of compositionality seems to demand. A key reason for distinguishing between substituted-for expressions and resulting sentence frames is to account for the productivity of language. A substitution can preserve all the inferences of the original sentence. A replacement, on the other hand, can inferentially strengthen or weaken a sentence. Distinguishing the two kind of syntactic roles and connecting them to symmetric and asymmetric roles in inferences allows for the production and evaluation of novel sentences. Having a syntactic criterion for identifying singular terms and predicates allows us to exploit the syntactic construction in evaluating novel utterances.

Indeed, in an appendix Brandom shows how to use the substitutional derivation of categories to define a functional derivation of categories.[43] Key to this functional derivation is the ability to clearly distinguish sentence frames from grammatical terms playing the role of being substituted for. Finally, the key to <u>that</u> is “The Argument” that grammatical terms which are sub-

42 Ibid., 378-381.
43 Ibid., 404-9.

stituted for are restricted to those whose occurrence is inferentially symmetric and that substitutional frames are necessarily asymmetric.

The account of substituted-fors and sentence frames is motivated by a number of prima facie differences between substitution (which can only operate on expressions) and replacement (which can only operate on frames. First, frames are derivative syntactic categories in that they result from substitution. In that sense they are better thought of as linguistic patterns than expressions. A frame, then, is not a constituent of a sentence put a product of analyzing it. This analysis reveals many occurrences of expressions that can be substituted for but only one frame resulting from such substitutions. Finally, replacing frames is different from substituting expressions because it requires us to keep track of argument places and cross references among them. Sentence frames have a fixed number of adicities which must be preserved in order to preserve well-formedness.[44]

Brandom argues that if a language allowed for subsentential expressions to syntactically play the role of being substituted-for but semantically correspond to asymmetric simple material substitution commitments, that language would lose the expressive ability of both conditionals and negations. Both conditionals and negations allow for the production of an inferen-

44 Brandom, Articulating Reasons, 132.

tially complementary sentence frame from any arbitrary sentence frame. If subsentential expressions that were substituted for were inferentially symmetrical, this would not be possible.

The argument relies on notions of the inferential strength and weakness of claims and complementary inferences. A claim, P, is inferentially stronger than another claim, Q, if every claim that is a consequence of Q is also consequence of P but there are some claims which are a consequence of the conjuctive P & not-Q. For example, “Wulf is a dog” is inferentially stronger than “Wulf is a mammal”. Everything one is committed to by claiming that “Wulf is a mammal” (perhaps including that he is a vertebrate, that he drank milk as an infant, that he has hair) one is also committed to by claiming that he is a dog. However, some claims that one is committed to by claiming that he is a dog (that he barks perhaps) do not also follow from a commitment that he is a mammal (cats are mammals but do not bark, for example).

For two sentence frames, say Px and P'x, P'x is inferentially complementary if whenever (Pa → Pb) and ¬(Pb → Pa) then (P'b → P'a) and ¬(P'a → P'b). One way to read this intuitively is by thinking of negating both sides of a conditional. Given a inferentially asymmetric conditional such as “if it is a dog then it is a mammal”, we can produce the inferentially complementary frame “if it is not a mammal then it is not a dog”. The claim is

that this feature of language would be lost if substituted-fors could have an asymmetrical inferential significance.

Imagine that singular terms could work in this inferentially asymmetric way. For example imagine we could infer from “Wulf is a mammal” to “Dulf is a mammal” but not vice-versa and that we can infer from “Dulf is a dog” to “Wulf is a dog” but not vice-versa. Clearly no symmetrical inferential role is governing these inferences but no asymmetrical role is governing them either. Substituting “Wulf” with “Dulf” either inferentially strengthens or weakens a sentence but not both. This precludes us from codifying inferences governing these substitutions because substituting “Wulf” with “Dulf” has opposite results depending on whether the frame is “x is a dog” or “x is a mammal”. This argument generalizes if we imagine, first that “Wulf” and “Dulf” behave in this complementary fashion for every predicate in which they appear, then that there is a general formula for creating an inferentially complementary frame for any given frame. Finally, note that any language which did not have such a general formula would not be able to form conditionals or negations because that is what conditionals and negations do.

Now, no atomist would take issue with singular terms playing this dual syntactic/semantic role. The issue is whether the substitutional notion of composition allows them to do so. But arguing that it does is only half the story. The other half is to

argue that a plausible notion of content can be bestowed on singular terms in this manner.

Brandom says, “[t]o introduce a singular term into a language one must specify not only criteria of application but also a criteria of identity, specifying which expressions are intersubstitutable with it”.[45] The fact that a singular term only yields symmetric inferences is actually a part of its semantic content. Thus the fact that a singular term is a singular term is part of the content of that singular term. This should also be true for a representationalist — surely the representations of objects are different in kind from representations of properties or relations. What puzzles the atomist, but not the inferentialist, however, is how singular terms that fail to denote can be representations and hence be concepts. This is not puzzling for the inferentialist because failing to denote does not preclude having an inferential role.

Instead, the inferentialist faces difficulty in explaining the introduction of singular terms for new objects. If a novel singular term names an object already known, a criterion of identity is trivially an identity statement between the two names. Coining a new singular term for a new object, however, requires giving at least one way, other than the newly coined name, for recognizing that object. According to this view, objects cannot

45 Brandom, Making it Explicit, 372.

be introduced by a simple baptism. Expressions which cannot be substituted for would not be recognized as singular terms. The idea of an object that can be referred to in only one way, Brandom maintains, is not the idea of an object at all.[46] Although atomism does not also require this, it should not seem farfetched. An object can be named. But if we are to be able to recognize that same object in different instances, there must also be an expressible criterion for recognition. This gives us, already, two ways of picking out that object.

Brandom's approach follows Frege's in Foundations of Arithmetic. In that work, Frege wishes to introduce numbers as unique objects. His concern is to justify entitlement to the use of the definite article when purporting to referring to an object. In order to justify that entitlement, it must be shows that some object falls under the concept and that only one object falls under it.[47] In order to do this, we must "fix the sense of an identity statement", which is equivalent to giving a criterion for recognizing that object.[48] It is not important that we be able to apply this criterion infallibly — what is important is that "the notion of correctness of identifications and discriminations must have been settled somehow."[49]

46 Ibid., 425.
47 Ibid., 415.
48 Ibid., 417.
49 Ibid., 416.

Frege introduces the method of abstraction for introducing new objects. Essentially, new terms can be introduced using an equivalence relation which is defined on existing terms. Frege's example is to introduce "the direction of" by appeal to the equivalence relation "is parallel to" defined on lines:

> the direction of a = the direction of b iff a is parallel to b

Frege imposes a strong condition on such identity statements. He believes that you must settle the truth-values of all identities of every object. If "the direction of a" and "Julius Ceasar" are objects, for example, then it must be settled whether "the direction of a = Julius Ceasar" is true. This strong requirement, however, is perhaps due to Frege's goal in The Foundations of Arithmetic of laying an absolutely rigorous foundation for mathematical proofs. In explaining linguistic meaning generally however, it seems fine to just accept that it is unsettled whether or not Julius Caesar is the direction a. In general discourse, it is not only acceptable that terms which purport to refer to a singular object fail to refer, it is sometimes desirable that they do so (as in the case of referring to fictional objects).

Instead, then, Brandom proposes a minimum condition on the introduction of a singular term: "it must have been settled that the occurrence of the putative term have some (symmetric) substi-

tution-inferential significance".[50] This is sufficient for the singular term to play a role a linguistic role and be passed on in communication regardless of whether it indeed refers to one object.

4.3 From Words to the World

Brandom's semantics reverses a number of intuitive aspects of atomism: contrary to what atomism suggests, anaphora is prior to deixis because in order for a concept to refer to an object the concept must be repeatable and the object must be recognized time and again; contrary to the idea that empirical vocabulary provides a basis for understanding the representational nature of thought, Brandom argues that "the ability to use ordinary empirical descriptive terms [...] already presupposes a grasp of the kinds of properties and relations made explicit by modal vocabulary";[51] contrary to the notion of intentionality inhering first in the mind of an individual, Brandom holds that (conceptually meaningful) intentionality comes through linguistic discourse. In a sense, then, a person only has an intentional grasp of the world if that person's linguistic peers attribute such a grasp to her. Grasp of the world here comes only after the use of language in a social setting, whereby we are challenged to provide reasons for our commitments. The holistic story ends where the atomist

50 Ibid., 424.
51 Brandom, Between Saying and Doing, 96-7.

story begins: with representation of the world. Since this story is admittedly less intuitive, the justification must be that the atomist order of explanation is impossible: deixis cannot be explained without anaphora; empirical vocabulary cannot be explained without modal vocabulary; and conceptual intentionality cannot be explained without the social use of language.

The notion that deixis requires anaphora in order to be conceptually meaningful is anticipated by the considerations that objects must be recognizable as the same in different contexts, and that linguistic tokens which pick out those objects must be repeatable. A demonstrative tokening by itself is generally unrepeatable because changes in perspective and changes in the state of the world can make a token “that x” pick out different x's at different times. We can only repeat a reference to an object picked out demonstratively because demonstratives are treated as the initiators of anaphoric chains. As such, demonstratives have substitutional-inferential significance.[52] If you are able to pick out “this book”, you (or even I) can repeat the reference with “that book” or “it”. The very act of picking out an object using a demonstrative and predicating something of it will make that reference repeatable.

There are also related arguments that demonstratives must be inferentially articulated in order to be able to pick out ob-

52 Brandom, Making it Explicit, 462.

jects. A demonstrative utterance must have a sortal at least implicitly attached in order for it to be meaningful. “Have you ever seen that before?” is ambiguous between a specific and general reading: have you ever seen Sam the three-legged dog versus have you ever seen a three-legged dog before? The atomist order of explanation presumes that utterances like “that is red” can meaningfully predicate without an inferential role being attached to “that”. The argument here, however, is “that” cannot even pick out an object to be predicated without an inferential role.

This forms considerations against the “direct reference” aspect of semantic atomism. The idea is that direct reference to the world, using deixis, can form the semantic building blocks for a theory of meaning. The, perhaps common-sense, intuition this exploits is that meaningful discourse can find its beginnings in the ability to pick out an object and predicate something of it. Under this view, the ability to demonstratively refer to an object is prior to any anaphoric chain which allows that reference to be repeated. Brandom's response is that this is an incoherent position since an object cannot be picked out unless it is done in a way which carries a repeatable, inferential significance.

Relatedly, the holistic semantics Brandom endorses turns the typical view of the relation between empirical and modal vocabulary on its head. The representational aspect of semantic

atomism places empirical, non-inferential reports at the base of meaning because this aspect of communication coincides with a view of conceptualization as representation. Infamously, however, placing this importance on empirical vocabulary creates a tension with the modal vocabulary of necessity and possibility. Ever since Hume, Brandom argues, empiricists have embraced a stark choice between either explaining modal vocabulary in non-modal terms or living without it.[53] In his own view, however, it is modal vocabulary which serves naturalist purposes, because modal vocabulary describes the laws which govern the natural world.

Semantic atomism presupposes that empirical discourse is intelligible antecedently of modal discourse. By contrast, Brandom endorses what he terms the Kant-Sellars thesis that the ability to use empirical descriptive terms presupposes a grasp on the law-governed relations expressed by modal vocabulary. This comes from Kant's insight that the application of all empirical concepts implicitly involves counter-factual supporting dispositional commitments: "One has not grasped the concept cat unless one knows that it would still be possible for the cat to be on the mat if the lighting had been slightly different, but not if all life on earth had been extinguished by an asteroid-strike."[54] The moral that Brandom draws is that if one knows how to use ordinary empirical vocabulary, then one also knows everything that is

53 Brandom, Between Saying and Doing, 96.
54 Ibid., 97.

needed in order to introduce modal vocabulary. Intuitively, by referring to objects in the world and predicating things of them we are presupposing a law-governed world.

The argument for the Kant-Sellars Thesis is that without it language users face an intractable updating problem with regard to their empirical beliefs. It is entailed by the following premises: every autonomous discursive practice must have some observational vocabulary; those who engage in discursive practice must distinguish between materially good/bad inferences; material inference is in general non-monotonic; at any given time, most of a subject's beliefs could only be justified as conclusions of inferences; and discursive practitioners must be epistemically responsible.[55] Since discursive practitioners cannot review all their beliefs in light of new belief, they must instead associate the new belief with a set of material inferences which could overturn that belief. This set of potential defeasors, in turn, define a range of counter-factual robustness associated with the new belief.

The idea is that commitment to such counter-factuals is already implicit in our ability to use observational vocabulary. If you can observe that "the cat is on the mat", you can implicitly know that if he is on the mat, he is not on the counter but that his being on the mat is irrelevant as to whether it is rain-

55 Ibid., 106-8.

ing (for example). Modal vocabulary allows us to make the use of these practices explicit.[56] The key is to show how modal vocabulary can be introduced on the basis of these counter-factual conditionals. The Humean problem of induction is here seen as a semantic problem. If the meaning of empirical, observational vocabulary is derived solely from representation, then modal vocabulary will be a mystery. However, if the meaning of empirical, observational vocabulary is determined by the practices implicit in employing that vocabulary, then modal vocabulary comes out simply as a description of those very practices.

These same arguments are applied to understanding normative modal vocabulary. In order to use empirical, descriptive vocabulary, we must also be able to do everything needed to introduce normative vocabulary.[57] Because the asserting and inferring linked together as mutually necessary aspects of meaningful communication, every autonomous discursive practice must include the practices of giving and asking for reasons.[58] Being able to give and ask for reasons, in turn, requires identify two kinds of normative status of claims: those to which one is entitled and those to which one is committed.[59]

The upshot of all this is that the empiricists difficulty with modal vocabulary is one which disappears when we allow that

56 Ibid., 109.
57 Ibid., 110.
58 Ibid., 111.
59 Ibid., 112.

the use of vocabulary in make inferences determines its conceptual content. Such inferential use already implicitly employs the normative practices and alethic beliefs which modal vocabulary makes explicit. We could not refer to the world unless we first believed that there was a law-governed world to be referred to.

Fortunately, for the possibility of analysis, both alethic and normative modal vocabularies can be founded on a basic notion of incompatibility. In Between Saying and Doing Brandom offers what he calls “incompatibility semantics” from which modal and empirical vocabularies — mutually necessary aspects of discursive practice — can be understood. It starts with the suggestion to “represent the propositional content expressed by a sentence with the set of sentences that express propositions incompatible with it”.[60] Incompatibility can be taken as a semantic primitive which itself can be understood in terms of our practical engagement in the world.

Finally, then, this explanation reverses the seemingly intuitive atomist order of explanation between language and intentionality. Atomism exploits from the seemingly plausible idea that individuals can employ language because they are first able to have beliefs about the outside world. The considerations above suggest that this is not so. Having beliefs about the world requires a conceptually articulated view of the world, and a con-

60 Ibid., 123.

ceptually articulated view of the world presupposes the practices which are already necessary for discursive practice. This should not be seen as an absurd denial that creatures cannot represent the world (in some basic way) before being able to describe it. Rather, it is a denial that creatures can represent objects in the world as concepts which can serve as the meanings of natural language. Only inferentially articulated concepts can serve as such meanings, and, in turn, refer to the world.

This amounts to an argument against the possibility of atomist semantics, despite the elegant story it tells of the connection which language users draw between words and the world. Brandom offers an alternative by distinguishing between two kinds of intentionality: practical intentionality which is a basic, pre-conceptual means by which a creature's behaviour is directed toward and constrained by the outside world; and discursive, conceptual intentionality that institutes semantic relations between symbols and the world.[61] He identifies four sequential steps in the pragmatic order of semantic explanation using these two notions of intentionality.

First, the most basic form of intentionality is understood in terms of feedback-governed practical transactions. These are cycles of "test-operate-test-exit" (TOTE) behaviour which exhibits both a differential response and a further response to the

61 Ibid., 190.

effects of that response. This is enough to differentiate between the behaviour of sentient organisms and something like a thermostat which merely exhibits a differential response. Creatures that exhibit TOTE interact with the environment in a goal-directed manner. This goal directedness is the most basic necessary condition of a mind which demonstrates an intentionality directed at the outside world since a mere differential response can be instantiated by mindless objects like thermostats or magnets.

Secondly, conceptual, discursive intentionality is understood as a special kind of feedback-governed practical engagement of the world which is mediated by relations of material inference and incompatibility. This idea of the conceptual replaces the notion that intentionality is manifested in distinct relations between representands and representings. The relevant relation, rather is between agents and the world.

The key to understanding this relation, in turn, is to look at the resulting meaning-use relation between normative and modal vocabularies, both of which can be understood in terms of two basic relations of incompatibility: alethic and deontic. On the one hand, feedback-governed interaction with the world presupposes alethic incompatibility. If some particular piece of space-time is red, it cannot also be blue. The particular kind of feedback-governed interaction with the world which results in conceptual intentionality requires a further adherence to deontic incompat-

ibility: if you take something to be a cat, you ought not also take it to be a dog.

This founds linguistic meaning on a pre-linguistic practice, but the bafflement at how language can get started on this story still stands. Conceptual intentionality and discursive practice mutually entail each other. How is it that creatures like us were able to move from basic feedback-governed practical engagement in the world to that special kind of conceptually articulated feedback-governed engagement which goes along with discursive practice?

Answering that question is, I believe, filling in the response to Fodor's compositional arguments against holism. It is, in a sense, unsatisfying to argue that substitutional methods of dividing up compositional content are equivalent to functional methods under the domain principle. Given a meaningful whole and some systematic way of dividing it into components, it is unsurprising that meanings can also be attributed to those components. How are those "meaningful wholes" to be given in the first place? The compositional story, by contrast, seems straightforward because those meaningful wholes are constructed from the meaningful components and the components get their meaning from basic representations. Brandom's criticisms show that that story is not as straightforward as it seems; but that, in itself, is not a vindication of a holistic semantics.

4.4 Holism's Plausibility

The discussion above blends aspects from the conceptually separable projects in Making it Explicit and Between Saying and Doing. As such, it gives us two ways of filling in a response to Fodor's criticisms from compositionality. One way is to draw together the semantics of inference, substitution and anaphora with Brandom's story about deontic scorekeeping and normative commitment/entitlement (admittedly neglected here) to show how discursive practitioners can use novel propositions in inferences and attribute commitments and entitlements based on them. The other way is to fill in the details of the incompatibility semantics mentioned above and explicate Brandom's claim that it provides recursive projectibility without compositionality.[62] Of these two ways, the latter is meant to be more general. It defines a semantics for modal logic by taking the notion of incompatibility as a semantic primitive. The former way, invoking inference, substitution and anaphora together with normative commitment and entitlement can be considered one of the possible ways of filling in the philosophical details of such a semantics.

Incompatibility semantics starts with the suggestion to represent the propositional content expressed by a sentence, p, with the set of sentences incompatible with it.[63] From that notion, Brandom defines two basic logical operators, negation and

62 Brandom, Between Saying and Doing, 133.
63 Brandom, Making it Explicit, 123.

conjunction. The negation of p, Np, is defined as intersection of the incompatibility sets of everything incompatible with p.[64] Similarly the conjunction of p and q, Kpq, is the set of sentences which are incompatible with the set {p,q}.[65] There are a couple of surprising features of incompatibility semantics. One is that although the logical operators are not anything like truth functions, they result in logical operators which behave according to the theorems of classical two-valued logic.[66]

However, they are not compositional. Although content p is associated with the set of everything incompatible with p, and likewise with q, the content of Kpq is not associated with the union of those two sets. Rather, Kpq is associated with the set of everything incompatible with the union of those two sets. This definition is intensional and non-compositional because computing the set of sentences incompatible with Kpq involves sentences other than p and q. It requires computing the incompatibilities of all the sentences which are either incompatible wit p or incompatible with q. However, this captures intuitions about the counter-factual robustness of incompatibility. For example, although being a blackberry is incompatible with being red and being ripe, it is not incompatible with either individually. Blackberries can be red and they can be ripe but they cannot be

64 Ibid., 127.
65 Ibid., 128.
66 Ibid., 128.

both. Even negation is intensional this way, because computing the semantic value of not-p requires the incompatibility of every sentence incompatible with p, hence it involves looking at a lot of other sentences which are not p.[67]

This semantics is clearly a violation of the principle of compositionality as Fodor means it. The argument, however, is that novel expressions can still be semantically evaluated, and hence that productivity, systematicity and learnability are possible because the semantics of complex expressions is still computable from the component expressions. Determining the semantic value of not-P requires first determining the semantic value of P. Semantic evaluation is, in that sense, recursive between levels. It is, however, finitely recursive and semantically grounded, because computing the semantic value of P yields all the information you need to evaluate not-P. Likewise with conjunction. Brandom offers a proof that the interpretive frame for a language with logically complex sentences can be reduced to the frame for syntactically less complex fragments of the language.[68] This amounts to a proof that a extension to a language created by introducing complex expressions composed of simple components already contained in the language can be given a semantic evaluation in a non-circular way.

67 Ibid., 135.
68 Ibid., 147.

This amounts to knocking down formal aspects of the compositionality argument. Chapter 1 anticipated as much. The further question, however, is whether those who employ language could plausibly evaluate novel utterances in a way akin to this. Recall that the plausibility of the atomistic view to produce novel evaluations was never really called into question. Instead, it was the ability to extend the atomistic view of meaning to an appropriate explanation of use that was called into question. The basic problem has, then, been reversed. The holistic approach builds pragmatist notions into meaningful content. I have discussed numerous arguments that such an approach provides a theory of meaning which is recursively computable in principle. The question now is whether it is psychologically plausible. This is unsurprising because all the support mustered for compositionality, from productivity, learnability, and systematicity, is essentially psychological in nature.

This question is best addressed, I think, by moving from consideration of the incompatibility semantics to the more detailed story offered in Making it Explicit. One connection between the two works is that the semantic notion of incompatibility which forms the basis of incompatibility semantics is itself to be understood on a prior, pragmatic notion of inference which employs the normative concepts of commitment and entitlement. Between Saying and Doing argues that all discursive prac-

tice presupposes these normative notions, and that they are sufficient for carrying out a number of projects and methods of analytic philosophy. Making it Explicit, on the other hand, draws out the details of a theory of language and mind which is based on this commitments.

An oft-repeated argument against the psychological plausibility of holistic semantics comes from the supposed incommensurability of holistic concepts. The concept "electron" is a popular example: suppose that Einstein and Rutherford both recursively evaluate the concept "atom" from a prior set of commitments and entitlements.[69] Clearly both see themselves as committed and entitled to entirely different propositions in which the concept "electron" occurs.

Although Brandom calls this the biggest challenge that has always faced semantic theories, his response is in fact a measured dismissal.[70] He argues that incommensurability only becomes mysterious for holism if one holds all of 3 views about communication: communication is a transfer of meaning rather than reference; communication happens according to a Lockean model of sharing the same idea; and a Cartesian view of the concepts shared as mental particulars rather than a Kantian view of them as rules.[71] Alternatively, Brandom views communication as a shared

69 Extending the incompatibility semantics discussed above from a propositional calculus to a prepositional one would determine the content of constituent concepts in this recursive way.
70 Brandom, "Inferentialism and some of its challenges".
71 Ibid., 670.

practice whereby people are able to track each others' inferential commitment and capacities by bringing together each others' notions of the applicability of certain concepts. Finally, also, there is the point that people can still refer to the same object even with very different beliefs about that object. At bottom, inferential semantics is a theory of sense and senses are meant to be modes of presentation which determine the reference.

The semantics of Making it Explicit address this incommensurability. The entire structure of inference, substitution, and anaphora is aimed at explaining how we are able to share reference without sharing senses. Substitution and anaphora allow us to refer to the same things without sharing many beliefs about those things, but to refer to the same things without sharing any beliefs about them. Since we can refer to the same objects, communication will surely allow us to find some inferences about those objects that we share. At least this is most obvious with objects revealed in a plain, empirical manner. Brandom's oft-employed example is communication with a Zoroastrian sun-worshipper about the sun. Although Brandom and the sun-worshipper mean different things when they refer to the sun, the fact that they both refer to the sun can be easily revealed by the fact that they will share at least some commitments about it. For example, whether you believe that the sun is a ball of gas or a god, you will also believe that if it is not in the sky it is dark.

Summary

The primary goal of this chapter was to explicate Brandom's inferential semantics in way as to show how he can accommodate linguistic productivity. The answer is two-fold. First, there is general framework for a sentential semantics which can project the meanings of novel sentences from Between Saying and Doing. Second, there is the specific semantics of inference, anaphora, and substitution, gleaned from Making it Explicit, which derives semantic content from the social acts of giving and asking for reasons.

§4.1 attempted to justify prioritizing sentences as the primary bearers of meaning. This justification comes in the form of criticisms of the representational aspects of atomism, inspired by considerations from the Philosophical Investigations.

§4.2 drew on the semantics from Making it Explicit to explain how inferences can confer semantic content. First, the conceptual content of sentences is identified by their inferential role. Next, substitution allows us to identify and define subsentential components on the basis of the contribution they make to the sentences inferential role. Finally, anaphora ensures that language users can share and communicate those concepts.

§4.3 tried to complete this semantic story but showing how it connects to the world. This discussion relied on Brandom's presentation of what he calls the Kant-Sellars thesis, which is

that modal vocabulary is required to describe how we can employ straightforward empirical vocabulary. Our conceptual connection to the world, in this theory, is through our actions in the world, rather than our representations of it. Two kinds of intentionality were identified as a partial answer to the question of how such a semantics can get started: a primary, pre-conceptual intentionality which arises out of our goal-directed interactions with the world, and a conceptualized , meaningful intentionality which arises from the justifications we give for those actions.

§4.4 finally addressed two of Fodor's complaints against inferential role semantics. I gave an explication of how the incompatibility semantics of Between Saying and Doing allows for the projection of novel meanings, even though it is not compositional. I also explained how it is that people can share a lot of the same concepts even if they do not attach identical conceptual roles to those concepts.

Conclusion

The metaphysical debate between semantic atomism and holism connects with empirical studies of natural language in two debates. One is about theories about L1 language learning, and the other is about theories about "protolanguage". Both such theories are about the origins of language. L1 language learning concerns the origins of language within the minds of individuals. Protolanguage concerns the historical origins of linguistic communication amongst humankind. Holism and atomism both suggest vastly different explanations in both domains. Also, both struggle with what might be called the miracle of language. Natural languages, with their syntax, often seem to be an all or nothing affair, somehow triggered amongst humans, who are then able to employ them with great efficiency.

Atomism suggests precisely the "Augustinian picture" of language which Wittgenstein criticizes. The origins of language here are pretty much pointing and naming. Chomsky's notion of an innate universal grammar works well with this. The child has an innate capacity to learn language — a general grammar starting point. All the child needs is to learn the names for things and adjust his innate grammar to match that of his peers. The mystery is then pushed onto the explosion of lexical concepts which occurs. As we saw in Fodor, one is tempted to treat concepts like "doorknob" and "curry" as innate, although it seems strange to

think how someone can be born with an innate tendency to call things doorknobs or curries, or even what all that amounts to.

The mystery happens earlier in the holistic explanation, and perhaps this is why it seems less intuitive. Perhaps you will concede, as I have maintained, that there is no logical implausibility to a two step semantics whereby the meanings of complex expressions are analyzed into components, from which new complex expressions can be constructed. But how is it that people can just grasp the meanings of some set of complex expressions in the first place?

I think that a debate in theories of protolanguage can shed some light on how this might be plausible. Alison Wray has argued that the first systems of holistic communications must have been holistic in nature.[1] She argues that most primitive communication was likely to have come in the form of preset holistic phrases, which served three functions: the successful manipulation of others, identify group membership, and holding the attention of others.[2] Wray argues that there are only two ways that a consistently retrievable meaning can be associated with a particular utterance: have a grammar which favours one interpretation, or come to a consensus that a certain sound will have a certain meaning.[3] Presuming that associating meanings with sounds predates a gram-

1 Wray, “Protolanguage as a holistic systemic for social interaction”.
2 Ibid., 55.
3 Ibid., 49.

mar, only the second option was available for the first human speakers.

The idea is that some complex sounds can be identified with a complex meaning. Most likely the first meanings to be communicated were demands such as “put the meat on the hot stone”. We know when such a meaning is grasped, because the person will have displayed the correct behaviour. At some point, the introduction of a grammar allowed us then to analyze these expressions into component meanings, and associate those component meanings with particular parts of the utterance. This, then, allowed us to create novel utterances, whose meaning was not conferred in advance my consensus.

The analytic system which produces novel sentences now plays a large role in much of our language. So much so, that it can seem that all our language is so used. There are, however, still holistic aspects to our language — idiomatic phrases, for example. Wray argues that holistic phrases are still important for successful communication because the analytic method of communicating new meanings, complex meanings is computationally expensive.[4]

I admit that this thesis has taken place in an enormous shadow of ignorance about how people actually do process natural language. I would be will to bet, however, that we use both com-

4 Ibid., 63.

positional and holistic methods to understand and communicate meanings.

If that is true, then the question is which ontology of meaning is best suited to both holistic and analytic processing methods. I think that the skeptical arguments from Plato and Wittgenstein suggest that maintaining a strictly atomist ontology, whereby simple concepts are independent of any other, is much more problematic than it first appears. Further, the holistic approach, although it has a lot of difficulty explaining how our meanings are connected to the outside world, is much more promising than it first appears.

Bibliography

Brandom, Robert B. 2008. Between Saying and Doing: Towards an Analytic Pragmatism. Oxford: Oxford University Press.

———2007. “Inferentialism and some of its challenges”. Philosophy and Phenomenological Research 74, (3) (May): 651-676.

———2000. Articulating Reasons. Harvard: Harvard University Press.

———1994. Making it Explicit: Reasoning, Representing, and Discursive Commitment. Harvard: Harvard University Press.

Chan, Wing-Tsit. 1963. A Source Book in Chinese Philosophy. Princeton: Princeton University Press.

Chomsky, Noam. 1997. “Language and problems of knowledge”, Teorema, 16(2), 5-33, 29 p.

———1992. “Explaining language use”, Philosophical Topics, 20 (1), 205-231, 27 p. Spring 1992.

———1959. “Review of Verbal Behaviour”. Language 35, (1) (January-March): 26-58.

———1957. Syntactic Structures. The Hague/Paris: Mouton.

Cowie, Fiona. 1997. “The logical problem of language acquisition” Synthese, 111(1), 17-51, 35 p. April.

Deshpande, Madhav, "Language and Testimony in Classical Indian Philosophy", The Stanford Encyclopedia of Philosophy (Spring 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/language-india/>.

Davidson, Donald. 1967. “Truth and meaning”. Synthese: An International Journal for Epistemology, Methodology and Philosophy of Science 17, (September): 304-323.

Dummett, Michael. 2006. Thought and Reality. Oxford: Clarendon Press.

———1973. Frege: Philosophy of Language. New York, NY: Harper & Row.

Fodor, Jerry A. 2008. LOT 2: The Language of Thought Revisited. Oxford: Clarendon Press.

———2004. “Having concepts: A brief refutation of the twentieth century”. Mind and Language 19, (1) (February): 29-47.

———2001. “Brandom’s burdens: Compositionality and inferentialism”. Philosophy and Phenomenological Research 63, (2) (September): 465-481.

———2001. “Language, thought and compositionality”. Mind and Language 16, (1) (February): 1-15.

———2000. “Reply to critics,” Mind and Language, 15(2-3), 350-374, 25 p. April-June

———1998. Concepts: Where Cognitive Science went Wrong. New York: Clarendon Press.

Fodor, Jerry, and Lepore, Ernie. 2007. “Brandom beleaguered”. Philosophy and Phenomenological Research 74, (3) (May): 677-691.

Frege, Gottlob. 1993. “On sense and reference” in A.W. Moore, Meaning and Reference. London, UK: Oxford University Press.

———1993. “Letter to Jourdain” in A.W. Moore, Meaning and Reference. London, UK: Oxford University Press.

———1980. The Foundations of Arithmetic. Oxford: Basil Blackwell.

———1963. “Compound Thoughts”. Mind Vol. 72 No. 285 (January): 1-17.

Grandy, Richard E. 1990. “Understanding and the principle of semantic compositionality”. Philosophical Perspectives 4, Action, Theory and Philosophy of Mind: 557-572.

Griffiths, Paul E. 2002. “What is innateness?”. The Monist vol. 85, no. 1: 75-85.

Harte, Verity. Plato on Parts and Wholes - The Metaphysics of Structure. Oxford: Oxford University Press, 2002. Oxford Scholarship Online. Oxford University Press. 20 December 2010

Hintikka, Jaakko. 1984. “A hundred years later: the rise and fall of Frege's influence in language theory”. Synthese 59, 27-50, 24 p. April 1984

Jackendoff, Ray S. 1992. Languages of the Mind. Cambridge, MA: MIT Press.

Jackman, Henry. 2005. “Descriptive atomism and foundational holism: semantics between the old testament and the new”. Protosociology, Vol 21: Compositionality, Concepts and Representations.

Janssen, Theo M V. 2001. “Frege, contextuality and compositionality”. Journal of Logic, Language and Information, 10(1), 115-136, 22 p. Winter 2001

Kripke, Saul A. 1982. Wittgenstein on Rules and Private Language. Cambridge, MA: Harvard University Press.

Lahav, Ran. 1989. “Against compositionality: the case of adjectives”. Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 57(3), 261-279, 19 p. November 1989

Levine, James. 2002. “Analysis and decomposition in Frege and Russell”. Philosophical Quarterly 52, (207) (April): 195-216.

Linsky, Bernard. 2003. “The metaphysics of logical atomism”. The Cambridge Companion to Bertrand Russell, Griffin, Nicholas (ed). Cambridge: Cambridge Univ Press.

Linsky, Leonard. 1992. “The unity of the proposition”. Journal of the History of Philosophy 30 (2) (April): 243-273.

Macbeth, Danielle. 2007. “Logical analysis, reduction, and philosophical understanding”. Croatian Journal of Philosophy, 7(21), 475-485, 11.

Meinong, A. 1960. “The Theory of Objects” in Realism and the Background of Phenomenology, ed. Roderick Chisholm. Glencoe, Ill: Free Press

Pagin, Peter. 1997. “Is compositionality compatible with holism?” Mind and Language 12, (1) (March): 11-33.

Pagin, Peter, and Westerståhl, Dag. 2010. “Compositionality II”. Philosophy Compass 5, (3) (May): 250-264.

———2010. “Compositionality I”. Philosophy Compass 5, (3) (May): 250-264.

Panjvani, Cyrus. 2008. “Rule-following, explanation-transcendence, and private language”. Mind 117, (466) (April): 303-328.

Partee, Barbara H. 2009. “Formal semantics, lexical semantics, and compositionality: the puzzle of privative adjectives”. Philologia 7:7-19

———1988. “Semantic facts and psychological facts”. Mind and Language, 3, 43-52, 10 p. Spring 1988.

Pelletier, F.J. Forthcoming. “Holism and Compositionality” in W. Hinzen, E. Machery, M. Werning The Oxford Handbook of Compositionality (Oxford UP).

———2000. "Did Frege Believe Frege's Principle?" Journal of Logic, Language, and Information 10: 87-114.

———1994. "The principle of semantic compositionality" Topoi 13: 11-24.

Perry, John. 1994. “Davidson's Sentences and Wittgenstein's Builders”. Proceedings and Addresses of the American Philosophical Association (68), (2): 22-37.

Jacob, Pierre, "Intentionality", The Stanford Encyclopedia of Philosophy (Fall 2010 Edition), Edward N. Zalta (ed.), URL =<http://plato.stanford.edu/archives/fall2010/entries/intentionality/>.

Plato. The Collected Dialogues of Plato. Translated by Benjamin Jowett. Intelex Corporation (online resource).

Russell, Bertrand. 1985. The Philosophy of Logical Atomism. La Salle: Open Court.

Ryle, G. 1990. “Logical atomism in Plato's Theaetetus”. Phronesis: A Journal of Ancient Philosophy 35, (1): 21-46.

———1947. “Knowing how and knowing that: the presidential address”. Proceedings of the Aristotelean Society 46: 1-16.

Szabó, Zoltán Gendler. Forthcoming. “The case for compositionality” in W. Hinzen, E. Machery, M. Werning The Oxford Handbook of Compositionality (Oxford UP).

————. Problems of Compositionality. Unpublished Ph.D. dissertation.

Tarski, Alfred. 1944. “The semantic conception of truth and the foundations of semantics”. Philosophy and Phenomenological Research 4, (March): 341-376.

Wittgenstein, Ludwig. 2010. Tractatus Logico-Philosophicus. Translated by C.K. Ogden. Project Gutenburg.

————2001. Philosophical Investigations. Translated by G.E.M. Anscombe. Oxford: Blackwell Publishing.

Zadrozny, Wlodek. 1994. “From compositional to systematic semantics”. Linguistics and Philosophy 17: 329-342.