ABSTRACT

Title of Document: ON THE REPRESENTATION OF OBJECTS IN THE VISUAL SYSTEM

Dimiter Iavorov Kirilov, Ph.D., 2012

Directed By: Professor Georges Rey, Philosophy Department, UMCP

There is evidence in the psychological literature for representations of objects (Pylyshyn's visual indexes) that refer to and track, not properties, but what in our sort of world typically turn out to be individual physical objects. I am concerned with how such representations acquire their content.

Two strategies for accounting for the content of representations are a) representations of particulars refer to the entity that caused them; and b) representations of particulars refer to the entity whose properties are represented by the visual system. The first strategy faces the "which link" problem: since any one of the links in the causal chain leading to the token representation counts as a cause of the token representation, no particular link is individuated as the referent. I examine a recent proposed solution to this problem (Fodor's counterfactual triangulation) and conclude that it fails to determine whether the referent of a visual index is an object, as opposed to a state of affairs, or an event.

The problems with the first strategy are a reason to explore the second strategy: representations of objects refer to the entity whose properties are represented by the visual system. I adopt Fodor's asymmetric dependency account (ADA) of intentionality to account for how representations of properties get their content. Fodor's account is chosen not because it is free of problems, but because it has the structure of a theory that promises to deal with many of the classic problems that befall informational semantics (e.g. the disjunction problem).

Since ADA is designed to work for causal relations between properties and not for causal relations between particulars, it cannot, by itself, account for how representations of particulars get their content. So I suggest that ADA be supplemented with conceptual role semantics to account for the logico-syntactic roles of representations of particulars. In particular, I suggest that to represent objects the visual system requires the capacity to form and store in memory definite descriptions containing: a) predicates referring to spatio-temporal relations; and b) temporal indexicals.

ON THE REPRESENTATION OF OBJECTS IN THE VISUAL SYSTEM

By

Dimiter Iavorov Kirilov

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Ph.D.
2012

Advisory Committee:
Professor Georges Rey, Chair
Peter Carruthers
Erin Eaker
Allen Stairs
Howard Lasnik

Preface

This dissertation examines whether and how Fodor's asymmetric dependency account of intentionality (ADA) can be used to ground the contents of representations in the "early visual system". Before continuing, it is important to define exactly what "early visual system" means here in order to specify the type of representations that are going to be the focus in this dissertation:

I have adopted Pylyshyn's usage of "early visual system" in this dissertation which, with few exceptions[1], corresponds to the general use of the term. Before expounding on the main traits of the definition, it is important to note that Pylyshyn defines "early visual system" functionally, as opposed to neuroanatomically. One of the results of defining it this way is that the computations of the early visual system are not confined to those of the primary visual cortex, as is sometimes thought[2]. The definition that Pylyshyn (1999) provides goes as follows[3]:

---

[1] For example, according to Pylyshyn's usage, focal attention is outside (and prior) to the early visual system, which departs from the way the term is used in neuropsychology (*cf.* Pylyshyn (1999, p. 3, ff. 2)).

[2] *Cf.* Julesz (1991, p. 740, *italics* mine): "Conceptually defined, "early vision" should be identical to pure bottom-up visual processes depicted without being influenced by the top-down stream of semantic information. *Neurophysiologically defined, "early vision" should correspond to the first neural processing stages in the retina and the visual cortex.* Psychologically defined, "early vision" should encompass a range of perceptual phenomena that can be experienced by humans in the absence of higher cognitive and semantic cues."

[3] The following is a condensed version of the definition Pylyshyn provides in Pylyshyn (1999, p. 6-7).

1. A proper part of the process of visual perception involves a uniquely visual system called “early vision”.

2. The early vision system involves the computation of most specifically-visual properties, including 3D shape descriptions.

3. The early vision system carries out complex computations, many of which involve top-down processing[4]. That is to say that the interpretation of parts of a stimulus may depend upon the interpretation of other parts of the stimulus, resulting in global-to-local influences.

4. The early vision system is encapsulated from cognition: it is cognitively impenetrable. That is to say that it cannot access relevant expectations, knowledge, and utilities in determining the function it computes.

5. Cognition intervenes in visual perception either by allocating attention to locations or properties *prior* to the operation of early vision, or by recognizing and identifying patterns *after* the operation of early vision. It does not interfere *in* the operations of early vision.

[4] What “top-down” means for Pylyshyn is different from the meaning of the term in other authors’ writing (*cf*. ff. 2 and Julesz (1991, p. 740)). For Pylyshyn (1999, p. 6) top-down influences within early vision are influences of visual interpretations computed by early vision on other visual interpretations, separated either by space or time. What authors like Julesz (1991) mean by top-down process is closer to what Pylyshyn (1999, p. 6) calls “cognitive penetration”: Cognitive penetration originates outside the visual system and affects the content of visual perception in a meaning-dependent way. When a system is cognitively penetrable, then the function it computes is sensitive to the organism’s goals and beliefs.

Based upon this definition, the focus of this dissertation is going to be on representations produced by an autonomous visual system whose processes are not affected by background beliefs. The structure of these processes is innate—not influenced by the environment. Thus, how the representations in question get their content is not dependent on background knowledge, but on innate contstraints (what are called "natural constraints"[5]) that are built in the visual system. Thus, the goal of this dissertation is to figure out what some of the necessary constraints need to be in order for a representation to refer to a particular object.

I have chosen to examine how Fodor's account can help determine what some of the necessary constraints need to be in order for a representation to refer to a particular object. Fodor's account is chosen not because it is free of problems, but because it has the structure of a theory that promises to deal with many of the classic problems that befall informational semantics.

[5] *Cf.* Pylyshyn (1999, p. 28): "Embodying a natural constraint is different from drawing an inference from knowledge of the world (including knowledge of the particular constraint in question) in a number of ways. (a) A natural constraints that is embodied in early vision does not apply and is not available to any processes outside of the visual system (e.g., it does not in any way inform the cognitive system)…(b) Early vision does not respond to any other kind of knowledge or new information related to these constraints (e.g., the constraints show up even if the observer knows that there are conditions in a certain scene that render them invalid in that particular case)."

Applying Fodor's account and modifications of it to the representations in the early visual system allows me to take a position with regards to the question in psychology as to what is required for representation of particulars in the early visual system. In particular, I am concerned with Zenon Pylyshyn's claim that visual indices (FINSTs) refer unmediated by conceptual representations of location to objects in the external world.[6] In contrast, I argue that representations of particulars are mediated by conceptual representations of spatial relations. Furthermore, I argue that the capacity to express definite description is necessary to represent particular objects.

Pylyshyn's view presupposes a strategy for accounting for reference where tokens of symbols represent the particulars that caused them. As I explain in Chapter 1, such strategies have trouble specifying the link, in a causal chain connecting a particular with a token symbol, that is the symbol's referent[7]. Furthermore, such strategies, by themselves, have difficulty accounting for how properties are represented: allowing that they account for representations of particulars, these strategies need to provide a principled manner of generalizing from representations of particulars to representations of properties. It is not clear what this principled manner would be in a strategy that is built around *only* causal relations between particulars. This suggests that such strategies be modified—e.g. by adding teleosemantic features or by also appealing to causal relations between properties. I argue in Chapter 1 that the end result is likely to be at least as complicated account as the one that I will

[6] *Cf.* Pylyshyn (1988, 2001, 2007).

[7] This is known as the "which link" problem.

pursue in this dissertation: a strategy for accounting for content built around nomic causal relations between properties. This provides a reason to explore the latter strategy as well, which I do in this dissertation.

In the chapters that follow I explore whether and how one instance of this strategy—Fodor's asymmetric dependency account (ADA)—deals with one of the main problems raised against it (its inability to explain how particular objects are represented) and whether and how it can help determining what is required to represent particular objects in the early visual system. The overall argument in the dissertation can be informally summarized as follows:

1. ADA is a promising account (Chapters 1-2)
2. ADA cannot account for how representations of particular objects get their content (Chapters 2-3)
3. Buttressing ADA with a conceptual role semantics is a promising strategy for solving the problem in (2) (Chapter 3)
4. Conceptual role semantics is not as hopeless as some (e.g. Fodor) think it is in dealing with its main objections (Chapter 4)
5. The account in (3) claims that definite descriptions are required to represent particular objects.
6. *If* the ability to form definite descriptions presupposes conceptual representations, then the account in (3) is contrary to Pylyshyn's claim that representations of objects are not mediated by conceptual representations of features at locations

In Chapter 2 I present ADA, the motivations for it, some of the main changes it has gone through, and the problem of representing particular objects that it faces. Fodor designed ADA as a type of informational semantics account because of his skepticism about conceptual role semantics—the account according to which the content of a symbol is defined in terms of its inferential relations with other symbols. This account Fodor takes to fatally suffer from the holism problem—the charge that conceptual role semantics spells doom for psychology as a science aiming to provide generalizations over people's mental states. The holism charge is that a) one cannot identify a subset of a symbol's inferential relations as the meaning-constitutive ones in a non *ad hoc* manner[8]; b) a symbol's content would then have to be defined through *all* of a symbol's inferential relations which would spell doom for psychology. This bleak scenario would follow because no two people with different minds would share the same symbols with the same contents. This means that psychology would fail to provide laws and generalizations. Since psychology *does*

[8] Fodor takes this for granted because he is persuaded by Quine's attack on the analytic/synthetic distinction in "Two Dogmas of Empiricism". Analytic statements are true in virtue of the meaning of their constituent concepts. They can be used to define the meaning of a term. Synthetic statements are true in virtue of how the world is. The analogue of analytic statements in conceptual role semantics is a proper subset of a symbol's inferential roles—the meaning-constitutive ones. The analogue of the synthetic statements is all the other inferential roles. But if the analytic/synthetic distinction is bankrupt, then conceptual role semantics would not be able to appeal to the analogous distinction between a symbol's meaning-constitutive inferential roles and the rest. This would make it a non-starter.

provide laws and generalizations, conceptual role semantics must be mistaken and so a different approach is needed—an informational semantics one.

Informational semantics, however, suffers from the disjunction problem—the problem of determining which one of the many causes of a symbol is its referent. Fodor's solution is ADA: according to ADA a symbol refers to the property with which it enters in a nomic relation on which all the other nomic relations the symbol enters in depend, but not *vice versa*.[9] Fodor stresses that ADA is *not* a theory of mentality.[10] It is an account of how representations of *properties* get to have the intentional content they do. Importantly, it is not an account of how representations of *particular objects* get their content. The reason is that *nomic* relations exist *only* between properties and not between particular objects. Fodor offers a different and unrelated account—his triangulation account—to explain how representations of particulars refer. One could try to simply clump the two accounts together, but one would need to provide an explanation that does not presume what the content of a given symbol is, as to why the content of that symbol should be accounted for by one and not the other of the accounts. Examining whether this approach could work I leave for another time. In this dissertation I will explore what seems to be a simpler

---

[9] *Cf.* Fodor (1987, 1990).

[10] *Cf.* Fodor (1990, p. 130): "Even if it's true that intentionality equals information plus robustness [(ADA)], it wouldn't have to follow that information plus robustness is sufficient for *mentality*. Sufficient conditions for being in a state with intentional content needn't also be sufficient conditions for having a belief or a desire or, indeed, for being in any other *psychological* condition".

approach: whether and how ADA can be used to ground the contents of representations not only of properties but of particulars as well.

In Chapter 3 I argue that ADA, when supplemented with conceptual role semantics, can account for how representations of particular objects refer. The idea is to use conceptual role semantics to account for the logical and non-logical roles of representations of particular objects in the visual system. The relations of particular interest are those between the representations of particular objects and representations of spatio-temporal relations. The representations of super-determinate spatio-temporal relations that I appeal to—$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$—co-vary, in the way Fodor suggests, with the property of being a spatio-temporal relation at time t such that: a) it involves the objects x, l, n, and r forming a pyramid with sides of length d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx}; and b) r, l, and n form the base of the pyramid and r and l are my eyes, while n is my nose:

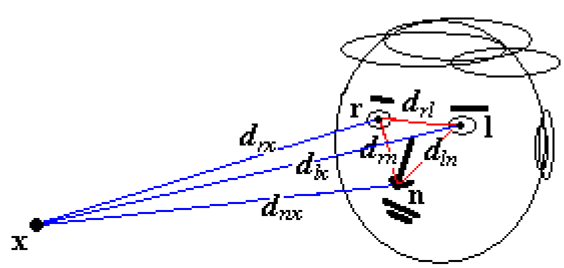

Figure 0.1.: This is an instance of a configuration involving four objects x, l, n, and r forming a pyramid with sides of length d_{rx}, d_{lx}, and d_{nx} and a base of a triangle with sides d_{rl}, d_{rn}, and d_{ln}. The configuration instance is instantiated so that l, r, and n are the left, right eye, and the nose of "Pinocchio" here.

From now on, for brevity, I will use "'C_j" for the predicate "'$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$'".

Importantly, the content of the predicate 'C_j' is not defined entirely *á la* ADA, but is inferentially and functionally related to a first-person reflexive term ('I') and an indexical for time ('t'). It is this characteristic that complements and helps ADA here: in virtue of the predicate being related to 'I' and 't' it doesn't pick out just any pyramidal super-determinate configuration of type C, instead it picks out a particular instance of it: *the* one that is set on me at t. This makes the referent of the predicate 'C_j' a trope.[11] 'C_j' is nomically related to the property of being a pyramidal super-determinate configuration but is *not* nomically related to a trope. Instead, the trope is *picked out* in virtue of the 'C_j''s nomic relations with the property of being the super-determinate spatial configuration C_j *and* in virtue of its inferential and functional relations with 'I' and 't'. The picking out of a trope is cashed out as the satisfaction relation that obtains between Fregean definite descriptions and their referents. In short: when 'I' and 't' in the predicate '$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, O)$' are saturated, then the referent of this predicate is a trope.

[11] Tropes, also known as property instances, concrete properties, and abstract particulars can be schematized as follows (*cf*. Funkhouser (2006, p. 14)): "[(O, *t*), P]. This is to be read as "O's having/being P at *t*." O is an object or spatial location, *t* a time (span), and P a property type. Such property instances occur if, and only if, O *really is* P at *t*."

I suggest that we use the predicate ‘$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, O)$’ in the following definite description in virtue of which the symbol ‘x’ gets to refer to a particular object:

$$(\boldsymbol{E}x)((y)(C(y, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{ry}, d_{ly}, d_{ny}, t, O) \leftrightarrow (x = y)))$$

Put crudely, this reads: right now x is the one and only particular that has the relational spatial property of being at the specified distances with respect to my nose and eyes.

This approach makes use of conceptual role semantics for more than the logical apparatus needed for the definite description. Conceptual role semantics is also used to define the content of ‘I’ (the first-person reflexive term) and is perhaps needed also for the content of the temporal predicate ‘t’.

Since there are serious problems against conceptual role semantics accounts, I devote Chapter 4 to meeting the key objections. Following up on Rey’s comparison of Horwich’s and Fodor’s accounts, I suggest that by taking the asymmetric dependencies inside the mind (i.e. by applying them to the roles between symbols), one can identify the meaning-constitutive roles of a symbol: the meaning-constitutive ones are the ones on which all the rest depend, but not *vice versa*. This would allow one to solve the analytic/synthetic distinction problem and because of this it would allow one to avoid the holism charge levied against conceptual role semantics.

After defending conceptual role semantics, I incorporate it in an informational semantics account *á la* Horwich's basicality strategy for accounting for the content of symbols. According to this strategy the content of a symbol is determined by the relation in which a symbol enters that is explanatorily basic: the one on which all other relations involving the symbol depend on, but not *vice versa*. The basic relations in question can be external, internal, or a mixture of both. For example, many primitive symbols in the sensory modalities will have a referent that is accounted for entirely through external (nomic) relations. Logical operators would be accounted for entirely through internal (inferential) roles. Empty names and predicates (e.g. 'Zeus', 'miracle', and 'triangle') and complex predicates might have their content accounted for through both external and internal relations. The content of representations of particulars would be accounted for through internal (inferential) relations with symbols whose content is in turn accounted for at least in part through external (nomic) relations. Thus, the content of representations of particulars is accounted for through both nomic and inferential relations.[12]

[12] Note that this leaves open whether the content of representations of particulars *is*, in part, the inferential relations. That is, it is left open whether representations of particulars have narrow content. What is said in the main text only states that the internal relations play a role in determining the content of representations of particulars, but something can determine the content of 'x' without *being* the content of 'x' (e.g. Kaplan's definite descriptions that individuate a 'dthat''s content). Logical operators, in contrast, have inferential roles that *are* their content.

In Chapter 5 I apply the account of how particulars get represented to the problem of representing particular objects in the early visual system: according to the account above of how particular objects get represented a representation of a particular object gets its content in virtue of definite descriptions. This means that the possession of quantifiers, predicates, argument variables, and names is required to represent particular objects. If the capacity to express definite descriptions is taken to imply the possession of conceptual representations (these being the predicates and argument variables in the definite descriptions), then Pylyshyn would be mistaken to hold that FINSTs are not mediated *via* conceptual representations.

Dedication

To One Tireless Advisor

Table of Contents

List of Figures

Chapter 1: Nomic and Singular Causal Strategies for accounting for content

1.A. Introduction:

In this chapter I explain why nomic accounts of content like ADA are chosen to begin with. Consider the representations in the early visual system: they are visual indices (demonstratives) and representations of properties. Strategies for accounting for content centered on causal relations between symbols and *particular objects*—what I call "singular causal strategies of content"—are the most obvious approach to take in explaining how visual indices refer. Strategies for accounting for content centered on causal relations between symbols and *properties*—what I call "nomic strategies of content"—are the most obvious approach to take in explaining how representations of properties refer. In what follows I argue that singular causal strategies of content face a significant problem—the "which link" problem. This problem constitutes a reason to examine how promising nomic strategies of reference are in explaining how representations of particulars refer. In the following chapters I will explore how much this promise gets fulfilled.

The plan for this chapter is as follows: Section 1.B is where I draw the distinction between singular causal strategies of content and nomic strategies of content. In Section 1.C I analyze the most significant problem singular causal strategies of content face—the "which link" problem. The implication is that this problem is a reason to explore other strategies: specifically, nomic strategies of content. In Section 1.D I argue that even ignoring this problem, singular causal strategies of content end up being about as complicated as the nomic strategy that will be examined in this dissertation. Again, the implication is that this is a reason to explore the nomic strategy as well, as I do in this dissertation.

1.B. Causal strategies for accounting for content:

In this section I distinguish between singular causal strategies of content and nomic strategies of content. I ignore teleosemantic strategies of content, even though they also fall under the type of causal strategies of content[13]. The reason is that they, like ADA, have trouble accounting for how representations of particulars refer—it is extremely implausible that I have a mechanism whose function is to represent Bob, for example. After all, Bob wouldn't have been around when the mechanism would

[13] *Cf*. Neander (1995) where she argues that teleosemanticists should be concerned only with those properties that *cause* an underlying genotype to be selected: "According to the definition favoured in this paper, the function of a trait is to do whatever it did in ancestral creatures that was adaptive and *caused* the underlying genotype to be selected" Neander (1995, p. 114, *italics* mine).

have been selected for. Thus, teleosemantic theories, like ADA, need to be supplemented in order to account for representations of particulars. In fact, if ADA is mistaken and teleosemantic theories are right, then the solution I propose in Chapter 3 will still work after ADA is replaced with a teleosemantic theory. As a result, I will not argue about how ADA and teleosemantic theories match up[14]. Instead, I will concern myself with causal theories that, *prima facie* at least, show promise in accounting on their own for how representations of particulars refer.

In what follows I distinguish between strategies of content based on causal relations between *particular objects* and strategies based on causal relations between *properties*. But causal relations exist between particulars *and* between properties. Why center a causal strategy for accounting for the content of representations on relations involving *only* particular objects or on relations involving *only* properties? Why not have a strategy where if a symbol is caused by a property, then it represents the property, and if it is caused by a particular object, then it represents the object? The reason stems from the fact that the causal relations between properties are generally taken to be nomic relations—lawful co-variation between properties[15]. As will be pointed out in the next chapter, I will assume, following Fodor, that such relations "cover" causal relations between particular objects. This means that when we have causal relations between particular objects we also have causal relations

[14] I believe that ADA is better, but discussion of that is left for another time. *Cf.* Rey (1997) for an excellent survey of the main problems facing teleosemantic theories.

[15] *Cf.* Armstrong (1997).

between properties, "covering" the causal relations between the objects. The same applies in the opposite direction if one assumes, as I will do in this dissertation, that no property can exist uninstantiated by a particular object.[16] This means that whenever a symbol is tokened: a) the token of the symbol is caused by a particular object; and b) there is a law that relates properties of the token with properties of the particular object. It is then *ad hoc* to claim that with some symbols it is the particular object which causes their tokens, that is their content, while with other symbols it is the property of the particulars which cause their tokens, that is their content. Thus, causal strategies should be split into two types: singular causal strategies[17] of content which focus on causal relations between particular objects and nomic strategies of content which focus on causal relations between properties.

The prototypical example of singular causal strategies of content is Kripke's (1980) account. The following illustrates Kripke's approach: when a baby is born, one points at the baby and names it 'Bob'. Whoever uses 'Bob' to refer to the baby is connected by a causal chain that runs through communications with people who use 'Bob' to refer to the baby and understandings to use 'Bob' in the same way, all the way through the point when one pointed at the baby an baptized it 'Bob'. Now, this

---

[16] *Cf.* Armstrong (1997) for a persuasive defense of this position.

[17] I have borrowed the term "singular causal relation" from Armstrong (1997) where he uses it to distinguish causal relations between particulars from causal relations between properties (the latter specify laws). *Pace* nominalists, I take this distinction to be plausible. At the very least, the distinction does not seem to be as controversial as the expulsion of universals from the ontological universe.

approach is merely a strategy and is not meant by Kripke to be a theory of content.[18] As Kripke (1980, p. 97) points out, it appeals to intentional terms like "pointing", "naming", "communicating", and "understanding" which means that intentionality has not been reduced.[19] In what follows, I argue that there are substantial problems for trying to account for the content of representations of particular objects *only* through non-intentionally specified causal relations between particular objects.

The paradigm example of nomic theories of content is Dretske (1982). Dretske's account and its problems is what gave rise to ADA. On Dretske's account a symbol's content is what the symbol carries information about. What a symbol carries information about is what the symbol enters in a lawful (nomic) causal co-relation with. Dretske argues[20] that there is a learning period during which a symbol is correlated with a certain cause. After the learning period is over, other things than the cause can trigger the tokening of the symbol. Since they would not be what triggered the symbol during the learning period, the symbol would misrepresent them, even if it carries information about them.

Before I continue, I want to point out that the reason it is ADA, rather than Dretske's account, that was chosen for this dissertation is that I was persuaded by

---

[18] Devitt & Sterelny (1999) is an example of someone who attempts to give a singular causal *theory* of content.

[19] *Cf*. Kripke (1980, p. 97): "[my account] takes the notion of intending to use the same reference as a given".

[20] *Cf*. Dretske (1982, p. 193).

Fodor's criticisms of Dretske's account: Fodor (1990, p. 41) roundly criticizes Dretske's approach because a) the learning period is specified in an *ad hoc* manner; b) Dretske's account applies only to *learned* symbols; and c) it ignores the relevant counterfactuals. The point in (c) is that Dretske's account assumes that symbols carry information about their referents. That is, it assumes that there are nomic relations between the two. But nomic relations are counter-factual supporting. This means that they cover not only what actually happens at time t, but also what *would* have happened at time t, *if* circumstances were different. Thus, even if at time t a cat caused 'cat', if the light were weaker and small dogs were around, it would have happened that a small dog would have caused 'cat' at time t. Thus, 'cat' enters in nomic relations with (and therefore carries information about) both cats and small dogs at dark. If this is the case, then what *actually* happens at time t (or in the learning period) is not as important as what was nomologically possible to happen if circumstances were different. Thus, Dretske's account does not explain why 'cat' refers to cats as opposed to small dogs in the dark.

1.C. Representing particulars in singular causal strategies:

In this section I analyze what I take to be the main problem for singular causal strategies of content: the "which link" problem. I discuss a solution proposed by Fodor (2008)—triangulation—and argue that it fails to solve the problem.

On singular causal strategies of content a symbol's referent is the particular object which, through a causal chain, causes a token of the symbol. This is the most obvious way to account for representations of particular objects. However, it also faces an equally obvious problem: what has been called the "which link" problem. Suppose that a token of a symbol is linked with a causal chain to a particular object. The causal relation between the two is, in most cases, a complex one. It involves many intermediate causal relations. In other words, in most cases, there are many intermediate links each of which bears a causal relation (direct or indirect) with a token of the symbol. That makes every link be a particular object which, through a causal chain, causes a token of the symbol. The "which link" problem is the problem of individuating a particular link as the symbol's content.

Fodor (2008) presents a way of solving the "which link" problem—through triangulation:

> Imagine there is not just the actual Adam with the perspective that he actually has, but also a counterfactual Adam ('Adam2') who is, say, three feet to the actual Adam's right. Adam2 has a (counterfactual) perspective on the (actual) visual scene; one that differs from Adam's perspective in accordance with the usual (i.e. the actual) laws of parallax. Assume that Adam2 tokens a representation of the same type that Adam does. Draw a line that starts at Adam2's token and represents its causal history (i.e. the causal history that Adam's token would have had if Adam had been at the position that Adam2 occupies in the counterfactual scenario). Call this Adam2's line. The metaphysical problem is:

> given the two causal histories, solve for the referents of the tokens. RTM allows us to do so. It says that the two tokens have the same referent iff Adam's line and Adam2's line intersect at a link; and that their referent is the link at which they intersect. (Fodor (2008, p. 213))

According to a charitable interpretation of this account, the link in the causal chain that is the referent of a token symbol is the *first* intersection of the causal chain with a counterfactual causal chain that leads to the same token symbol, but from slightly different counterfactual perspective.[21] This proposal has three key elements: a) intersection of causal and counterfactual chains; b) the counterfactual chain has a slightly different perspective; and c) the referent is the first intersection. The intersection of the actual and the counterfactual chains is a causal link at a spatio-temporal location. The causal link is shared by both chains. According to the triangulation account we move the observer and as a result, the actual and the counterfactual chains pass through different space-time points up to the space-time

[21] As I pointed out in Section 1.B, there are two types of causal relations: nomic ones (between properties) and singular ones (between particular objects). Only the latter ones can intersect in a spatio-temporal manner (which is what is required to individuate a particular object) because only they occupy spatio-temporal regions. They occupy spatio-temporal regions because they involve specific individuals. For example, a causal relation between two particulars is located in the line segment where the particulars are at the ends. The nomic relations on the other hand, being relations between properties, while instantiated in particulars, are not located anywhere—just as properties are not located in a specific place (this, of course, assumes that properties are different from classes of particulars—i.e. they are more than the members in their extension). Not being located anywhere trivially denies them the ability to intersect in a spatio-temporal manner with each other.

location of the link where they intersect and merge (they continue merged back to the Big Bang). The referent is not the Big Bang because the referent is the first intersection and the Big Bang is a later intersection.

Now I will present three arguments as to why the triangulation account, under this intuitive interpretation, cannot work:

1.C.I. Interference effects:

The first problem for Fodor's triangulation account is that the first intersection need not be the one we would intuitively label as the "right" one. We get such cases when the causal chains are instantiated by light beams. The reason we get them is because of a *fundamental* property of light—it interferes with itself. The upshot, as I explain below, is that the first intersection of the actual and the counterfactual causal chains is the link where the interference happens and that is not the intuitively "right" referent. In more detail:

Suppose that a laser is directed at a transducer[22] and that all the photons from the laser end up at the transducer. Suppose that there is no other light present and

---

[22] A transducer is a detector of the sensory system that maps ambient energy onto mental representations. *Cf*. Fodor (2008, p. 187): "Computation (like, for example, thinking) takes mental representations onto other mental representations. Transduction (like, for example, registering impinging redness) takes ambient energy onto mental representations. In the usual case (barring

finally suppose that the laser is instantiating a causal chain that *via* the transducer leads up to a token of a symbol ('x').[23] A problem emerges once we consider lasers and interference effects:

If a laser shines upon a barrier with a pair of slits upon it, then an interference pattern obtains:

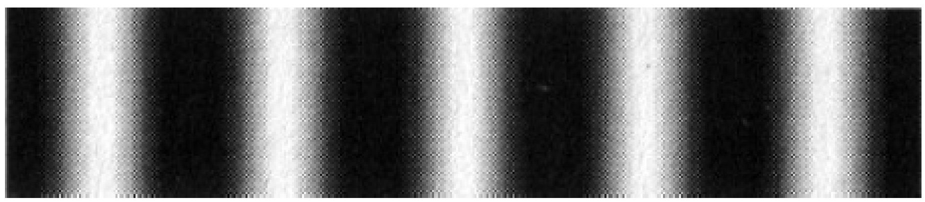

Figure 1.1.: This is the interference pattern that obtains when a laser beam is shone on a barrier with two parallel slits on it.

---

hallucinations and such) perceptual integration starts with the registration of sensory information. Were there no transducers, perception couldn't get started."

[23] Were the transducer to be moved, the actual causal chain would no longer lead up to the token 'x'. Moreover, were it the case that only the transducer was moved and nothing else changed, then there would no longer be any (counterfactual) causal chain leading to a (counterfactual) token of 'x'. This might seem to suggest that the triangulation account would not work here, because in the counterfactual conditions we would lack the counterfactual causal chain necessary for the triangulation. But the obvious solution is to counterfactually change a bit more than merely the position of the transducer: we can counterfactually change also the direction of the laser. The idea is that, whatever the source of the laser, we can assume that in the counterfactual circumstances the laser shines upon the transducer at its new (counterfactual) location. Then, whatever the source of the laser is, the actual and the counterfactual chains will intersect at the source.

The same interference pattern obtains (over time) even if one photon per day passes through the slits. There have to be at least two slits for the pattern to obtain and the pattern does not obtain if one slit is closed. This observation suggests that there is something that interferes with the photons so as to cause them to change direction once they are out of the slits.[24] So we have a picture like this:

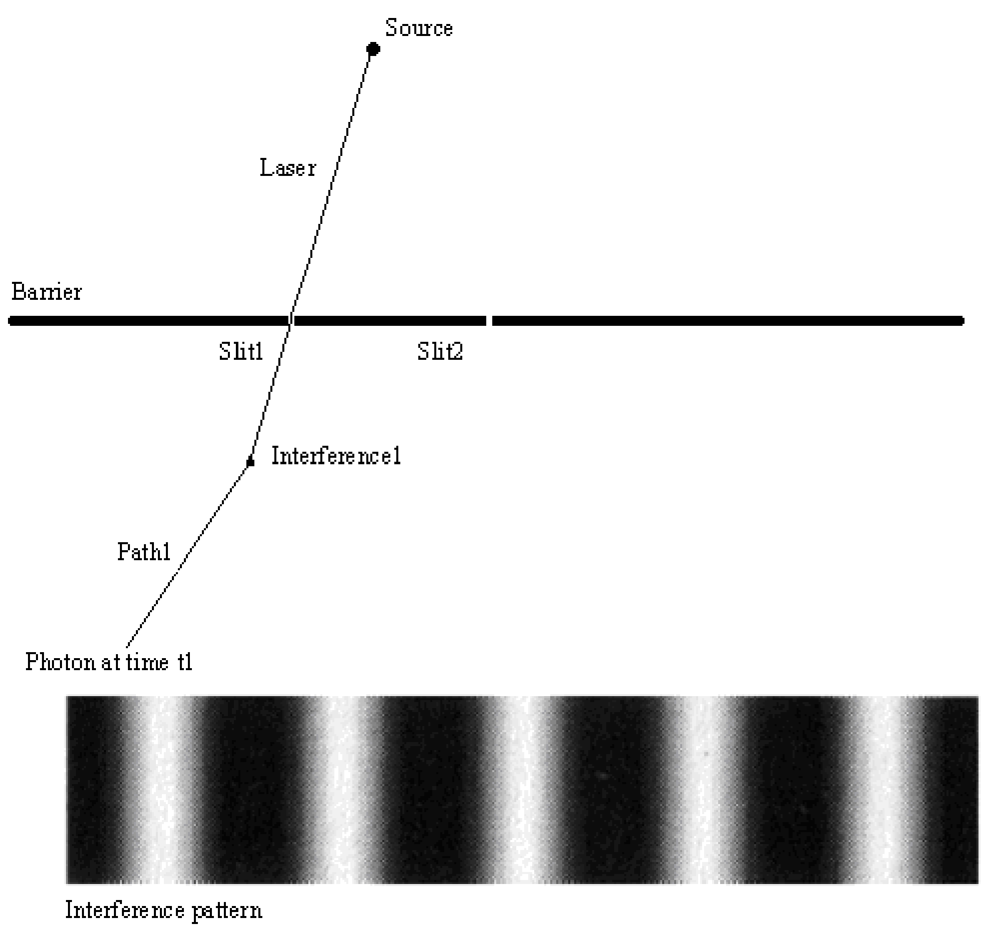

[24] *Cf*. Deutsch (1997, p. 41-42).

Figure 1.2.: Schemata of a source emitting a laser, photons of which pass (one per day) through one of the slits in a barrier. Since a second slit is open, there are interferences, which explain the interference pattern.

Now consider the following diagram:

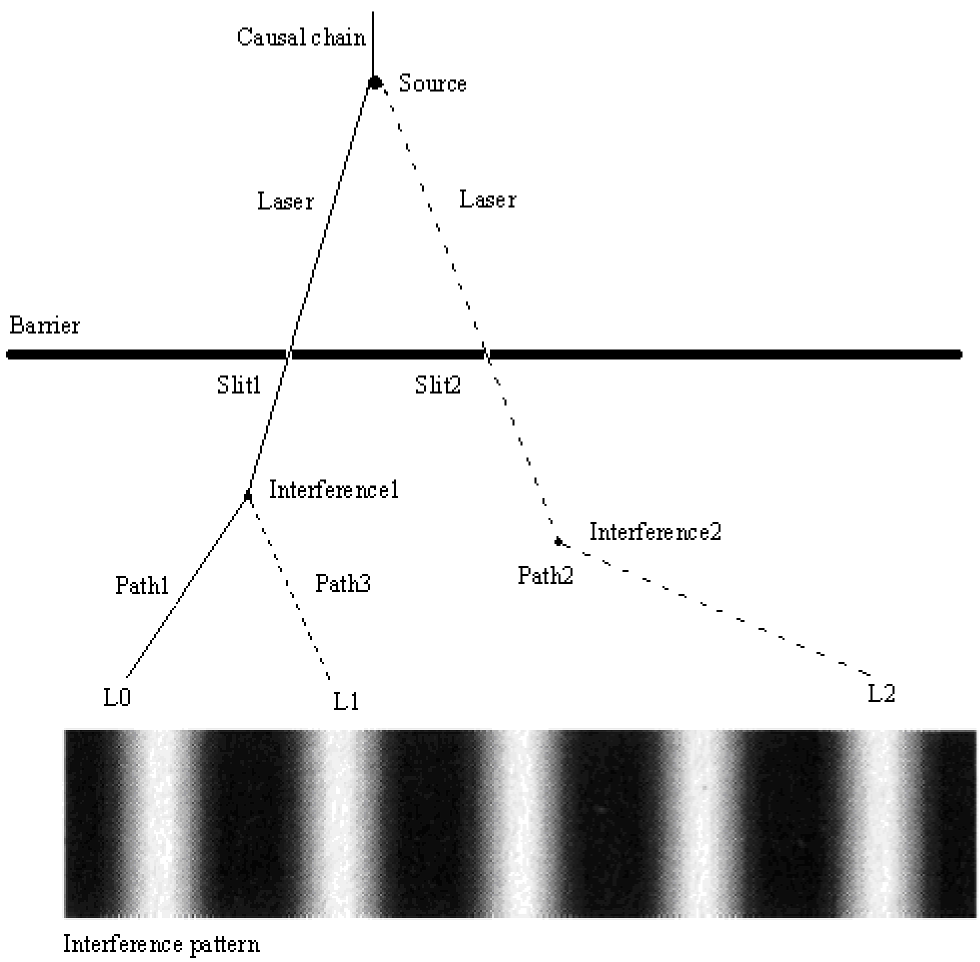

Figure 1.3.: Schemata of a source emitting a laser, photons of which pass (one per day) through one of the slits in a barrier. Since a second slit is open, there are interferences, which explain the interference pattern. The striped lines denote the

counterfactual paths of photons. L0 denotes the actual location of a transducer, while L1 and L2 denote two counterfactual locations of the transducer.

Suppose that a transducer is placed at L0—in the path of a photon that has been interfered with. Let's apply the triangulation account here: if the transducer were to be placed at L2 and if a photon were to travel down that path, then we would have an actual and a counterfactual causal chains (they would be bent a little, but that matters not) that intersect at Source, which is the answer that we intuitively want. So far, so good. However, suppose the transducer were to be moved in location L1. Suppose that in this counterfactual scenario the photon that moved along Path1 now moves along Path3. Then, the two causal chains would intersect at the point of interference (Interference1) and so Interference1 would be the referent of 'x'. First of all this seems wrong. Furthermore, another problem is that depending on where we choose to place the transducer, we get two different intersections of the causal and the counterfactual chains: Source and Interference1. This means that the triangulation account does not determine a unique referent of 'x'.

Note that while I assumed that the point of interference for the photon that passed through Slit1 stays the same, we don't need to hold the point of interference constant: in the counterfactual scenario the point of interference might be at a different location than in the actual scenario—say, several centimeters after or before the actual interference point. The conclusion above still follows as long as the photon in the counterfactual scenario travels along the same path before the slit: the actual and the counterfactual chains will still intersect prior to Source and this will still give us a different intersection than Source:

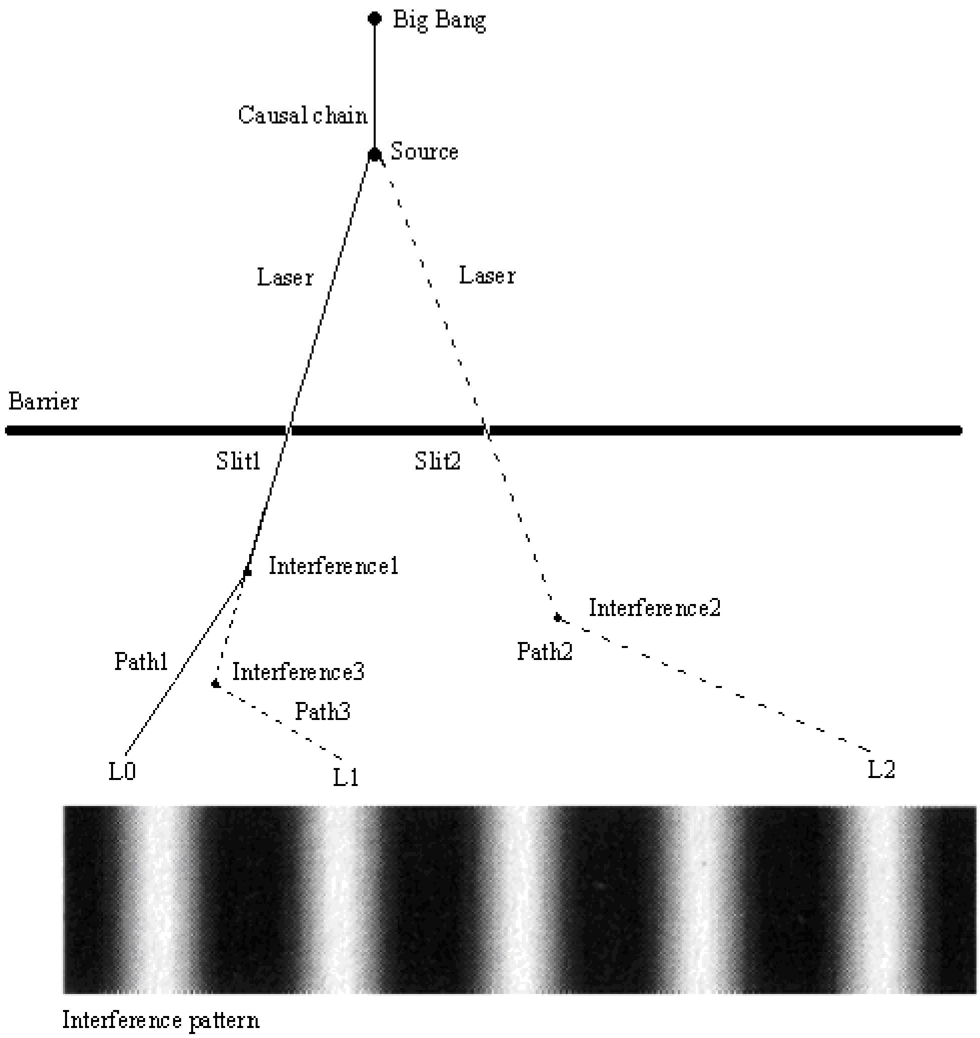

Figure 1.4.: Schemata of a source emitting a laser, photons of which pass (one per day) through one of the slits in a barrier. Since a second slit is open, there are interferences, which explain the interference pattern. The striped lines denote the counterfactual paths of photons. L0 denotes the actual location of a transducer, while L1 and L2 denote two counterfactual locations of the transducer.

To avoid the above problem, one might try to postulate that the counterfactual chain travels along a different path from the actual one, excepting the path after the first intersection. They will merge at the referent of ‘x’ and continue as a single causal chain back to the Big Bang. This would *seem* to eliminate the counterfactual chain that starts at Source and continues to L1 because it travels on part of the path of the actual chain that starts at Source and continues to L0. Then, purportedly, we would be left only with the counterfactual chain that starts at Source and continues to L2 for triangulation and the problem above would seem to be solved. The problem with this approach, however, is that Path1 and Path3 satisfy the above constraint, but still produce an intersection other than Source—at Interference1:

Consider the fork that starts at Source and has prongs with tips at L0 and L1. These prongs split at Interference1. We have another fork—one that starts at the Big Bang and has prongs with tips L0 and L2. These prongs split at Source. According to the above constraint, the counterfactual chain must travel along a different path from the actual one excepting the “handle” of the fork (from Source to Big Bang). Since the counterfactual chain that ends up at L1 travels part of the route of the actual chain, it should be dismissed. But if the “handle” of the fork (from Source to Big Bang) is allowed to be an exception, then there is nothing wrong with taking the route from Source to Interference to be a “handle” and treat it as an exception, as well. Then, the prongs with tips at L0 and L1 travel along different routes (Path1 and Path3), excepting the path after the first intersection, and satisfy the above constraint.

One could object that the problem above assumes that the "right" link can be a link at a point of light interference—a *microscopic* event. However, the objection continues, our vision is not designed to detect microscopic events but *macroscopic* ones. This would seem to deal away with the problem above. However, it is not obvious that eyesight is not designed to detect microscopic events. Frogs eyesight is so sensitive that they can detect single photons[25]. More importantly, it is possible for systems to refer to particular objects at a level below the macroscopic one. The move above would leave such systems unexplained.

1.C.II. Single-slit diffraction:

A second problem with the triangulation account can be seen once we consider single-slit diffraction. When light passes through a single slit it can be observed to spread out (or "fray")[26]:

[25] *Cf*. Deutsch (1997, p. 34) and Donner (1989).

[26] *Cf*. Deutsch (1997, p. 38-39).

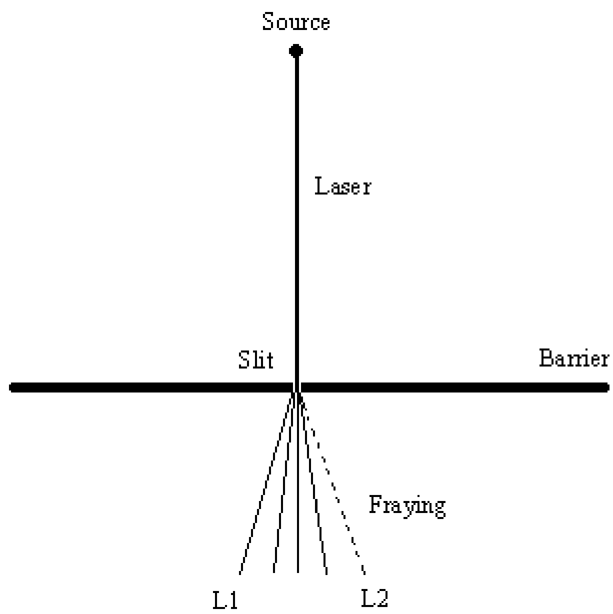

Figure 1.5.: Schemata of a source emitting a laser, photons of which pass through a slit in a barrier. The striped line denotes a counterfactual path of a photon. L1 is the location of the actual observer. L2 is the counterfactual location of the observer.

Now, consider an observer at L1 and let's apply the triangulation account: were she to move, say to L2, then she would be in the path of a fraying light beam that intersects the actual light beam at the location of the slit. The actual and the counterfactual light beams intersect and converge at Slit. By the triangulation account then, it is the photons at the location of Slit which are the referent of a token symbol 'x'. However, as Cole points out[27], "we can refer to things further back on the causal

[27] *Cf.* Cole (2009, p. 442). In his thought experiment he uses sound waves passing through a key hole, but the point is the same: "[W]e can refer to things further back on the causal chain than any perceptual

chain than any perceptual intersection". That is, if I am standing at L1 and observing the light beam, I am tokening 'x' and using it to refer not to the location of Slit, but to a point behind it. Thus, the triangulation account gives the wrong result.

One could object that the case above is a "weird" case: Fodor's triangulation account is meant to be an account that explains how, *ceteris paribus*, we refer to particular objects. Cases of occlusion, like the one above, are to be dismissed *via* the *ceteris paribus* clause. A response to this objection is as follows: *ceteris paribus* conditions are there to weed out cases where it should be expected for a theory not to apply. Perhaps in the case of light fraying, we can grant that it is a weird case of representing particular objects: the hole has to be very small so that we can talk about light beams actually intersecting at it. Such scenarios are not common enough to need to be covered by a theory of how particular objects are represented. However, as in Cole's example, cases where we *hear* people through holes in closed doors are surely not uncommon. Sound waves passing through holes as big as keyholes *spread out* and causal chains instantiated by sound waves can be said to intersect at such holes. Given that this is not an uncommon scenario, I take it that it would be a scenario that we would like an account of how particular objects are represented to explain. Thus, it would be a weakness for the triangulation account to try to sweep it under the *ceteris paribus* carpet.[28]

---

intersection. Suppose e.g. you and I listen to a conversation through a keyhole. We can be thinking about the speakers, even though our perceptual causal chains intersect at the keyhole."

[28] Another problem raised by Cole is the "cat assassin" problem (Cole (2009, p. 442)). The essence of the problem is that the counterfactual observer might not see what the actual observer sees. Thus,

1.C.III. Objects, tropes, and states of affairs:

The triangulation account is an account of how particular objects get represented. It takes for granted that only objects enter in singular causal relations. As a result, when it individuates a singular causal link, the link can only be an object. Now, objects enter in singular causal relations *in virtue* of the properties they have. That is the same as saying that nomic relations between properties "cover" the singular causal relations that objects enter in. Yet another way to put the point is that properties play a causal role in singular causal relations between objects. So when we

---

suppose that I stand in front of a cat assassin (while looking at my cat) and so prevent him from killing my cat. Suppose that behind my cat is the mother of my cat. I am tokening 'x' as a result of a causal chain that goes through my cat, through my cat's mother, back to the Big Bang. Were I to move, however, then the cat assassin would vaporize my cat, and I won't see it, but will instead see its mother. Thus, the counterfactual chain that causes counterfactual-me to token 'x' would not contain my cat but my cat's mother. The actual and the counterfactual chain would then intersect and converge at the mother of my cat, which would make it, and not my cat, the referent of 'x'. This, as Cole points out, is the wrong result, because, "I am thinking about my cat, not [about] something that in nearby worlds would come into view." Furthermore, this shows that depending on how we choose the counterfactual scenario we get different results: my cat and my cat's mother. This means that the triangulation account does not determine a unique referent of 'x'. But one could avoid this problem by pointing out that, intuitively, we can restrict the counterfactual scenarios to those relevant for triangulation—the ones where I move, but the cat assassin does not strike. How "relevant for triangulation" is specified here is unclear, but the intuition seems to be that in triangulating we allow only what instantiates the causal chains to "move" and we "freeze" the rest of the world. This would leave the cat assassin "frozen" and would deal with Cole's example.

are individuating *via* triangulation *a* causal link in a singular causal chain, why are we specifying only an object and not also the properties that the object has? One answer is that we are individuating a particular *at* a space-time location. Objects can have space-time locations. Properties don't. But now, it seems that a lot hinges on what metaphysical view we are committed to:

Suppose that one denies that properties are universals and appeals instead to tropes—super-determinate properties instantiated at space-time locations. On this view, instead of a property being causally relevant to a singular causal relation that an object enters in, we have a trope, instantiated at the space-time location of the object, being causally relevant. Triangulation, then, triangulates two particular entities at the same space-time location—an object and a trope at the space-time location of the object. Both play a causal role in the causal chain leading to a token symbol 'x'. So, which one is the referent of 'x'? One can argue that since both are playing a causal role in the singular causal chain, both, together, should be the referent of 'x'. This is treating them as an event or, alternatively, a state of affairs (a particular object instantiating a property at a space-time location). But if one goes this way, then one has given an account of how states of affairs (or events) are represented, not how objects are. Note that the two are different, since objects, unlike states of affairs (and events), can move through space-time, while the latter can't. This means that *if* one makes the metaphysical moves above, then triangulation ceases to be the account of representation of particular objects that we are looking for. Granted, the metaphysical moves may be deemed "sketchy" due to the appeal to tropes. But such moves *are* defended in the literature by nominalists and, in any case, the point is merely that if the triangulation account is an account of how objects are represented, then it needs to be elaborated to prevent "sketchy" moves like the ones above being made.

Before I move on, let me briefly address another option here: instead of appealing to actual and counterfactual causal chains and their intersections—an ingenious but metaphysically problematic idea—appeal to intersections of actual causal chains. That is: take the causal chains leading to the eyes (technically we need three to triangulate an object in 3-D, but I will ignore this complication) and the first place where they intersect is the referent of 'x'. The proposal is intuitive, since our visual system is able to triangulate objects through convergence and information about which retinal receptors are activated.[29] However, the obvious problem here is that this account does not explain systems that can refer to objects without

[29] *Cf*. Palmer (1999, p. 205): "[A] source of information about [absolute] depth comes from eye convergence: the extent to which the two eyes are turned inward (toward each other) to fixate an object. The eyes fixate a given point in external space when both of them are aimed directly at the point so that light coming from it falls on the centers of both foveae simultaneously. Since each fovea has only one center, only one point can be precisely fixated at any moment. The crucial fact about convergence that provides information about fixation depth is that the angle formed by the two lines of sight varies systematically with the distance between the observer and the fixated point". *Cf*. Palmer (1999, p. 338): "An object's position on the retina relative to the center of the retina specifies its direction relative to the observer's direction of gaze because each retinal receptor signals light coming into the eye from one particular direction… information about retinal position is preserved in the early stages of visual processing via the many retinotopic maps that preserve relative position in the 2-D sheets of cortical cells."

triangulation through actual causal chains: imagine a system that has one eye and takes information about azimuth and altitude from the location of the eye within the eye orbit and integrates this information with information about distance from the ciliary muscles.[30, 31]

[30] Distance can be registered through sensors that register accommodation: *cf*. Palmer (1999, p. 203-4): "Accommodation is the process through which the ciliary muscles in the eye control the optical focus of the lens by temporarily changing its shape. It is a monocular depth cue because it is available from a single eye, even though it is also present when both eyes are used…[T]he lens of the human eye has a variable focusing capability, becoming thin to focus light from faraway objects on the retina and thick to focus light from nearby ones…. If the visual system has information about the tension of the muscles that control the lens's shape, then it has information about the distance to the focused object."

[31] Rey (p. c.) has pointed out to me that there may be a further problem with singular causal theories of content: If, as singular causal theories would have it, the demonstrative 'x' gets its content based on what causes it, then 'x' would be individuated by the particular entity that caused it. This means that for an indefinite number of causes of demonstratives, we need an indefinite number of demonstratives: for example, if we translate this to visual indexicals (FINSTs), this means that a FINST that is caused by $apple_1$ is a different FINST than one caused by $orange_1$ and is different from one caused by $apple_2$. But then, throughout our lives we are faced with different causes of FINSTs in our visual system. This means that we would need an indefinite number of FINSTs to deal with the indefinite number of visual stimuli. *Prima facie*, this may sound implausible not only because of the required proliferation of symbols, but also because minds are assumed to be instantiated in computational systems with *finite* alphabets of symbols. Having a mind with an indefinite number of symbols seems to threaten this assumption. However, the assumption applies only to how minds are *instantiated* in computational systems—i.e. the requirement is that there is a finite number of elements with syntactic/inferential/functional roles. This requirement can be met while assigning to the same element (individuated purely by its syntactic/inferential/functional role) different contents, based on contexts (as is the case with visual demonstratives).

1.D. Complexity in representing particulars and properties:

In this section I argue that the nomic approach that I am pursuing in this dissertation is not more complicated than the singular causal approach. Both approaches need to appeal to at least two more strategies for accounting for content. The fact that the approach that I am proposing is about as complicated as the singular causal approach, is a reason for considering it—a reason that will stand regardless of whether the singular causal approach manages to solve the "which link" problem, or not.

Suppose that one chooses singular causal relations to ground the relations between symbols and external entities. Thus, suppose, as the proponents of singular causal strategies of content hold, that representations of particular objects get their content in virtue of entering into singular causal relations with the particular object that is their cause. There is now the problem of accounting for how representations of properties get their content. Obviously, predicates cannot get their content in the same way as representations of particular objects do, since that would make them representations of particular objects, as opposed to representations of properties. So there must be another way and this way can appeal only to representations of particular objects. That is, unless modifications of the singular causal strategy are allowed. How do we get from representations of particular objects to representations of properties? Perhaps we can *generalize* over the represented objects? But to generalize over represented objects we must first be capable of representing the properties of these objects so that we can cluster the representations of objects based on those representations of properties. But, obviously, to appeal to representations of properties in trying to explain how representations of properties get acquired is to beg the question. Without representing the objects' properties, however, there seems to be no principled way in which to generalize over the represented objects *if* we stay within the limits of singular causal strategies of content. There are moves that can be made here, though:

One could supplement singular causal strategies of content with teleosemantics and have the latter account for the content of primitive representations of properties[32]. Alternatively, teleosemantics may be used to solve the *ad hoc*

[32] Devitt & Sterelny (1999, p. 162), who defend a singular causal theory of content, make a similar move both for the reason pointed out in the main text and also to deal with the *qua* problem with respect to properties: "We are attracted by a less ambitious use of teleology to explain meaning. Instead of taking biological functions to determine the contents of *thoughts* we take them to determine the contents of more basic representational states, *perceptions*. Perceiving a rabbit as a rabbit is a matter of being in a state with the biological function of representing a rabbit. An interesting thing about this idea is that it does not *replace* the historical-causal theory of reference fixing, it *supplements* it. That theory, it will be remembered, suffered from the *qua*-problem [with respect to *properties*]: In virtue of what is a particular grounding of 'rabbit' a grounding in rabbits rather than mammals, vertebrates, or whatever? The present idea offers a teleological answer: the grounding is in rabbits

problem: if one can distinguish between representations of particular objects and representations of properties based on their function (i.e. through teleosemantics), then one can define their content through an appeal to the respective causal relation.

Another reason to appeal to teleosemantics is to account for misrepresentation. Representations are taken to be different from natural signs in that representations can *mis*represent, while natural signs cannot be *mis*effects of their causes. Thus, a theory of representation has to be a theory of misrepresentation, as well. Demonstratives are representations and a theory of them should give an account of how they misrepresent. But singular causal strategies of content do not seem capable of doing so. On singular causal strategies of content demonstratives would simply refer to what caused them, leading to the unpalatable conclusions that a) anything that causes a demonstrative is its content; and b) demonstratives do not misrepresent.[33] Consider what happens when a token of a representation of a particular object is caused by the pokings of a neurosurgeon—being the cause of the

---

because it involves a perceptual state that has the function of representing rabbits. The teleological theory of perception becomes an essential part of the theory of groundings."

[33] Kaplan (1989, p. 490) argues that hallucinations are an instance where demonstratives misrepresent. But Kaplan's demonstratives are mediated through intentions. This would make Kaplan's demonstratives mediated through sensory percepts (as Levine (forthcoming) argues). When the percepts misrepresent (as in cases of hallucination), so does the demonstrative. The question then becomes whether these sensory percepts can misrepresent. As far as singular causal strategies are concerned, an unmediated sensory demonstrative (e.g. a visual index (Pylyshyn's FINST)) would face the problems in the main text.

representation, then the representation would refer to the pokings—surely an unintuitive result. More importantly, whatever the cause of the token of a representation of a particular object, it is going to have a referent—the cause. This means that it will never misrepresent. Here it looks like another theory of content, like teleosemantics, is needed to account for the *normal* conditions when a demonstrative is tokened. Then the two conclusions above can be avoided—pokings by a neurosurgeon will not fall under the normal conditions under which a representation of a particular object is tokened and when they *are* the cause of the representation, it can be claimed that the representation misrepresents.

Now, we have seen that a singular causal strategy needs to appeal to teleosemantics to account for representations of properties and misrepresentations. But that is not all. Singular causal strategies also need to appeal to conceptual role semantics to account for logical operators. Since the content of logical operators is taken to consist of types of inferential relations between symbols, conceptual role semantics is standardly taken to provide the best account of their meaning.

The above suggests that in appealing to singular causal strategies one would end up appealing to teleosemantics (for representations of properties and to account for misrepresentation) and to conceptual role semantics (for logical operators). Now, the approach that I suggest pursuing also appeals to three strategies—it has nomic, singular causal, and conceptual role semantics elements. In Chapter 3 I argue that this approach is sufficient to account for how represenations of particular objects get their content and for how they misrepresent. Since it has nomic elements it is obviously capable of accounting for representations of properties. Since it also has conceptual role semantics elements, it can account for logical operators. As a result, in that it doesn't appeal to more strategies than the supplemented singular causal approach, my approach deserves equal consideration.

1.E. Conclusion:

I have argued that singular causal strategies of content face a significant problem in the "which link" problem. The latter cannot be resolved by an appeal to triangulation (whether actual or counterfactual). The troubles of singular causal strategies with the "which link" problem constitute a reason to pursue other strategies, which is what I do in this dissertation. Furthermore, in order to account for representations of properties, misrepresentation, and logical operators, the singular causal strategies seem to be in need of being supplemented with teleosemantics and conceptual role semantics. Now, the approach that I will explore in this dissertation also has three elements—nomic, singular, and conceptual role semantics' ones. This makes my approach equally complex and this is therefore a reason to pursue its potential as well, as I do in this dissertation.

Chapter 2: Fodor's Theories of Content

2.A. Introduction:

In this dissertation I examine how viable Fodor's asymmetric dependency account (ADA) is, as a theory that explains how the mind connects with the world. As I explain in this chapter, ADA falls under nomic strategies of content—strategies centered on lawful causal relations between properties and symbols. The reasons for choosing to explore the potential of nomic strategies were laid out in the last chapter where I argued that singular causal strategies of content (strategies centered on causal relations between particulars) face the "which link" problem, which cannot be resolved by an appeal to triangulation (whether actual or counterfactual). In this chapter I first present the motivations behind ADA (Section 2.B), then ADA itself (Section 2.C), and finally the problem for it that will be of concern in this dissertation—the problem of accounting for representations of particular objects (Section 2.D).

2.B. The motivations behind ADA:

ADA is designed by Fodor as a key step in the quest for naturalizing psychology. Crudely put, if a property P can be naturalized, then P is a property that "it is possible for physical things to have" (Fodor (1995, p. 5)). Fodor assumes that every non-basic state and non-basic law can be naturalized to some basic states and basic laws. Thus, if basic states and laws are the states and laws of the stuff of which the universe is made, then Fodor assumes that non-basic states and laws supervene[34] or reduce[35] to basic ones. Psychology is a non-basic science. As such, if Fodor's naturalization assumption is right, then psychology's laws must be naturalizable. That is, a) the states that its laws subsume must be such that it is possible for physical things to have them; and b) the psychological nomic relations must be implementable in mechanisms specified through properties and laws from the basic sciences. These implementing mechanisms are ones "in virtue of whose operation the satisfaction of a law's antecedent reliably brings about the satisfaction of its consequent" (Fodor (1995, p. 8)). The way the mechanism $M_1 \rightarrow M_2$ (where M_1 and M_2 are basic states and M_1 is sufficient for M_2) implements the non-basic science law $L_1 \rightarrow L_2$ (where L_1 and L_2 are non-basic states and L_1 is sufficient for L_2)[36] is by L_1 being sufficient for M_1 and M_2 being sufficient for L_2 (*cf.* Fodor (1995, p. 10)). Then M_1 naturalizes L_1, M_2 naturalizes L_2, and $M_1 \rightarrow M_2$ naturalizes $L_1 \rightarrow L_2$.

---

[34] A property A supervenes on a property B iff there can be no change in A without a change in B.

[35] A property A reduces to a property B iff A is identifiable with B. Importantly, reduction is not the same as elimination: while A might reduce to B, still both A and B are real.

[36] For simplicity, it is assumed that $L_1 \rightarrow L_2$ is not multiply realized.

Fodor argues that the laws in psychology are intentional—they are about propositional attitudes (beliefs that Px, desires that Px, etc.). His argument is that intentional laws form the best psychological explanations and no other alternative explanations are even remotely plausible:

> That people (and, surely, other higher organisms) act out of their beliefs and desires, and that, in the course of deciding how to act, they often do a lot of thinking and planning, strikes me as maybe empirical in principle but surely not negotiable in practice. (Fodor (1995, p. 3-4))

If psychology is naturalizable, as Fodor believes, then intentional laws must somehow be naturalized. This would involve naturalizing the following properties of intentional states:

a) Thought is *productive* (there is an indefinite number of thoughts that can be thought) and *systematic* (e.g. being able to think 'if p then q' entails the ability to think 'if q then p').

b) Intentional states can be causally responsible for our behavior.[37]

[37] This deserves a bit of elaboration. Some, like Davidson (1970), believe that the laws that cover intentional causes have to be strict laws (with no *ceteris paribus* clauses). The reason provided is that if they are not strict laws (i.e. if they are like the laws of psychology), then they could not be causal laws, and then one would not be able to explain how intentional states have causal powers (a causal law is

c) Thought processes could be rational—they can take the form of valid arguments.

d) Propositional attitudes have semantic properties—they are about things (e.g. a belief that the sun is yellow is about the sun and the property of being yellow) and can be true or false.

---

one where the causes are nomologically sufficient for the effects). This argument entails that mental states, if they are to be covered by causal laws, are *identified with* physical properties (since only physical properties enter in strict causal laws). The problem is that many, not just Fodor, believe that mental states and symbols are multiply realized. That is to say that a) any disjunct of a disjunction of lower-order properties is sufficient for the instantiation of the higher-order property that is the symbol; and b) the instantiation of the symbol is sufficient for the instantiation of the disjunction, but not for the instantiation of any of the disjuncts (this definition of multiple realization is paraphrased from Fodor (1995, p. 11)). For example, the property of being a neural state of type P can be sufficient for the instantiation of symbol 'x' while the instantiation of 'x' is sufficient only for the instantiation of the disjunction [the property of being a neural state of type P, or the property of being a circuit state of type Q, or the property of being…] but not for any of the disjuncts. Now, if mental states are multiply realized, then they cannot be *identified with* particular physical properties as Davidson's argument entails. Fodor's reply (*cf*. Fodor (1990, Ch. 5)) is to deny the premise that if a law is not strict, then it cannot be causal. He argues that *ceteris paribus* laws can be causal laws *when* their *ceteris paribus* clauses are satisfied. In this dissertation I will assume that something like Fodor's response to Davidson is on the right track.

If the naturalization is to be accomplished successfully, then these properties of intentional states above have to be naturalized in a way that does not presuppose any intentional properties.

The natural way to explain how thought is productive and systematic is to treat intentional states as sentences in a language (*cf.* Fodor (1990, p. 18-19)). Natural languages are productive because of a finite set of words (their vocabulary) an indefinite number of sentences can be constructed. So, if we postulate mental representations (symbols[38]) in the mind and rules for combining them, then we can explain the productivity of thought—out of a finite set of symbols an indefinite number of thought-sentences can be made. The systematicity of thought can be similarly explained in the same way it is explained for natural languages—being able to think 'if p then q' entails that one possesses the symbols 'p', 'q', and the implication operator '→'. This in turn entails the ability to think 'if q then p'. Since Fodor explains the naturalization of the productivity and systematicity of thought through symbols with semantic and syntactic properties, this commits Fodor to a) a

[38] I use the words "symbol" and "mental representation" interchangeably. Being a symbol is a *property* of the symbol's tokens (instantiations). A neural state can be an example of a token of a symbol. Treating symbols as mental representations is contentious, however, since symbols are discrete and there are arguments that there are analog mental representations (e.g. the fineness of grain of experience argument). However, since nothing in my dissertation hangs on whether mental representations can be analog, I will treat mental representations as symbols.

representational theory of mind and b) a "language of thought"[39] view (*cf.* Fodor (1990, p.16)). In other words, Fodor is committed to treating thought as possessing characteristics (syntax and semantics) that natural languages also possess.

All of this, of course, leaves one the task of naturalizing the symbols and syntax. Symbols are individuated by their semantic and syntactic properties. Thus, the task of naturalizing symbols is the task of explaining how physical entities can have semantic and syntactic properties. The semantic properties of a symbol are those that determine its meaning. Fodor's view on how physical entities can have semantic properties will be explained when I get to Fodor's asymmetric dependency account in Section 2.C. The syntactic properties of a symbol are those that determine which combinations with other symbols it can enter in. In this regard, Fodor says that a physical entity's syntax is "one of its second-order physical properties…its syntactic structure [is] an abstract feature of its (geometric or acoustic) *shape*" (Fodor (1990, p. 22)). The idea is that a symbol's syntax supervenes[40] on a *physical property* like shape. As an illustration, consider a set of building blocks with different shapes

[39] More precisely, the language of thought is the language that our minds use to think. As any language, the language of thought is a system composed of a vocabulary of symbols (representations) and rules for the formation and transformation of the latter. *Cf.* Fodor (1975).

[40] Fodor actually says that "to all intents and purposes, syntax *reduces* to shape" (Fodor (1990, p. 22, *italics* mine)). This would mean that syntax is *identified with* shape. However, if syntax is "an abstract feature *of*…shape", it cannot be then *identified* with shape. As a result, in my opinion here syntax is better treated as *supervening* on shape. This in no way interferes with Fodor's other points and can be taken as a helpful amendment.

where, because of its shape, a block can combine with some, but not other blocks. Here it is a block's shape that determines the block's syntactical properties and it is the blocks' shapes that determine the system of rules for combination of blocks that is the syntax of the construction set. Note that different physical entities can have the same syntax when they have the same shape (this is part of the reason why symbols are multiply realized). Importantly, syntax supervenes on a disjunction of sets of shapes. For example, two different sets of building blocks can have the same rules of combination, despite having members with vastly different forms. The result is that the members of one set of physical entities can have different shapes from the members of another set, but both sets can have the same syntax (this is also part of the reason why symbols are multiply realized).

Another of the characteristics of intentional states, of intuitive and explanatory plausibility, is that intentional states have causal powers: if I am thirsty, then my *belief that* there is water in the fridge will *cause* me, *ceteris paribus*, to go to the fridge. If intentional states are naturalized into symbols, then the causal powers of intentional states must be explained through the causal powers of symbols. We saw above that the syntax supervenes on shape. Shape is a physical property and as such determines some of the causal powers of symbols. Among those causal powers are the powers to combine with other symbols. A symbol has these 'syntactic causal powers' in virtue of its syntax, which, as we saw above, is an abstract feature of its shape. Fodor argues that it is a symbol's syntax that would explain how the semantic properties of a symbol can be causally sufficient for our behavior.

To see this, the first step is to appeal to a fact from formal logic that is owed to Gödel—namely that (in first-order logic) there is a correspondence between semantic truth and syntactic provability. This means that semantic relations among symbols can be, as Fodor says, "mimicked" (Fodor (1990, p. 22)) by syntactic ones. Türing has further shown that a machine that operates by transforming symbols can be structured in a way in which its syntactic transformations of symbols correlate with semantic relations between the symbols. Namely, the relations of inference between symbols in a valid argument can be mimicked by a machine whose transformations of symbols are governed by the symbols' syntactical properties.[41] This explains in non-intentional language how intentional states can be rational and it shows how the semantic properties of symbols can have causal powers—when they correlate with the syntactic properties of symbols in an appropriate fashion. In a few words: mental processes are computationally implemented (*cf.* Fodor (1995, p. 14)).

The remaining problem in the naturalizing of intentional states is to explain how symbols have semantic properties. This is the topic of the next section:

---

[41] *Cf.* Fodor (1990, p. 22-23): "the machine is so devised that it will transform one symbol into another if and only if the symbols so transformed stand in certain *semantic* relations; e.g., the relation that the premises bear to the conclusion in a valid argument. Such machines—computers, of course—just *are* environments in which the causal role of a symbol token is made to parallel the inferential role of the proposition that it expresses."

2.C. The asymmetric dependency account:

2.C.I. Holism, the analytic/synthetic distinction, and the disjunction problem:

The remaining problem in the naturalizing of intentional states is to explain how symbols have semantic properties. According to Fodor, one misguided way to do so is to try to define the semantic properties of symbols through their inferential or functional relations with other symbols. This is the way of conceptual role semantics where a symbol's meaning is identified with the symbol's role in a computational system. Fodor, however, argues that this can't be the right way to go: he points out that ever since Quine's "Two Dogmas of Empiricism" there has been substantial doubt that an analytic/synthetic distinction can be made.[42] With respect to conceptual role semantics, this means that one cannot define in a principled manner the meaning-constitutive relations of a symbol, *unless* the meaning-constitutive relations are *all* of the relations a symbol enters in. But since it is practically impossible for two people to share all of their beliefs, desires, etc., it follows that no two people ever share all of their inter-symbol relations. That is, if all of a symbol's relations are meaning-constitutive, it follows that no two people ever share thoughts with the same content.

---

[42] A statement is analytic if it is true in virtue of the meaning of its constituent parts (e.g. "bachelors are unmarried"). A statement is synthetic if it is true in virtue of how the world is (e.g. "bachelors are party-goers"). Quine has argued that this distinction cannot be substantiated. For an extensive review of the arguments see "The Analytic/Synthetic Distinction" by Rey in *Stanford Encyclopedia of Philosophy*.

But that would make psychology dead in the water as a science, because, after all, psychology is about generalizations over the intentional states of different people. If no two people can share their intentional states, then psychology cannot be done. This problem is known as the "holism problem".

In order to avoid the holism problem Fodor foregoes conceptual role semantics and defends informational semantics. According to informational semantics a symbol's meaning supervenes on the symbol's causal relations: 'dog' means DOG because dogs cause 'dog'. The attraction of informational semantics is that they are atomistic—'dog' can be the only symbol in the universe and it can still mean DOG if it is caused by dogs. This entails that on informational semantics a symbol's meaning is independent from its relations with other symbols. According to Fodor, this is a positive feature of informational semantics, because it makes such theories immune to the holism problem. However, while informational semantics accounts avoid the holism problem, they face an intimidating problem of their own—the "disjunction problem":

The disjunction problem is to explain a property of our mental representations that Fodor calls "robustness" (Fodor (1990, p. 91))—a symbol can be caused in many different ways, some of which are not meaning-constitutive. Thus, 'dog' can be caused by cats on dark nights, while keeping its meaning DOG (and not switching to DOG OR CAT ON A DARK NIGHT, for example). The problem for informational semantics of distinguishing the meaning-constitutive causes of a symbol from the non-meaning-constitutive ones is known as the "disjunction problem". Avoiding or solving this problem is what is required if one is to use informational semantics to naturalize intentional states. This is what Fodor attempts to do through his own version of an atomistic informational semantics account:

2.C.II. The current version of ADA:

Fodor's informational semantics theory is now widely called his "asymmetric dependency account" (ADA). It has endured a few changes since its introduction in *Psychosemantics*.[43] Below is the most current version of it:

The symbol 'x' means the property X if:

i) There is a nomic relation between the property of being an instantiation of X and the property of being a tokening of 'x'[44].

ii) Tokenings of 'x' are "robust", i.e. they could be caused by instances of Y, where Y ≠ X.

[43] Fodor (1995) drops the condition in his (1990) that "[s]ome 'X's are actually caused by Xs" (Fodor (1990, p. 121)). *Cf*. Fodor (1995, p. 90): "What a thought represents is largely independent of its *actual* causal history if the informational version of externalism is true. Thoughts of cats are thoughts of *cats* not because cats *do* cause them but because cats *would* cause them under circumstances that may be largely or entirely counterfactual." See also the Appendix B in Fodor (1995)).

[44] *Cf*. Fodor (1990, p. 120-121) and (1987, p. 164, ff. 6).

iii) For all the instances of Y, where Y ≠ X, if an instance of Y causes a tokening of 'x', then its doing so asymmetrically depends on (i), i.e. on the nomic relation between the property of being an instantiation of X and the property of being a tokening of 'x'.

iv) The dependence in (iii) is synchronic.

The first thing to note is that ADA specifies a sufficient, but not a necessary condition for 'x' meaning X. Thus, the theory is safe from objections that 'p' means P, but without satisfying ADA's conditions. Symbols can have their content because they satisfy conditions other than ADA's. Second, note that the nomic relations (between the property of being an instantiation of X and the property of being a tokening of 'x') are instantiated in singular causal relations (between instantiations of X and tokenings of 'x') where these singular causal relations can involve causal chains with links that are tokens of symbols other than 'x'. However, while these tokens of symbols help form the singular causal chains that instantiate the nomic relations, they do not play a role in individuating 'x''s content. The latter is individuated *only* through the nomic relation between the property of being a tokening of 'x' and a property of being an instantiation of X. The particulars that help instantiate this nomic relation are not meaning-constitutive. This makes ADA an atomistic account—an account where a symbol's content is defined independently of the symbol's relation with other symbols. That ADA is an atomistic account means that it is possible that some physical entities can have semantic properties without being related to other physical entities with semantic properties.

Now for an explanation of the conditions: Condition (i) involves nomic relations. Thus, a word on the ontological assumptions of Fodor with respect to nomic relations is in order. He assumes that nomic relations involve *only* properties. This is evident from the fact that to avoid the problem of representing particular objects (see next section), Fodor does not argue that there are nomic relations between particular objects, but suggests that every individual has a corresponding *property* of being that individual.[45] Furthermore, nomic relations are counterfactual-supporting for Fodor[46]. That they are so specified means that even if no instances of X and no tokenings of 'x' entered in singular causal relations with each other, there could still be a nomic relation between the property of being an instantiation of X and the property of being a tokening of 'x'.[47] The nomic relations not in the basic sciences are also *ceteris*

---

[45] Also, see Fodor (1990, p. 93): "*Ontologically* speaking, I'm inclined to believe that it's bedrock that the world contains properties and their nomic relations; i.e., that truths about nomic relations among properties are deeper than—and hence are not to be analyzed in terms of—counterfactual truths about individuals."

[46] *Cf*. Fodor (1990, p. 93): "I assume that if the generalization that Xs cause Ys is counterfactual supporting, then there is a "covering" law that relates the property of being *X* to the property of being a cause of *Y*s: counterfactual supporting causal generalizations are (either identical to or) backed by causal laws, and laws are relations among properties."

[47] That they are so specified also means (according to Fodor) that even if no instances of X and no tokenings of 'x' *existed*, there could still be a nomic relation between the property of being an instantiation of X and the property of being a tokening of 'x'. This is important since it may allow for

paribus laws. This means that they resist exceptions when those exceptions are due to a failure of the *ceteris paribus* clauses. As a result, when the *ceteris paribus* clauses of a law relating the property of being X with 'x' are broken, tokenings of 'x' could be caused by instances of Y (where Y ≠ X) and instances of X may fail to cause tokenings of 'x', without such cases being counterexamples to the law.[48] Finally nomic relations "cover" singular causal relations where the latter obtain between particular objects. In Fodor's words:

> [S]ingular causal statements need to be covered by causal laws. That means something like:
>
> 4. Covering principle: If an event *e*1 causes an event *e*2, then there are properties *F*, *G* such that:
>
> 4.1. *e*1 instantiates *F*
>
> 4.2. *e*2 instantiates *G*

ADA to account for how we represent uninstantiated properties like being an instance of a unicorn. *Cf.* Fodor (1990, p. 100-101): "I take it that there can be nomic relations among properties that aren't instantiated; so it can be true that the property of being a unicorn is nomologically linked with the property of being a cause of "unicorn"s *even if there aren't any unicorns*. Maybe this cashes out into something like *"there wouldn't be nonunicorn-caused "unicorn" tokens but that unicorns would cause "unicorn" tokens if there were any unicorns*. And maybe *that* cashes out into something like: *there are nonunicorn-caused "unicorn" tokens in worlds that are close to us only if there are unicorn-caused "unicorn" tokens in worlds that are close to them*."

[48] *Cf.* Fodor (1990, p. 152-154).

and

4.3. “*F* instantiations is sufficient for *G* instantiations” is a causal law

> When a pair of events bears this relation to a law…the individuals are each *covered* or *subsumed* by that law and…the law *projects* the properties in virtue of which the individuals are subsumed by it…[W]hen an individual is covered by a law, it will always have some property in virtue of which the law subsumes it. (Fodor (1990, p. 142-143))

These assumptions will be taken for granted in this dissertation, *pace* nominalists (who do not believe that properties (as universals) exist), Humean bundle theorists (who treat particulars as bundles of properties), and regularity theorists (who treat nomic relations as mere constant conjunctions). Now for the rest of the conditions:

What distinguishes natural signs from symbols is that symbols can *misrepresent*. Or, as Fodor puts it, symbols are “robust”. A natural sign is always a sign of its cause—whatever the cause is. It is merely a register, as the effect, of whatever caused it to come into existence. Thus, a natural sign can never be mistaken—it can never be a mis-effect—and this is what makes it different from a representation. Thus, condition (ii) is to account for the difference between natural signs and symbols—to account for the ability of symbols to misrepresent. But here, of course, is where the disjunction problem arises: if a symbol gets its content from the nomic relations it enters in *and* if the symbol can misrepresent, then this means that only some of the nomic relations are going to be veridical (i.e. meaning-constitutive). That is why condition (iii) is needed. Condition (iii) says that the way to distinguish between the meaning-constitutive nomic relations involving 'x' and the other ones is to look at the counterfactuals: the meaning-constitutive nomic relations involving 'x' are the ones, but for which the others wouldn't exist. Thus, if there was no nomic relation between dogs and 'dog', then there wouldn't be a nomic relation between dog-looking-cats and 'dog', but not *vice versa*.

ADA's final condition states that the asymmetric dependence is *synchronic* as opposed to diachronic. Thus, for example, one's *present* disposition to apply 'dog' to dogs must not depend on any *present* dispositions to apply 'dog' to cats, and one's *present* disposition to apply 'dog' to cats must depend on one's *present* disposition to apply 'dog' to dogs (*cf.* Fodor (1987, p. 108-9)). This condition is included to deal with cases where one's disposition to apply 'x' to X is acquired entirely through non-instances (Y, where Y ≠ X). For example, the disposition to apply 'horse' to horses may be acquired entirely through horse-looking-cows.[49] Once the disposition to apply 'horse' to horses is acquired, the disposition to apply 'horse' to horse-looking-cows would depend on the disposition to apply 'horse' to horses. However, the present disposition to apply 'horse' to horses also depends on the *past* disposition to apply 'horse' to horse-looking-cows. Thus, we don't have a case of *asymmetric* dependence, but we still have error (application of 'horse' to horse-looking-cows),

[49] *Cf.* Fodor (1987, p. 109).

contrary to ADA which explains error through asymmetric dependence. The fourth condition deals with this case because it is a case of *diachronic* dependence of one's *present* disposition to apply 'horse' to horses on the *past* disposition to apply 'horse' to horse-looking-cows.

2.C.III. Rey's fifth condition:

Rey complements Fodor's ADA account (p. c.) with a fifth condition:

v) The asymmetric relation in (iii) is one on which other asymmetric relations involving a law between 'x' and X depend, but not *vice versa*.

This condition helps ADA cope with a problem raised by Rey (p. c.) of multiple asymmetric dependencies involving the same symbol. The problem goes as follows: suppose that because of a nomic relation between 'x' and women there is a nomic relation between 'x' and people with skirts, but not *vice versa*. Now it is possible that because of the nomic relation between 'x' and people having skirts, there is a nomic relation between 'x' and people with baggy pants, but not *vice versa*. We now have a case of 'x' being involved in *two* asymmetric dependencies between nomic relations. The problem is to decide in a principled manner which one is the meaning-constitutive one. Rey's fifth condition is such a principled manner. The suggested condition takes nothing of the externalist spirit of ADA and helps ADA deal with this problem.

2.D. The problem of representing individuals:

Multiple objections have been raised for ADA, as to whether it succeeds to solve the disjunction problem. I will not analyze here whether these objections are successful or not, but will assume that ADA, or some modification of it, might offer a plausible way to solve the disjunction problem. [50] What I will analyze in this

[50] ADA faces the problem of distinguishing between necessarily co-instantiated properties (examples of necessarily co-instantiated properties are rabbit and undetached-rabbit-part. They are not co-extensive (the part is not identical to the whole) and co-occur in every world in which one of them exists. Fodor's asymmetric account would not distinguish between representations of them precisely because they never appear separated. A special case of this problem is distinguishing between necessarily un-instantiated properties (some examples of necessarily un-instantiated properties are (being a) round-square, miracle, and monster) (*cf*. Rey (1997, p. 250)). First, Fodor's asymmetric account would not be able to account for them because necessarily un-instantiated properties do not enter in nomic relations. Second, even if we can talk about nomic relations with the null set, then ADA would have the unpalatable consequence that representations of miracle and monster would have the same content. Another problem for ADA is that it risks gratuitous lockings (gratuitous lockings are referential relations of symbols that have no cognitive significance (*cf*. Rey (1995, p. 6-7)). For example, 'cow' might enter in asymmetrically dependent nomic relations with cows and cow-looking-horses. But it might also enter in asymmetrically dependent nomic relations with pokings-by-a-neurosurgeon-when-potassium-levels-are-high and pokings-by-a-neurosurgeon-when-sodium-levels-are-low. However, the latter asymmetrically dependent nomic relations are cognitively *in*significant—their presence would suggest that 'cow' is ambiguous, but since it seems obvious that it is not, they

dissertation is whether and how ADA can account for how representations of particular objects get their content. It has been argued that ADA cannot account for representations of particular objects:

Adams & Aizawa (1997, p. 274-5) point that ADA cannot account for how representations of particulars get their content. ADA is based on nomic relations and Fodor himself assumes that there are no nomic relations between particular objects. Nomic relations are assumed to be between properties. This leaves ADA as incapable of addressing the question of how representations of particular objects get their content. One way to help ADA here is to appeal to properties such as the property of [being an instance of Aristotle] and to nomic relations between the property of being an instance of Aristotle and the property of being a token of 'Aristotle'.[51] But, as Adams (2003, p. 156) objects, this makes 'Aristotle' mean a property that can be shared by clones of Aristotle—entities that are numerically different from (the original) Aristotle, but share with (the original) Aristotle the property of being an instance of Aristotle. What makes it the case that 'Aristotle' means the original, as opposed to the clones? Fodor could insist on treating the property [being an instance

---

should be ignored ('cow' better be not ambiguous between cows and pokings-by-a-neurosurgeon-when-potassium-levels-are-high). These problems have already been widely discussed (*cf*. Fodor (1987, 1990) and Loewer & Rey (1991)). Since a conclusive argument for or against ADA has so far not been provided, for the purposes of examining whether and how ADA deals with the problem of representing particulars I will simply assume that it, or some version like it, has the tools to deal with all the other problems.

[51] *Cf*. Fodor (1995, Index B, p. 118).

of Aristotle] as unique for an individual. But this is dubious: as Adams (2003, p. 156) points out, "why would anyone ever have thought that individuals do not feature in laws?" Allowing such properties and laws between such properties entails a vast profligacy of laws—surely not the result we want if we are after a lean ontology.

A third way to try to avoid this problem is to note that it is not really a problem for ADA—ADA specifies only sufficient conditions for 'x' meaning X. Thus, the fact that it cannot account for how individuals get represented does not show that there is something wrong with it. One, however, might object that not being able to account for how representations of individuals get their content makes ADA lack in explanatory power and that *if* another strategy for accounting for content (e.g. a strategy that appeals to both broad and narrow content) does better, then one should pursue the latter strategy and renounce purely atomistic externalist strategies like ADA. This is in effect what I will argue in Chapter 3.

A fourth way to try to avoid this problem is by arguing that representations of particular objects are not in fact needed: Quine (1960b, p. 343-347) argues that we can transform a sentence with singular terms/arguments/variables into a sentence without one. To this purpose six combinators are used:

a. *Derelativization*: (Der P)$x_1 \ldots x_{n-1}$ if and only if there is something x_n, such that $Px_1 \ldots x_n$;

b. *Major Inversion*: (Inv P)$x_1 \ldots x_n$, if and only if $Px_n x_1 \ldots x_{n-1}$;

c. *Minor Inversion*: (inv P)$x_1 \ldots x_n$, if and only if P$x_1 \ldots x_{n-2} x_n x_{n-1}$;

d. *Reflection*: (Ref P)$x_1 \ldots x_{n-1}$, if and only if P$x_1 \ldots x_{n-1} x_{n-1}$;

e. *Negation*: (Neg P)$x_1 \ldots x_n$, if and only if not P$x_1 \ldots x_n$;

f. *Cartesian Multiplication*: (P x Q)$x_1 \ldots x_n$, $y_1 \ldots y_n$, if and only if P$x_1 \ldots x_n$, and Q$y_1 \ldots y_n$.

An example (from Robin Clark, "Variables, Interpretations and Quine-like Combinators", p. 2):

a. ***E***x***E***y(Pxy ***** Qxyx)

b. (P x Q)xyxyx by Cartesian Multiplication, f

c. (Inv (P x Q)xxyxy by Major Inversion, b

d. (inv (Inv (P x Q)))xxyyx by Minor Inversion, c

e. (Inv (inv (Inv (P x Q))))xxxyy by Major Inversion, b

f. (Ref (Inv (inv (Inv (P x Q)))))xxxy by Reflection, d

g. (Der (Ref (Inv (inv (Inv (P x Q))))))xxx by Derelativization, a

h. (Ref (Der (Ref (Inv (inv (Inv (P x Q)))))))xx by Reflection, d

i. (Ref (Ref (Der (Ref (Inv (inv (Inv (P x Q))))))))x by Reflection, d

j. (Der (Ref (Ref (Der (Ref (Inv (inv (Inv (P x Q))))))))) by Derelativization, a

As Robin Clark (p. 3-4) points out, what Quine actually shows is that we can eliminate only the *syntactic* contribution of variables. The *semantic* contribution remains intact. That is to say that symbols corresponding to variables in a language do not need to be explicitly tokened/written. Whatever their syntactic contribution to a sentence is, it can be substituted *via* the use of Quine's combinators. But this does not entail that we can eliminate the *semantic* contribution of variables. The semantic contribution is needed by Quine's own admission to define the arity of predicates. So we still need variables. In Robin Clark's words:

> Variables play a dual role in first order logic. First, they act as place-holders which indicate the arity of the predicate. Second, multiple occurrences of the same variable may indicate identity of reference. Quine himself notes that ". . . the essential utility of variables is that they mark positions." In particular, consider Quine's Cartesian Multiplication operator:
>
> > we can express our example 'Bxy and not Fwxz' as a single predication '(B x Neg F)xywxz'. (In reading this we have to know, of course that 'B' is two-place and F' is three-place.)
>
> The combinatory calculus developed by Quine eliminates the second of the above functions of variables. Although variables are not explicitly present in a combinatory expression, we need arity markers for predicates in order to properly interpret Cartesian Multiplication. We conclude that semantic variables have not really been eliminated since the arity of the predicates must be known, as Quine says explicitly when discussing his Cartesian Multiplication. (Clark, Robin ("Variables, Interpretations and Quine-like Combinators", p. 4))

The upshot is that, even if we can get rid syntactically of variables, we still need their semantic contribution. This means that attempting to deal away with representations of particular objects *via* Quine's combinators in order to avoid the problem of representing particular objects is not a viable option.

2.E. Conclusion:

In this chapter I have overviewed Fodor's asymmetric dependency account (ADA): Fodor designs ADA to be a theory of content that solves the disjunction problem without being vulnerable to the holism problem. In the rest of this dissertation I will assume that ADA, or some modification of it, offers a plausible solution to the disjunction problem. At the end of this chapter I pointed out the problem with ADA that will be the impetus for the rest of the dissertation: the problem of representing individuals. This problem will be analyzed in the next chapter where I will argue that to explain how particular objects are represented, the logico-syntactic roles of representations of particular objects need to be taken into account, and for this ADA needs to be complemented by conceptual role semantics.

Chapter 3: Configurations and Particulars for ADA

3.A. Introduction:

In Chapter 1 I argued that a) singular causal strategies face the "which link" problem; and b) nomic strategies are about as complex as singular causal strategies. The upshot was that the nomic strategies also deserve consideration. In the previous chapter I presented an example of such nomic strategies—Fodor's ADA—and explained how it solves the disjunction problem. But I also pointed out that ADA faces the problem of representing individuals. In this chapter I examine whether and how ADA can deal with this problem. I argue that to explain how particular objects are represented a) the logico-syntactic roles of representations of particular objects need to be taken into account, and for this ADA needs to be complemented by conceptual role semantics; and b) an appeal to the temporal and first-person indexicals must be made, and for this ADA must be supplemented with singular causal relations.

In Section 3.B I briefly overview ADA and the problem of representing particulars. In Section 3.C I present the notions of configurations, determinates, and determinables and then I tentatively propose in Section 3.D that with the help of super-determinate configurations and conceptual role semantics, ADA can account for how representations of particular objects get their content. The idea is to use conceptual role semantics to account for the logico-syntactic roles of representations of particular objects. The role of representations of particular objects is that of an argument in a definite description. The argument is concatenated with a predicate referring to a super-determinate spatio-temporal configuration instantiated at a space-time region—a trope. The configuration involves four objects (x, l, n, and r) forming a pyramid with sides of length d_{rx}, d_{lx}, d_{nx}, d_{rl}, d_{rn}, and d_{ln}. This configuration is so determinate that there is only one particular at a time that can be at the pyramid's vertex. I use this observation to argue that representations of particular objects get their content through definite descriptions of the sort: right now x is the one and only particular that has the relational spatial property of being at the non-base vertex of the super-determinate configuration described above, where this configuration has its base positioned on my eyes and nose. I also point out in Section 3.D that this view makes use of conceptual role semantics for more than the logical apparatus needed for the definite description. It is also used to define the narrow content of 'I' and is needed for the narrow content of the temporal indexical 'now'. I also argue that singular causal relations are needed to account for the broad content of the first-person and temporal indexicals.

3.B. ADA and the problem of representing individuals:

Fodor (1987, 1990, 1995) developed ADA to show that nomic theories of content are capable of solving the disjunction problem. According to ADA it is solved as follows: the nomic relations that do not involve a symbol's referent depend on the nomic relations that involve the symbol's referent, but not *vice versa*. In more schematic form:

The content of representation 'x' means the property X if:

i) There is a nomic relation between the property of being an instantiation of X and the property of being a tokening of 'x'.

ii) Tokenings of 'x' are "robust", i.e. they could be caused by instances of Y, where Y ≠ X, while still referring to X.

iii) For all the instances of Y, where Y ≠ X, if an instance of Y causes a tokening of 'x', then its doing so asymmetrically depends on (i), i.e. on the nomic relation between the property of being an instantiation of X and the property of being a tokening of ‘x’.[52]

Thus, ADA is an externalist account grounded in nomic relations between properties and symbols. Fodor makes the following ontological assumptions with regards to nomic relations: a) nomic relations involve *only* properties; b) they are counterfactual-supporting; c) the nomic relations not in the basic sciences are also *ceteris paribus* laws; d) nomic relations "cover" singular causal relations where the latter obtain between particular objects. These assumptions have been explained in the previous chapter and will be taken for granted.

Given Fodor's assumption that nomic relations involve *only* properties, it follows that particular objects do not enter in nomic relations. Since particular objects do not enter in nomic relations, it is not clear how ADA can account for representations of particular objects.[53] Fodor (1995) has argued that it can, by appealing to properties such as the property of [being an instance of Aristotle] and to nomic relations between the property of being an instance of Aristotle and the property of being a token of ‘Aristotle’. But, as Adams (2003, p. 156) objects, this makes ‘Aristotle’ mean a property that can be shared by clones of Aristotle—entities that are numerically different from (the original) Aristotle, but share with (the

[52] Fodor's fourth ("synchronic") condition and Rey's fifth condition (about dependencies between asymmetric relations) have been omitted, since they will not be of concern in this chapter.

[53] *Cf*. Fodor (1995, Index B, p. 118) and Adams & Aizawa (1997, p. 274-5).

original) Aristotle the property of being an instance of Aristotle. What makes it the case that ‘Aristotle’ means the original, as opposed to the clones? Fodor could insist on treating the property [being an instance of Aristotle] as unique for an individual. But this is dubious: as Adams (2003, p. 156) points out, “why would anyone ever have thought that individuals do not feature in laws?” Allowing such properties and laws between such properties entails a vast profligacy of laws—surely not the result we want if we are after a lean ontology. For the purposes of this chapter I will assume that Adams & Aizawa are right and ADA, by itself, cannot account for representations of particular objects.

3.C. Super-determinate configurations:

Since ADA is based on nomic relations and nomic relations cannot exist between particular objects, ADA, by itself, cannot account for how particular objects get represented. In Section 3.D I will present how ADA, with some help from conceptual role semantics, singular causal relations, and representations of super-determinate spatio-temporal configurations, can account for how particular objects get represented. But first, in Section 3.C.I, I explain the notion of super-determinate property. Then, in Section 3.C.II I go over issues involving their representation and the role of such representations in producing representations of particular objects.

3.C.I. Determinates and determinables:

Before I begin with spatio-temporal configurations—a brief word on determinates and determinables: the paradigmatic examples are *colored*/*red* and *red*/*scarlet*. *Red* is a determinate of *colored* and a determinable of *scarlet*. One aspect of the determinate/determinable relation that Funkhouser (2006, p. 548-9) provides is: "for an object to have a determinate property is for that object to have the determinable properties the determinate falls under *in a specific way*". For example, being red is a specific way of being colored and being scarlet is a specific way of being red. Funkhouser presents two "truisms" about the determinate/determinable relation: First, the relation between the determinate and the determinable is transitive: if *scarlet* determines *red*, and *red* determines *colored*, then *scarlet* determines *colored*. Conversely, if *colored* is a determinable for *red*, and *red* is a determinable for *scarlet*, then *colored* is determinable for *scarlet* too. Second, the transitive chain of determinates and determinables "does not go on forever". Thus, there are determinables that do not determine anything and determinates that are not determinables to anything. The former Funkhouser calls "super-determinables" and the latter, "super-determinates". I am going to adopt his usage. The examples of super-determinables and super-determinates that he gives are *colored* and *Coca-Cola red*, respectively.

Funkhouser also points out that every determinable has "determination dimensions". Those are the dimensions along which the determinates vary. For example, with respect to *triangular*, the determination dimensions are the lengths of the sides and the angles between the sides. Determinables and their determinates share what Funkhouser calls “non-determinable necessities”: features that the determinables and their determinates share but which do not allow for variation. For example, the non-determinable necessities of *triangular* are “3-sided, closed, plane figure”. I will now apply these notions to spatio-temporal configurations:

3.C.II. Representing spatio-temporal configurations:

In Section 3.C.II.1 I start presenting my account how particular objects are represented, which appeals to spatio-temporal configurations. Being complicated it naturally invites the objection: well, even if this works, why not do it in this simpler way? In Section 3.C.II.2 I examine one such simpler way to account for representations of spatial relations, and explain why I do not think it is actually simpler.

3.C.II.1. The configurations approach:

A quick bit of terminology, before I proceed: space points are *particulars*. Spatial relations (i.e. configurations[54]) are relational *properties* of particulars. An

[54] Configurations are not identical to shapes. A shape is a non-relational property of an object’s boundaries and is defined in terms of these boundaries, while a configuration is a relational property of objects and is defined in terms of the objects’ positions with respect to each other.

example of a spatial relation is the one that obtains between two objects separated by 5 m.

Configurations, being properties, can enter in nomic relations. As a result, ADA can account for representations of them (regardless of whether they are super-determinate or not), without having to appeal to representations of the particulars that compose the configurations.[55] The notion of configuration will end up doing a lot of work in accounting for representations of particular objects. For starters, note that since ADA can account for representations of super-determinate configurations, it can account for representations of super-determinate configurations such as the following super-determinate configuration:

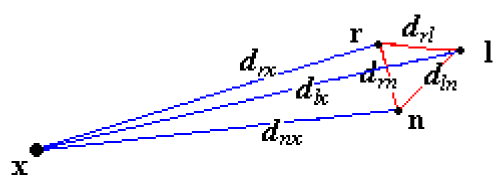

[55] While a particular plays a causal role in the tokening of a representation of a property, it need not be represented for that. According to ADA, a token ‘P’ represents P if the property of being a tokening of ‘P’ enters in the requisite asymmetric nomic relations with P and other properties. That there be a nomic relation between ‘P’ and P requires that instantiations of P cause tokens of ‘P’. Since ADA is the assumed account of reference and not singular causal accounts, merely entering in a singular causal relation can be held as not being sufficient for being a representation of the cause. Therefore, a token ‘P’ can represent P without also representing the object instantiating P that caused the token ‘P’.

Figure 3.1.: This is an instance of a super-determinate configuration involving four objects x, l, n, and r forming a pyramid with sides of length d_{rx}, d_{lx}, and d_{nx} and a base of a triangle with sides d_{rl}, d_{rn}, and d_{ln}.

The represented property here is: [being a configuration involving four objects (x, l, n, and r) forming a pyramid with sides of length d_{rx}, d_{lx}, d_{nx}, d_{rl}, d_{rn}, and d_{ln}], where d_{rx}, d_{lx}, d_{nx}, d_{rl}, d_{rn}, and d_{ln}, are specific lengths. This gets us a lot of what we need to account for how representations of particulars get their content:

If the base of the pyramid is instantiated in a particular creature in a particular way, then only one particular object at a time can be at the vertex of the pyramid. Thus, if the base is instantiated in me so that its three vertices are my two eyes and nose (so d_{rl}, d_{rn}, and d_{ln} would be the lengths between the eyes and the nose), then only one particular object at a time can be at the vertex.[56] We can have the following super-determinate configuration then: [being a configuration such that:

a) it involves the objects x, l, n, and r forming a pyramid;

b) the sides of the pyramid are of lengths d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx};

[56] Note that this does not mean that we need to represent the eyes and the nose or their respective locations. Just like I can represent a property, without representing the particular that caused the representation of a property, so I can represent a configuration, without representing the particulars that caused the representation of a configuration. Note that none of this denies that the particulars and their locations can be causally responsible for the token representations.

c) r, l, and n form the base of the pyramid and r and l are O's eyes, while n is O's nose].[57]

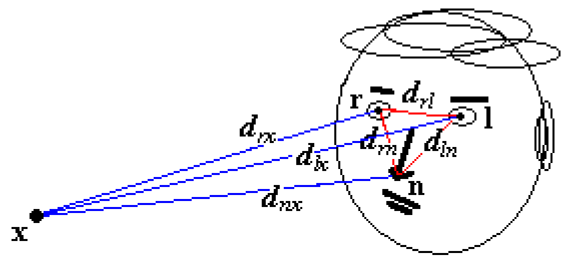

Figure 3.2.: This is an instance of a configuration involving four objects x, l, n, and r forming a pyramid with sides of length d_{rx}, d_{lx}, and d_{nx} and a base of a triangle with sides d_{rl}, d_{rn}, and d_{ln}. The configuration instance is instantiated so that l, r, and n are the left, right eye, and the nose of "Pinocchio" here.

Representations of this configuration ('C_j'[58]) are problematic for ADA because the configuration involves a particular person—O. I have assumed with Fodor that only properties enter in nomic relations. This raises a dilemma for ADA: either the configuration is denied the status of a property because it is specified

[57] I am slightly oversimplifying here: note that there can be two particulars that enter in the pyramidal configuration above—one in front of the observer and one behind her (a mirror reflection of the pyramid around its base). The way to discard the second particular (the one behind the observer) is to add a fourth particular in the base (one that is not in the same plane as the eyes and the nose). I will not do that for presentation purposes—the spatial configuration would become much more complicated than the intuitive pyramidal one I am using.

[58] "'C_j'" is short for "'C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)'".

through a particular, or it is treated as a property but of a very special kind: a trope[59], a property instance like being scarlet at this space-time point. The first horn of the dilemma commits us to denying that ADA can account for representations of the super-determinate configuration above. The second horn commits us to tropes. But if tropes are to help ADA, then they have to enter in the nomic relations of ADA. However, tropes entering in nomic relations is problematic—the same doubts about particular objects entering in nomic relations apply to tropes.[60]

Suppose, however, that we complement ADA with the first-person and temporal indexicals. This then would be enough to explain how the trope above is represented. The idea is to have a predicate 'C_j' that satisfies the following two conditions: a) it nomically covaries with super-determinate configurations like [being a configuration such that it involves four objects (x, l, n, and r) forming a pyramid where the respective distances are d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx} and where l, n, and r are the left eye, right eye, and nose of some or other observer]; and b) it is inferentially and functionally related to a first-person reflexive term ('I') and an indexical for time ('t'). In virtue of 'C_j' being related to 'I' and 't' it doesn't pick out just any pyramidal super-determinate configuration, instead it picks out a particular instance of it: *the*

[59] Tropes, also known as property instances, concrete properties, and abstract particulars can be schematized as follows (*cf.* Funkhouser (2006, p. 14)): "[(O, *t*), P]. This is to be read as "O's having/being P at *t*." O is an object or spatial location, *t* a time (span), and P a property type. Such property instances occur if, and only if, O *really is* P at *t*."

[60] Note that even though it is problematic to allow for tropes to enter in causal relations that are nomic, what is not problematic is allowing tropes to enter in singular causal relations.

one that is set on me, now. This makes the configuration a trope. Note that there is no requirement that 'C_j' gets nomically related to the trope, as there would be if ADA were to account for how the trope is represented by itself. Instead, the trope is *picked out* in virtue of the 'C_j''s nomic relations with the property of being the super-determinate spatial configuration C_j *and* in virtue of the 'C_j''s inferential and functional relations with 'I' and 't'. The picking out of a trope is cashed out as the satisfaction relation that obtains between Fregean definite descriptions and their referents.[61] In short: when 'I' and 't' in 'C_j' are saturated, then the referent of 'C_j' is a trope.

[61] Where "q" and "g" denote tropes, "R(g, I, t)" denotes the three-place relation of [being a trope centered on me at some or other time t], and "C_i" denotes a particular super-determinate pyramidal configuration, the definite description is as follows: $(\boldsymbol{E}q)(g)((R(g, I, t) \& C_i(g)) \leftrightarrow (q = g)))$. This description need not be explicitly represented in the mind, of course, and perhaps isn't. When it is implicit, '$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$' is simply related to 'I' and 't'. Then a token of '$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$' refers to the super-determinate pyramidal configuration with which the predicate is nomically related that is based on the tokening organism (the referent of 'I') at the moment of tokening (the referent of 't'). *Importantly*, I am sympathetic to the objection that merely being related to a) the super-determinate pyramidal configuration; and b) related to 'I' and 't', does not make '$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$' a representations of a trope. *If* this is true, then I can argue (as I do for particular objects) that we need explicitly represented definite descriptions for '$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$' to refer to a trope. But, since my main focus in this dissertation is to account for how objects are represented, I can be charitable and grant that while we need explicitly represented definite descriptions for representations of objects, we do not need them for tropes. Either way, my point about representations of objects stays the same.

So, an organism representing this configuration at time t would use the following predicate: ‘$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$’. Here ‘x’, ‘l’, ‘n’, and ‘r’ are variables with the logico-syntactic role of arguments, but which, importantly, do not need to have referents for the predicate to be tokened. The variables ‘d_{rx}’, ‘d_{lx}’, ‘d_{nx}’, ‘d_{rl}’, ‘d_{rn}’, and ‘d_{ln}’, range over numbers in a sequence. The variable ‘t’ is an indexical for time and ‘I’ is a first-person reflexive term, which I will show later can be made implicit and dropped.[62, 63]

This proposal for representing super-determinate configurations at a space-time location may appear implausible: a) there is what seems to be an *ad hoc* reference to a creature’s eyes and nose; b) the proposal entails that different creatures would not share the same spatial representations; and c), since ‘x’, ‘l’, ‘n’, and ‘r’ have the same functional roles, ‘x’ can just as well be the nose or one of the eyes—a problem, if I want to claim, as I do later on, that ‘x’ refers to the object at the tip of the pyramid. I take those in turn:

[62] Note, by the way, that this proposal is not committed to an infinite number of nomic relations between the pyramidal configurations and predicates, since the visual system does not have infinite resolution. Furthermore, the proposal does not entail that our visual system uses an infinite number of non-complex symbols, despite implementing variables that range over numbers. That is because only two non-complex symbols (‘0’ and ‘1’) are enough to represent any number.

[63] Note that all six of the distance variables are needed here in order to account for a specific pyramid. For a particular organism, however, after a certain age the distances between the eyes and the nose remain set, which means that just three variables would be enough—those that specify the distances between the eyes and nose to the observed particular (we need all three of them to account for irregular pyramids). This means that ‘d_{rl}’, ‘d_{rn}’, and ‘d_{ln}’ are implicit and can be dropped.

With respect to (a): The proposal appeals to configurations involving a creature's eyes and nose. Even ignoring the question of how to locate them[64], why are these chosen and not other organs? Well, the answer is that I chose them for simplicity's sake—a three-sided pyramid is the simplest one. But it would have been more accurate to choose a pyramid with a complex base that has sixteen vertices, where the vertices are the two foveas[65], the two ciliary muscles[66], and the twelve muscles (six per eye) that control the position of the eyes in their orbits. Why the foveas and the ocular muscles? Because I suspect that for retinocentric and head-centered vision it is only these sixteen organs that play a role in locating an object. Were we concerned with a bat, for example, it would be the ears and the muscles that control them. Thus, the appeal to these organs is not completely *ad hoc*—they are appealed to because they are instrumental in detecting spatial relations.

With respect to (b): The proposal above entails that different creatures would not share the same spatial representations. For example, bats will represent different configurations than cats, since bats will represent configurations formed in part by their ears, whereas cats will represent ones that consist (in part) by the cats' eyes. Furthermore, the pyramid's base would need to be smaller for smaller creatures like

[64] I am presuming that there must be some principled way to do so. After all, the psychologists that study retinocentric vision, for example, do appeal to a center in the eye—the fovea.

[65] The fovea is a part of the retina where the photoreceptors are more densely packed.

[66] The ciliary muscles are the two muscles (one per eye) that change the shape of the eye's lens to permit focusing on an object.

mice, since their eyes and nose are located more compactly. This means that larger organisms will represent different configurations than the ones represented by smaller organisms. And it get's worse: wide-eyed-Bob's spatial concepts would be different from narrow-eyed-Andy's because Andy's eyes would be located differently than Bob's. This makes it seem that the current proposal runs against the intuitive view that different creatures can share spatial concepts.

There are two types of spatial concepts that we have—ones referring to spatial properties and ones referring to spatial points. With regards to relational spatial properties like distance, azimuth, and altitude—they can be represented without appealing to specific distances between the organism's perceptual organs. For example, we can talk about being 5 m away from the midpoint of the line between a creature's perceptual organs for space.[67] This means that the above problem does not arise for such properties. What about space points? The account that I give in Section 3.D.II of how particular objects are represented can easily be modified to account for how space points are represented. For present purposes, all that is important to know

[67] For azimuth we have: being on a line L_1 that a) starts off at point M where it intersects line L_2 (where L_2 i) passes through the midpoint (M) of the line running through the organism's visual sensory organs (L_3); and ii) is in the plane of the visual sensory organs and the vestibulary organs (P)); and b) is in P and "leaning" towards the sensory organ closer to the stimulus. For altitude we have: being on a line L_1 that a) starts off at point M where it intersects line L_2 (where L_2 i) passes through the midpoint (M) of the line running through the organism's visual sensory organs (L_3); and ii) is in the plane of the visual sensory organs and the vestibulary organs (P)); and b) is in the plane that contains L_2 and is perpendicular to P; and c) "leans" towards the sensory organ closer to the stimulus.

is that in that account I appeal to definite descriptions containing representations of tropes like the pyramidal ones above. Now, if different organisms don't share these representations, they will have different definite descriptions. However, as long as the organisms represent space points *de re*, as opposed to *de dicto*, this does not matter—the representations of space points will have the same content, despite this content being picked out *via* different descriptions.

With respect to (c): In the proposal above since 'x', 'l', 'n', and 'r' have the same functional roles, 'x' can just as well be the nose or one of the eyes—a problem, if I want to claim, as I do later on, that 'x' refers to the object at the tip of the pyramid. To avoid this problem we can stipulate that 'l', 'n', and 'r' carry information about the sensory organs, while 'x' doesn't. Note that this does not mean that 'l', 'n', and 'r' represent the *particular* organs. It just assumes that they carry information about the property of being a sensory organ. Now, 'x' will be different in that it carries no information about the organs. This will make it semantically different from 'l', 'n', and 'r', and will individuate it as referring to an entity that is not a sensory organ (i.e. an entity that is not on the base of the pyramid).

I have presented an account of how super-determinate configurations can be represented. I also argued that representations of super-determinate configurations can be an essential step in accounting for representations of particulars. The key lies in the observation that some configurations are so determinate that only one particular at a time can stand at a designated vertex. The super-determinate configuration I have appealed to is one that a) involves the objects x, l, n, and r forming a pyramid where the respective distances are d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx} and where l, n, and r are the left eye, right eye, and nose of some or other observer. I now examine the details of this proposal:

3.C.II.2. Coordinate systems:

Just as *shape* is a super-determinable and *square, triangle, circle,* etc. are determinates of it, so too *configuration* is a super-determinable and *square-constellation*, *triangle-constellation*, and *circle-constellation*, etc. are its determinates. One example of a super-determinate of *configuration* is the configuration of two particulars separated by 5 m. Being types of properties, super-determinates can enter in nomic relations. This means that ADA can account for representations of them. This allows for ADA to account for representations of super-determinates of *distance* (*distance* being another determinate of *configuration*).

Now suppose that the predicates 'P', 'Q', and 'R' are representations of super-determinates of *distance*. These are 3-place predicates whose arguments refer to the objects between which the distance obtains and the length of the distance: $Pxyd_{xy}$.[68] Wouldn't then the argument of which all three predicates are predicated at the same time ('$Pxyd_{xy}$ & $Qxmd_{xm}$ & $Rxrd_{xr}$') refer to a point in 3D space? And if so, doesn't

[68] Importantly, the arguments can have only the syntactic role of place holders—they do not need to have referents for the predicate to have a referent.

this solve the problem of representing particulars for ADA? The answers to these questions are two qualified "yes"-s. A particular space point could be represented by 'x' when '$Pxyd_{xy}$ & $Qxmd_{xm}$ & $Rxrd_{xr}$', only if certain conditions are satisfied: one must know a) the distances d_{xy}, d_{xm}, d_{xr} (represented by '$Pxyd_{xy}$', '$Qxmd_{xm}$', and '$Rxrd_{xr}$'); b) the directions of these distances[69]; and c) how y, m, and r are related with each other (at a time). Now, knowing how four points are related with each other (and assuming they are not all in the same plane) is sufficient to specify a coordinate system[70] in 3-D. Thus, if we introduce a fourth point, o (for origin), then we can have o, y, m, and r specify a coordinate system. The introduction of o would allow to avoid mentioning the directions of the distances d_{xy}, d_{xm}, d_{xr}, because the information about their direction would be implicit in the information about the distances d_{xy}, d_{xm}, d_{xr}, d_{xo} and how o, y, m, and r are related with each other. The upshot is that '$Pxyd_{xy}$ & $Qxmd_{xm}$ & $Rxrd_{xr}$' would not individuate a point but '$Pxyd_{xy}$ & $Qxmd_{xm}$ & $Rxrd_{xr}$ & $Sxod_{xo}$' would, if one also knows how o, y, m, and r are related with each other (at a time) and not all four of them are on the same plane:

[69] This qualification is needed to distinguish between mirror images of configurations comprised of x, y, m, and r.

[70] Directions are included in the specification of a coordinate system.

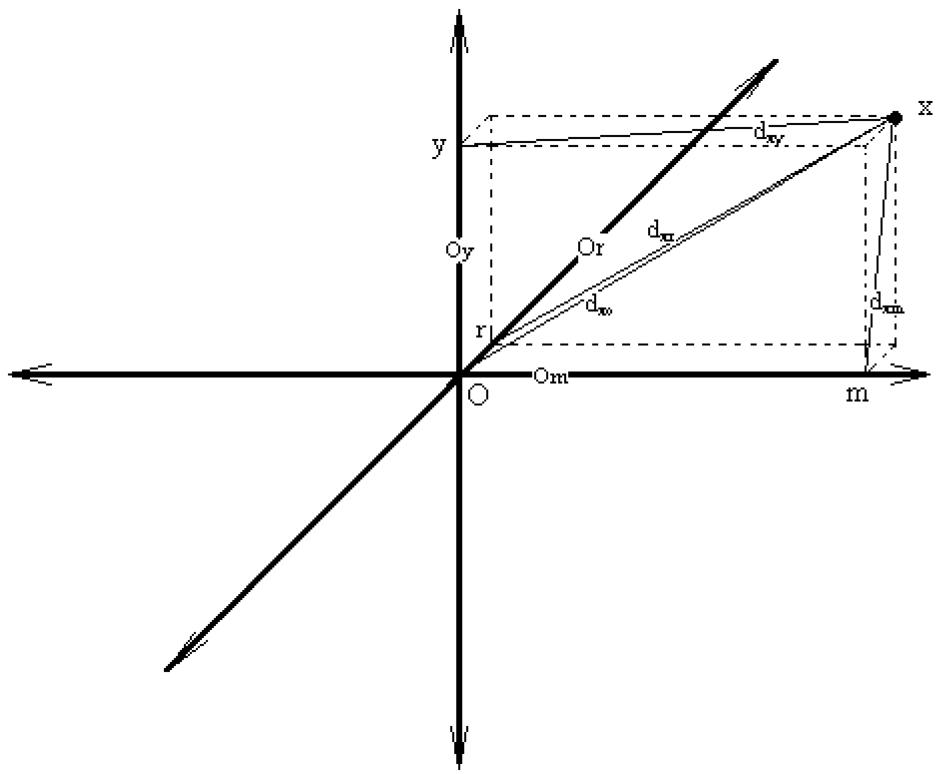

Figure 3.3.: One can individuate a particular x if one knows a) the distances d_{xy}, d_{xm}, d_{xr}, d_{xo} (represented by ‘$Pxyd_{xy}$’, ‘$Qxmd_{xm}$’, ‘$Rxrd_{xr}$’, and ‘$Sxod_{xo}$’); b) how o, y, m, and r are related with each other (at a time); and c) not all four of them are on the same plane.

Above I said that ‘$Pxyd_{xy}$ & $Qxmd_{xm}$ & $Rxrd_{xr}$ & $Sxod_{xo}$’ would individuate a particular if one also *knows* how o, y, m, and r are related with each other (at a time). But note what we get when we specify how o, y, m, and r are related with each other and their respective distances to x: we get a configuration, much like the one I have used above (the only difference is that the current one has five (when x is included), as opposed to four members). The upshot is that the coordinate system approach devolves into the approach that I pursue in this chapter or is so close to it that it doesn't really matter which one is picked.

3.D. Definite descriptions, narrow content, and representing particulars:

I have argued that whereas ADA cannot account for representations of particular objects, it can account for representations of super-determinate configurations, and, when supplemented with the first-person and temporal indexicals, it can account for representations of tropes like the pyramidal ones above. I will now argue that to represent particular objects, more is needed than just such representations, namely: logical operators sufficient to build definite descriptions. I will also argue that with respect to the content of indexicals one needs to appeal to a) theories of narrow content; b) singular causal relations. Before I present my proposal in Section 3.D.II.1 of how objects are represented, I will provide the theoretic background that underpins it in Section 3.D.I. In Section 3.D.II.2 I simplify it by pointing out how some of the symbols appealed to in Section 3.D.II.1 can be left implicit and dropped. In Section 3.D.III I apply the proposal to the problem of tracking particular objects.

3.D.I. Theories of content:

What my proposals boil down to is that the reference of a representation of a particular object is determined through a definite description. The definite description contains the predicate 'C_j', which, as explained above, refers to my pyramidal tropes in virtue of being functionally related to indexicals like 'I' and 'now'. Importantly, I am not making the claim that the definite description is the *sense* of the representation of a particular object. Thus, while I will say that tokens of 'C_j', 'I', and 'now' *causally* mediate the tokening of a representation of a particular object, I will not claim that the definite description containing 'C_j', 'I', and 'now' is *constitutive* of the content of a representation of a particular object.

That the reference of a representation of a particular object is determined through a definite description means that my proposal is not Russellian in that Russellian definite descriptions are quantifier expressions and as such do not refer (even under the guise of ordinary names). Russell's definite descriptions contain constituents with whose referents we are acquainted, where to be acquainted with an entity is to refer to it in a manner unmediated/unconstituted by the meanings of descriptions or of other concepts[71]:

> All propositions intelligible to us, whether or not they primarily concern things only known to us by description, are composed wholly of constituents with which we are acquainted. (Russell (1917, p. 128))

[71] *Cf.* Russell (1917, p. 108): "I say that I am *acquainted* with an object when I have a direct cognitive relation to that object, *i.e.* when I am directly aware of the object itself".

According to Russell, the only terms with which we are acquainted are sense-data, universals (which for Russell are concepts)[72], and ourselves. Like Russell's definite descriptions, my definite descriptions contain terms whose referent is not mediated/constituted by the sense of a description: the predicates 'C_j' that refer to the pyramidal tropes.[73] However, whereas for Russell concepts refer to types of sense-data, in my case predicates refer to external entities.

One can distinguish Russell's definite descriptions from Fregean definite descriptions[74] where none of the terms refer directly. Jeshion calls the theory that subsumes such descriptions "Pure Fregean Descriptivism":

> Pure Fregean Descriptivism is the theory according to which all aboutness – all linguistic reference and all intentional contents – is secured by means of an individual "fitting" or "satisfying" or "being in the extension of" a certain mental representation of it… A

[72] Universals for Russell are predicates that have in their extension sense-data. *Cf.* Russell (1917, p. 111): "Awareness of universals is called conceiving, and a universal of which we are aware is called a concept. Not only are we aware of particular yellows, but if we have seen a sufficient number of yellows and have sufficient intelligence, we are aware of the universal yellow…. And the universal yellow is the predicate in such judgments as "this is yellow," where "this" is a particular sense-datum."

[73] I pointed out before that I would be sympathetic to a view that takes definite descriptions to determine the content of 'C_j'. Importantly, this would *not* mean that the content of 'C_j' would be mediated/constituted by the sense of the description—that a description determines a content does not mean that the description has to be part of the content.

[74] I am ignoring here the question whether Frege himself thought of his definite descriptions in this manner.

singular term possesses <u>*purely* descriptive semantic content</u> just in case the reference of all of its contained terms is determined satisfactionally. (Jeshion (forthcoming, p. 5))

In the definite descriptions that I use there are terms whose referent is *not* determined entirely by their sense/character/role. These terms are the predicates ‘C_j’ that refer to the super-determinate pyramidal tropes. ‘C_j’ refers to a trope not only because of being related to ‘I’ and ‘now’, but also because it is *nomically* related to pyramidal super-determinate configurations. Thus, it does not possess a purely descriptive semantic content. In this respect my definite descriptions are unlike Fregean definite descriptions.

That the reference of a representation of a particular is determined through a definite description makes my proposals sound Fregean (for Frege, definite descriptions are the senses[75] of terms and they determine the terms' referents (if there are such)) or Kaplanian (for Kaplan the rigidifier ‘dthat’ has as referent the denotation of a definite description[76]). Importantly, they are not Fregean: on Fregean views the sense of a term (“something like a description in purely qualitative terms”[77]) is its meaning (propositional content), whereas in my account of how particular objects are represented, the meaning of a representation of a particular object is only its referent. This makes my view Kaplanian, since for Kaplan the meaning of a demonstrative is

[75] The sense of a term is its mode of presentation. The referent of a term is what the term is about.

[76] *Cf*. Kaplan (1989, p. 521).

[77] *Cf*. Kaplan (1989, p. 485).

merely the referent and not the sense of the description that picks out the referent. Below I briefly explain my reasons for taking a Kaplanian line:

One reason for not treating the meaning of representations of particular objects as (pure) Fregean sense is to avoid the problem raised by Strawson that if meaning were just Fregean sense, then we would not be able to distinguish between particular objects in massively qualitatively identical (real world) duplications. We can use as rich a description as we want, but the possibility remains that the universe contains another particular object that fits it just as well. Similarly, if meaning were just Fregean sense, then the terms of one definite description would be defined through other definite descriptions whose terms would be defined (directly or mediated by other definite descriptions) through the first definite description. This gives us only information about how the terms' referents are functionally related, but, however complex the functional relation is, it is always possible that more than one set of entities can satisfy it. Thus, the terms' referents would not be individuated.

As I said above, I take a Kaplanian line with respect to representations of particular objects: representations of particular objects have a character that determines a referent without the character being narrow content. The reasons for taking a Kaplanian line contra Frege and contra 2-factor theorists of content who argue that a representation of a particular object has both broad and narrow content are purely strategic. In this dissertation I just want to make the points that: a) we need the logic of definite descriptions to represent particular objects; and b) the content of representations of particular objects can be accounted for by appeal to representations of spatio-temporal tropes. If I appeal for a 2-factor theory of content, then I can make the points (a) and (b) above, but also (c): the *content* of representations of spatio-temporal tropes is *constitutive* of the content of representations of particular objects. While I do plan to defend this view in the future, I do not need to appeal to it to defend (a) and (b).[78]

[78] One reason for appealing to a 2-factor theory of content for representations of particular objects is that visual demonstratives can be tokened when there is nothing in the visual field (as with Macbeth's dagger). Here, appeal to narrow content would help explain the content of such hallucinations. Were I to appeal to a 2-factor theory of content, then my definite descriptions could account for cases where there is no object in the visual field to serve as the referent of 'x'. In this case some property is misrepresented as the property of being a Spelke object (the latter is represented by a predicate in the definite description). My definite descriptions could even account for cases where there is no space-time location in the visual field to serve as the referent of 'x' (image yourself on the edge of the universe looking out). In this case the 'C_j' predicate is mis-tokened: some property is misrepresented as the property of being a pyramidal space-time trope even though there is nothing (not even a space-time point) being the tip of the pyramid. However, it is not obvious that Kaplanian accounts cannot deal with cases of hallucination. They can hold that in cases of hallucination 'x' is an empty term and there is no singular thought (i.e. Russellian proposition) (*cf*. Corazza and Whitsey (2003)). The question is whether this will suffice to explain the rationality of the intentional states and behavior that results from the purportedly non-existent singular thought. It may: one strategy is to appeal to the character of the empty terms to explain their cognitive significance (but see Recanati (1990, p. 707-712) for criticism of such a move when the first-person indexical is concerned). I leave this discussion for another time.

As I said above, my definite descriptions are Fregean in the sense that they specify the referent of a representation of a particular object (this is what Kaplan's definite descriptions also do for the referent of a 'dthat' term). But there are famous arguments against Fregean definite descriptions. In what follows, I briefly address these objections:

Two standard objections against Fregean definite descriptions are Perry's perspectivalist argument and Strawson's duplication argument: Strawson (1959, p. 20) points out that Fregean definite descriptions would not be able to distinguish between massively qualitatively identical (real world) duplications. "DK" can be as rich a description as one wants, but the possibility would remain that the universe contains another individual that fits it just as well as I do—one just needs to imagine that the universe is like a set of monitors each one of them receiving the same broadcast so that what goes on in our sector of the universe is mirrored by what goes on in the other sectors. This is due to the fact that any Fregean description leaves out the sort of information contained in thoughts with indexicals and demonstratives—anchoring information that would pick out an entity, time, and place. I call this Strawson's "duplication argument". The need for indexicals in definite descriptions is demonstrated also by Perry (1979):

The propositional content of Fregean definite descriptions is absolute in the sense of not being relative to a particular person at a time-place. Perry (1979) targets this claim by arguing that, by being absolute, Fregean definite descriptions would suffice only to provide an objective representation of an object that is independent of a subject's perspective on it. This means that any Fregean description would leave out the sort of information contained in thoughts with indexicals. On the Fregean picture the thought "I need to go shopping" is short-hand for the thought "DK needs to go shopping" where DK is a definite description that applies only to me. The problem raised by Perry is that the thought "DK needs to go shopping" will cause me to go shopping only if I also believe that *I* am DK. That means that indexicals are essential—if the Fregean picture is to explain cognitive significance, it needs to appeal to indexicals. I call this Perry's "perspectivalist argument".

I agree with Perry that indexicals are essential and that is why they are embedded in my definite descriptions.[79] But my proposals would still suffer from the perspectivalist argument if the indexicals embedded in them are themselves Fregean definite descriptions in disguise (because these descriptions would not be specific enough to pick out a particular referent by the duplication argument). But if it is denied that the indexical are definite descriptions in disguise, then my proposals will be trivially immune to the perspectivalist argument. That is indeed my strategy in avoiding it: the narrow content of an indexical (its psychological character) is not a

[79] In that my descriptions are similar to Chalmers' primary intensions with their ordered triples of a world, an individual present in that world, and a time in that world. Schiffer (1981) also proposes to avoid problems like the one raised by Perry and Strawson by using definite descriptions *together* with the indexicals "I" and "now". Schiffer's definite descriptions are analogous to the approach of my proposal where the referents of representations of particular objects are specified through definite descriptions together with the indexicals "I" and "now".

definite description defined in purely qualitative terms, but the indexical's occurent inferential/functional role. Rey defines this role as follows:

> α is a FPRT [a first-person reflexive term] for agent x iff: (1-1) Whenever an input φ is received, x stores 'φα'; (1-2) Whenever x is in a mental state M, x is prepared to comp-judge a predication 'φα' that (ordinarily) gets released only when x is in M; (1-3) All preference states, and all basic action descriptions in x's decision system that lead up to action in a standard decision theoretic way are states and descriptions whose subject is α. (Rey (1997, p. 291))

The occurent inferential/functional role of an indexical qualifies as narrow content because it explains the generalization that subsumes my Twin Earth twins and me—that even though our token indexical thoughts have different referents, these tokens have the same type[80]—they are all about *our* selves[81].

Because of α's role, it is guaranteed that when DK tokens the thought "α needs to go shopping", DK will go shopping because α is the subject of DK's action descriptions (by 1-3) and α means DK. Were DK to token the thought "β needs to go shopping" (where β also means DK, but is not the subject of DK's action descriptions), then DK would not go shopping because DK does so only when α is the subject of DK's action descriptions. This deals with the perspectivalist argument, but

[80] *Cf*. Carruthers & Botterill (1999, p. 131-132).

[81] *Cf*. Rey (1998, p. 446).

leaves the question open as to why ‘α’ refers to *me* and not to a duplicate of me in a massively qualitatively identical duplication in the universe:

Strawson’s duplication problem is applicable to my proposal: I appeal to conceptual role semantics, in particular, Rey’s proposal, to define the *narrow* content (occurent inferential/functional role) of indexicals and an indexical’s role is mirrored not only in massively qualitatively identical duplications, but across people as well. However, because a token indexical can affect only the states of the token system within which it plays its role, the *broad* content (referent) of a token indexical would be the token system within which the token indexical plays its role. A token indexical plays its role in virtue of entering in *singular* causal relations with other tokens in a token system. So, a thought like “I am hungry” will be about me because the token ‘I’ a) has a specific inferential/functional role; and b) enters in a singular causal relation with me and not with a clone of me in a massively qualitatively identical part of the universe. That a particular system is identified in virtue of singular causal relations is very important because it means that appealing only to nomic causal relations would not help me explain how representations of particular objects get their content (the nomic relations (including the ones that implement the inferential/functional roles of symbols) will be the same for indexicals in all massively qualitatively identical duplications).[82]

[82] Note that we do not get the “which link” problem for first-person indexicals because we can just stipulate that the referent of a token indexical is the last link in the causal chain leading to the token—this being the system tokening it.

This concludes my discussion of Perry's perspectivalist argument and Strawson's duplication argument. In what follows I argue that the other well known arguments against Frege's descriptivist proposal do not apply to my own proposals:

Kripke argues that the descriptive notion of narrow content cannot account for the meaning of singular terms. His argument is broken down in three parts: the ignorance and error, rigidity, and unwanted necessity problems.[83]

The ignorance and error problem shows that associating a definite description with a term is neither necessary nor sufficient for the term to refer to the denoted referent: suppose that the only thing I know about Einstein is that he is a physicist. But [a physicist] is not a definite description and can be satisfied by more than one person, even though, purportedly, I can successfully refer to Einstein. Thus, it is not necessary for a term to refer to its referent that it be associated with a definite description. Furthermore, suppose that I associate [the physicist who invented the nuclear bomb] with 'Einstein' due to an error. Oppenheimer might be given the credit, not Einstein: when I utter, "Einstein is the physicist who invented the nuclear bomb", I utter something false. This means that even though the definite description identifies an individual, the name that the definite description is associated with might

[83] *Cf.* Kripke (1980, p. 3-15, 48-78). For discussion c*f.* Devitt & Sterelny (1999) and Reimer's "Reference" in *Stanford Encyclopedia of Philosophy*.

not refer to the same individual. Thus, associating a definite description with a term is not sufficient for the term to refer to the referent denoted by the description.

To see how the ignorance and error problem might (but does not) apply to my proposals consider visual indexicals (Pylyshyn's FINSTs). They refer to particular objects. One could argue that a) a visual indexical can refer to a particular object even if no definite description is associated with it; and b) a visual indexical can refer to a particular object even when it is erroneously associated with a definite description that refers to another particular. Thus, associating a definite description with a visual indexical is neither necessary, nor sufficient for the indexical to refer to the referent denoted by the description. The difficulty with such claims, however, is that they are so far just *hypotheses* and as such are in competition with hypotheses that deny the statements in (a) and (b) (this is what I do in my proposals). In the case of natural language, there are strong *natural language intuitions* that a) one can successfully refer to an individual, even when one has no definite description to associate with the individual; and b) the definite description associated with a name need not refer to the same individual the name does. Such natural language intuitions have to be accounted for and Fregean definite descriptions fail to do so. But with regards to the representations that connect the mind with the world there are no such intuitions to rely upon. Instead, the available empirical data can be interpreted in multiple ways. As a result, the application of the ignorance and error problem to my definite descriptions would be simply question-begging.

Two other problems that Kripke raises for the descriptive view are the rigidity and unwanted necessity problems. Kripke's rigidity problem stems from the fact that descriptions like [the most famous student of Plato] are not rigid: they can pick different people in different worlds. Names, like 'Einstein', however, are rigid: they pick the same person in every possible world in which that individual exists, whether or not the person in that world has all the properties that he has in the actual world. The reason for treating descriptions as non-rigid is to explain our intuition that it is true that Dionysius II could have been the most famous student of Plato. The reason for treating names as rigid is to explain our intuition that it is true that Aristotle could have been the *second* most famous student of Plato. Because of their different modal characteristics, this shows that names are semantically different from definite descriptions.

Similarly, Kaplan (1989, p. 518) argues that demonstratives should be rigid. Consider the sentence:

(1) He is the male being demonstrated at time t_1

It is possible that (1) is false. Yet, if Frege is right and the content of 'he' is "the male being demonstrated at time t_1", then (1) would express a necessary truth. Since it is possible that (1) is false, then the content of 'he' is not "the male being demonstrated at time t_1". That is, 'he' refers rigidly and its reference does not change across circumstances of evaluation.[84]

Prima facie, Kaplan's points would seem to apply to my proposals, since in my proposals the representation of a particular object is a visual demonstrative whose content is specified by a definite description. Kaplan shows that demonstratives are rigid designators—that is the content of a demonstrative does not change across circumstances of evaluation. This would suggest that that my descriptions would need to be rigidified in the manner in which Kaplan's dthat rigidifies demonstrative descriptions (p. 521). That is—the demonstrative representation of a particular would need to be treated as Kaplan's 'dthat'. However, in the early visual system, with which I am concerned, the rigidity problem does not apply, since the language of the early visual system deals with occurent perception and not with possible situations.[85]

---

[84] Kaplan also has a further argument (p. 513): the utterance of "he lives in Princeton" while pointing at Paul who indeed does lives in Princeton expresses the proposition p: "Paul lives in Princeton at time t_1". When evaluating the truth of p in counterfactual scenarios we are concerned with whether p is true. On the Fregean account, the utterance expresses the Fregean proposition q: "the male being demonstrated at time t_1 lives in Princeton". We can imagine counterfactual scenarios where p is true but q is false. This means they are not the same proposition and that in turn entails that the utterance above expresses a Russellian proposition (it is directly about Paul) because any Fregean proposition (that is indirectly about Paul) would not be the same as p by the above argument.

[85] Evaluations of counterfactuals is done later than tracking as well—possibly at the stage at which action schemata are produced. *Cf*. Carruthers (2006, p. 141-148).

That is to say that the language of the early visual system has no modal operators and as a result the rigidity problem does not apply here.[86]

A similar move works for addressing Kripke's unwanted necessity problem. The crux of this problem is that if the sense of 'Aristotle' is [the most famous student of Plato], then "Aristotle is the most famous student of Plato" should sound trivial. The reason it doesn't is because Aristotle could have been the second most famous student of Plato. This means that [the most famous student of Plato] cannot be the content of 'Aristotle'.

Kaplan (1989) shows how the problem applies to indexicals, as well: on Frege's account 'I' has as a sense "the speaker who is speaking". However, this would mean that utterances like:

(2) It is possible that I am not speaking

[86] Note that in my account of how particular objects are represented, I *am* in fact appealing to the Kaplanian 'dthat' analysis (for strategic reasons, as I explained above): according to my view a definite description specifies the content of a representation of a particular object without being constitutive of the content of the representation of a particular object. However, it is important to point out that Kaplan's argument does not work for the early visual system. This would allow me in the future to defend the view that a definite description *is* constitutive of the content of the representation of a particular object.

are always false, since “It is possible that the speaker who is speaking is not speaking” is always false. Which is false—(2) is true. Thus, Kaplan argues that “the speaker who is speaking” is not synonymous with ‘I’.[87] It can express the semantic rule that fixes the content of ‘I’ in a context (the character of ‘I’), but not the content of ‘I’. Then, the truth of (2) can be accounted for by fixing the referent in the actual context (the speaker) and evaluating (2) in circumstances of evaluation (possible worlds) where the referent does not speak.

But again—in the early visual system, with which I am concerned, the unwanted necessity problem does not apply, since the language of the early visual system deals with occurent perception and not with possible or counterfactual situations.

I have briefly addressed the standard objections to Fregean definite descriptions and explained why my account does not suffer from them. I now turn to the specifics of my proposal for representing particular objects:

3.D.II. Representing particulars through definite descriptions:

The essence of my proposals is that particular objects are represented through definite descriptions. Since there is evidence that particular objects are represented in the *early* visual system (Pylyshyn’s FINSTs), the definite descriptions must not

[87] See also Perry (1977).

incorporate concepts that only appear in higher cognitive capacities. I will now consider two proposals where different cognitive capacities are required for the representation of particulars. The first proposal is more intuitive for the theorist, but, as I will go and show in the second proposal, we can represent particular objects without some of the cognitive capacities the first proposal requires. It is the second proposal that I believe works better, if only because of its simplicity.

3.D.II.1. First proposal for representing particulars:

This proposal is more intuitive for the theorist because it appeals to *explicit* representations of a time and an observer: crudely put, given a time and an observer, only one particular object can enter in a super-determinate configuration with the observer's eyes and nose. In the next proposal, I will show that the representation of an observer can be made implicit. But first things first—take the following definite description:

$$(\boldsymbol{E}x)((y)(C(y, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{ry}, d_{ly}, d_{ny}, t, I) \leftrightarrow (x = y)))$$

To explain this definite description remember, from Section 3.C.II.2, that the predicate '$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$'[88] refers to a trope: being a configuration at time t such that: a) it involves the objects x, l, n, and r forming a pyramid where the respective distances are d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx}; and b) r, l, and

[88] I will refer to it as "'C_j'" for brevity's sake.

n form the base of the pyramid and r and l are my eyes, while n is my nose. In this super-determinate configuration x is at its non-base vertex and there can be only one such x. Note that this definite description can be modified to include a representation of the property S of being a Spelke object (bounded, unified, and persisting through time)[89] as follows:

$(\boldsymbol{E}x)((y)((C(y, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{ry}, d_{ly}, d_{ny}, t, I) \& Sy) \leftrightarrow (x = y)))$

The whole definite description then reads as follows: *right now* x is the *one and only* particular object that has the relational spatial property of being at the non-base vertex of the super-determinate pyramidal trope $C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$. Here is an illustration:

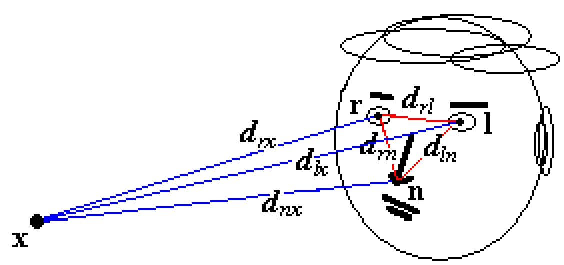

Figure 3.4.: This is Pinocchio looking at a particular object (x) in front of him. The particular object that Pinocchio is looking at is the one and only particular object that enters in a configuration with the left eye (l), the right eye (r), and the nose (n) of Pinocchio at time t, so that the respective distances are d_{rx}, d_{lx}, d_{nx}, d_{rl}, d_{rn}, and d_{ln}.

89 *Cf.* Spelke (1990).

This is reflected in the form of the definite description above: only one particular, given a time and an observer, would satisfy this relational spatial property.

In this definite description the symbol 'I' is a first person reflexive term. The variable 't' is a temporal indexical. The variable 'x' has the inferential/functional role of an argument[90].

In virtue of using the definite description above, the ontological assumption that only one entity can be at the tip of the pyramidal trope is implicitly represented.

---

[90] As to how this role may be specified, one may look to Strawson (1959, p. 167-170) and Strawson (1974, p. 103). I give his view merely as an example, while acknowledging that there may be problems with it: Strawson notes three ways that subjects and predicates differ from each other. First, while different predicate instances can be concatenated with one subject (e.g. 'Socrates' smile' and 'Socrates' frown' can be concatenated with 'Socrates'), only one subject can be the subject of a predicate instance (e.g. 'Socrates' smile' can be concatenated only with 'Socrates'). Second, with regards to predicates there are incompatibility groups, whereas there are no such groups with regards to subjects. An incompatibility group is such a range of predicates that if a predicate out of that range is concatenated with a subject, no other predicate out of this range can be concatenated with the subject. There are no such incompatibility groups with regards to subjects: a subject being the subject of a predicate does not preclude any other subject being the subject of that predicate. Third, with regards to predicates there are involvement relations, whereas there are no such relations with regards to subjects. Crudely put, an involvement relation is a relation between predicates such that one implies the other, but not *vice versa* (e.g. the determinate/determinable relation: being scarlet implies being red, but not *vice versa*). More precisely, there is a predicate P such that, for any subject x, Px entails Dx, where D is P's determinable. There are no such involvement relations with regards to subjects: there is no subject x such that, for any predicate P, there exists y such that Px entails Py.

If a definite description was not used, then the representing system would implicitly represent the world as a world where more than one entity can be at the tip of the configuration, which would mean that a particular entity would not be individuated.

In virtue of being embedded in the definite description as it is, ‘x’ represents a particular. But what makes ‘x’ not a mere representation of a particular but of a particular *object*? After all, tropes are particulars and so are states of affairs. Before I present what I think is the right answer, let me go through several answers that I do not think work:

One could argue that what makes the content of the argument ‘x’ different from that of the predicate ‘C_j’ (which refers to a trope) is that ‘C_j’ enters in asymmetric nomic relations with properties, while ‘x’ doesn’t. But we could easily imagine a situation where both enter in asymmetric nomic relations. One could instead argue that ‘x’ and ‘C_j’ have different contents because they get their content in different ways: ‘C_j’ does not enter in definite descriptions as an argument and as a result gets its content only through asymmetric nomic relations and inferential relations with indexicals. The argument ‘x’, on the other hand, enters in definite descriptions as an argument and as a result gets its content *also* through definite descriptions. This means that whereas the content of ‘C_j’ is constant (or robust), the content of ‘x’ is not—it can change based on what the definite description is. But that leaves it open that ‘x’ can refer at different tropes (or states of affairs) at different times—it doesn’t have to refer to different objects. Finally, one could argue that arguments relate to predicates in a way that *mimics* how objects relate to properties (as Strawson points out). Then, because ‘x’ is an argument, ‘x’ gets to have as content objects and not mere tropes, whereas because ‘C_j’ is a predicate, it gets to refer at most to tropes. I am actually somewhat sympathetic to this last approach. It is similar in structure to the point that because syntactic relations mimic semantic inferences the former can be taken to *mean* the latter. However, I still think that the approach is ultimately a problematic one because it seems plausible that other things than objects and properties can enter in relations identical to those between objects and properties. Pointing out that we can at least treat predicates as referring to properties (by appeal to ADA), and so eliminate some of the entities that enter in identical relations to those between objects and properties, does not resolve this problem because it seems possible that properties and other things than objects can enter in relations identical to those between objects and properties.[91]

I think the right answer as to why ‘x’ refers to an object and not merely to a trope (or state of affairs) has to do with the complex role of ‘x’ and in particular its relation to memory (see also Section 3.D.III.2): ‘x’ gets its content through definite descriptions. The definite description is stored in memory in a file (an “object file” similar to the one that Kahneman and Treisman (1992) propose). In the cases where the definite description that causes the tokening of ‘x’ changes from time t_1 to t_2 while the same token ‘x’ keeps being tokened, then the definite description is updated in the

[91] For example, note that Strawson’s three different ways that subjects and predicates differ from each other can be mimicked both by objects and properties and by states of affairs and properties.

object file. The previous description is not deleted but stored. This fact means that ‘x’ keeps its inferential relation to it. It is ‘x’ having both of these relations that explicitly represents that ‘x’ refers to the *same* particular at t_2 that it referred to at t_1. We have ‘x’ be related with two definite descriptions both of which define its content. That is—both definite descriptions individuate the *same* entity under different modes of presentation. Suppose that the entity in question were a trope. This would mean that the definite description containing ‘t_1’ and the definite description containing ‘t_2’ would refer to the same trope. But the definite description containing ‘t_1’ individuates a trope *at* t_1. Likewise, the definite description containing ‘t_2’ individuates a trope *at* t_2. These two tropes cannot be the same since a trope is a super-determinate at a spatio-*temporal* location—if the time is different, the trope is different. We have the consequence that the content of ‘x’ is two *distinct* entities: a trope at space-time location t_1 and a trope at a space-time location t_2. However, as stated above the content of ‘x’ as specified by the two definite descriptions has to be the *same*. Thus, ‘x’ cannot refer to particular that is a trope (or state of affairs[92]), but only to a particular that is an object. Thus, due to this updating and storing, ‘x’ represents the same particular object, and it represents it as being in different spatio-temporal relations relative to the observer.

[92] The same point applies for states of affairs: a state of affairs is a particular object instantiating properties, where, importantly, spatio-temporal relations are among the properties being instantiated. This means that if the spatio-temporal location of an object is changed, the state of affairs changes as well.

This, of course, leaves the problem of giving an account of how predicate logic, 'I', 't', and the numbers over which 'd_{rl}', 'd_{rn}', 'd_{ln}', 'd_{rx}', 'd_{lx}', and 'd_{nx}' range are instantiated in the brain. Here, at most, I can give a general sketch of an approach—an appeal to conceptual role semantics and singular causal relations—since accounting for them in detail would be beyond the scope of this dissertation.

Appeal to conceptual role semantics has been a favored method in explaining the meanings of logical operators. Take for example '&': the inferential role of '&' is such that there are dispositions to infer 'α' from 'α & β', 'β' from 'α & β', and 'α & β' from 'α' and 'β'. Because of its inferential role, we can say that the meaning of '&' is that of conjunction.[93] If this is plausible, then perhaps the same approach would work for quantification operators, the implication operator, and perhaps identity. Working out the details is a future project, though. Right now I only want to point out a plausible direction in which the meanings of logical operators can be accounted for.[94]

With regards to 'I', Rey (1997) suggests that if we help ourselves to conceptual role semantics, we can use the inferential roles of 'I' to define its meaning.[95] According to Rey, 'I' or what he calls "first-person reflexive term" or "FPRT" is defined as follows:

---

[93] There are problems with this approach, though, which I will discuss in the next chapter (Section 4.D.II).

[94] Discussion of the merits and problems of conceptual role semantics follows in the next chapter.

[95] On a conceptual role semantics account, the meaning of a symbol is its inferential role.

α is a FPRT for agent x iff:

(1-1) Whenever an input φ is received, x stores 'φα';

(1-2) Whenever x is in a mental state M, x is prepared to comp-judge a predication 'φα' that (ordinarily) gets released only when x is in M;

(1-3) All preference states, and all basic action descriptions in x's decision system that lead up to action in a standard decision theoretic way are states and descriptions whose subject is α. (Rey (1997, p. 291))

One might be suspicious of the individuation of 'α' because intentional terms such as "φ", "M", "decision system" are invoked. However, on Rey's view, conceptual role semantics doesn't have to avoid *every* intentional property. It just has to avoid, for fear of vicious circularity, those intentional properties it is trying to explain through inferential roles. Here, Rey grants that accounting for the meaning of "φ", "M", "decision system" might involve appeal to relations with external properties (as on an externalist account of content).

On Rey's view of FPRTs, 'α' means "receiver of present inputs, the instigator of outputs, and the subject of intervening mental states". That is not enough for 'α' to individuate a particular subject, though—after all the same functional/inferential role is had by tokens of 'α' in other people's minds as well. However, because a token 'α' can affect only the states of the token system within which it plays its role, the *broad* content (referent) of 'α' would be the token system within which the token indexical plays its role. 'α' plays its role in virtue of entering in singular causal relations with other tokens in a token system. So, a thought like "I am hungry" will be about me because the token 'I' a) has a specific inferential/functional role; b) enters in a singular causal relation with me and not with a clone of me in a massively qualitatively identical part of the universe.

In this proposal I adopt something like Rey's account of first-person reflexive terms, while accepting that this requires me to justify my appeal to inferential roles and conceptual role semantics (which I do in the next chapter). I do not adopt Rey's account as it is, because it is highly implausible that the first-person reflexive term which figures in the visual system is the same symbol that Rey talks about above—the one that we use to attribute mental states to ourselves. If there is a first-person reflexive term in the visual system, then it has roles that relate it to other symbols in the early visual system and to actions in the early visual system, such as tracking, saccading to, and focusing on a target. Thus, we need an account more or less as follows:

'β' is a *visual* FPRT in an agent x's visual system v iff:

(1-1) Whenever the input ψ is processed only by the modules for retino- or egocentric space, v stores 'ψβ';

(1-2) Whenever the input ψ is processed also by the modules for allocentric space, v stores 'ψ';

(1-3) All basic action descriptions in v that lead up to action through computational processes[96] are descriptions whose subject is 'β'.

This proposal implies that there is an explicit representation of a visual FPRT. However, as I will shortly show in the second proposal below, we can do without any explicit representation of first-person indexicals.

Giving an account of 'now' and of how time is represented, is a topic that I plan to pursue but which right now is beyond the scope of this dissertation. Now I will just wave at the work of Roberts (1998) and Gallistel (2009) who present evidence for biological clocks in a variety of creatures. In analyzing how an animal succeeds in representing time, Roberts argues that a biological clock consists of "something that changes (*pacemaker*) and something that measures the changes (*register*)". Without getting too specific, one can at least see how one can go about giving an account of how temporal spans are represented: the value of the *predicate* 'T_i' is the number recorded in the register of the changes of the pacemaker.[97] Changes

[96] This qualification is needed because some of the processes of the visual system are non-computational. For example, selecting the most active input for processing does not require computation but only lateral inhibition.

[97] Perhaps the fact that 'T_i' ranges over numbers in a sequence can be used to represent the fact that time is directional. As I explain below, a set-theoretic definition of numbers is built on the successor function and perhaps the asymmetricity of the successor relation can be used to account for the asymmetricity of time. This could be used to treat 'T_i' not merely as a representation of temporal spans (durations), but as a more abstract representation of time. Since the visual system only needs to

of value of this number are nomically correlated with changes of the pacemaker. Now, if the changes of the pacemaker are nomically correlated with temporal changes[98] and assuming that nomic relations are transitive[99], then the change of value of 'T_i' is nomically correlated with temporal changes. Assuming that the temporal changes and the changes of value of 'T_i' are incremental, then the value of 'T_i' would nomically co-vary with the length of temporal spans. This would suggest that an externalist account like ADA which accounts for content through nomic relations can explain how temporal spans are represented. It would define the content of 'T_1' as that super-determinate *temporal* configuration consisting of two temporal points on

represent temporal spans and not time to compute velocities and to distinguish past from present instantiations of the same object, I would not pursue here the much more difficult question of how time is represented.

[98] Can we take it as obvious that the changes of the pacemaker nomically co-vary with temporal changes? If they don't, then 'T_i' doesn't represent temporal spans but just the number of changes of the pacemaker. Suppose that time can flow with different speeds in between the changes of the pacemaker. This would mean that the changes of the pacemaker do not nomically co-vary with temporal changes. If this is correct, then we would perhaps need narrow content to argue that, since the organism *uses* 'T_i' to predict an object's future position (given its position and velocity), then 'T_i' is a representation of perceived temporal spans (as opposed to just being a representation of the number of changes of the pacemaker). I leave this issue for another time.

[99] Armstrong (1997, p. 234-235) denies that non-superveneint nomic relations can be transitive but allows that supervenient nomic relations can be. Tooley (1977, p. 697) does allow for transitive nomic relations. Examination of this dispute is left for future research on the representation of time.

the line of time (a temporal span) with which 'T_1' enters in a nomic relation ('T_1' stands for a particular value of 'T_i')[100].

What has been said so far seems to cover how temporal spans get represented, not how 'now' (or 't') gets its content (and the content of 't' needs to be accounted for 't' plays a role in my definite descriptions). But based on the above, we are in a good position to explain how 't' means NOW: 't' is always concatenated with the current 'T_i'. 'T_i' is shorthand for a particular predicate (e.g. 'T_1') representing a particular temporal span. In virtue of always being concatenated with the *current* 'T_i', 't' represents the current time—i.e. it means NOW. The concatenation '$T_i t$' is like the moving arrow of a clock—it at once represents how much time has passed and what time it is now. If this is right, 'now' (or 't') would have the following inferential/functional role:

't' is a *visual* "now" in an agent x's visual system v iff:

(1-1) 't' is always concatenated with the current 'T_i' where 'T_i' is the number currently recorded in the register of the changes of the pacemaker;

(1-2) Whenever input ψ is received, v stores 'ψt';

Remember how the content of a particular token of 'I' was accounted for—both through functional/inferential relations and through singular causal roles. The

[100] Assuming, of course, that this nomic relation is explanatorily basic—it explains other nomic relation in which 't_1' enters, but not *vice versa*.

reason was that ‘I’’s functional/inferential role is had by tokens of ‘I’ in more than one person’s mind and so ‘I’’s functional/inferential role does not individuate a particular person. In the same way, ‘t’’s functional/inferential role is had by tokens of ‘t’ in people at different temporal locations. Thus, ‘t’’s functional/inferential role does not individuate a particular point in time. To do so we need to appeal to singular causal roles again: a particular token of ‘t’ enters in a singular causal relation with a particular register and pacemaker and this picks out an entity with a particular temporal location.[101] Thus, a particular token of ‘t’ has as its content the particular time of its tokening in virtue of a) its functional/inferential role that defines ‘t’ as a visual “now”; and b) a singular causal relation between the token of ‘t’ and a particular register and pacemaker that picks out a particular temporal location.

This leaves one to specify how numbers and identity are accounted for. This is another topic for research which cannot be explored here and about which I will only say the following: a set-theoretic definition of numbers is built on the successor function $S(k) = k + 1$.[102] If one were to show how the successor function can be defined in terms of inferential/functional roles, then an appeal to conceptual role

---

[101] This works only if the entity tokening ‘t’ has a temporal part—by picking out the entity tokening ‘t’ one also picks out its temporal part which in turn picks out a particular temporal location. However, the entity tokening ‘t’ might not have a temporal part, as on 3D accounts (where objects persist by existing at two or more different times and not by having temporal parts). Then picking out the particular entity tokening ‘t’ would not individuate a particular time. I reserve discussion of 3D accounts for future research.

[102] *Cf.* Peano’s axioms.

semantics would be a promising strategy in giving an account of numbers. Given that another function, such as the '&' function, is a paradigmatic example of how conceptual role semantics works, it is not implausible to believe that a successor function could also be provided in a similar manner[103]. So, if some symbols enter in asymmetric and transitive functional/inferential relations that can be described as successor relations, then these symbols would be promising candidates for representations of numbers. Perhaps something similar can work for the relation of identity (R): an inferential/functional relation between symbols can be described as R if a) the relation is reflexive, symmetric, and transitive; and b) if 'α' is R-related to 'β', then if 'Pα', then 'Pβ'.

3.D.II.2. Second proposal for representing particulars:

The first proposal requires that there be a complex representation tokened, complete with indexicals and quantifier, implication, and identity operators. Also, the first proposal depends on the indexical 'I' to individuate a particular human whose eyes and nose enter in a super-determinate configuration with the represented particular object. The reason 'I', as defined by Rey (1997), locks onto me is because 'I' a) has a specific inferential/functional role; and b) enters in a singular causal relation with me and not with a clone of me in a massively qualitatively identical part of the universe. But, perhaps, we can deal away with explicitly representing 'I'. The

[103] *Pace* Fodor (2004) who argues that conceptual role semantics cannot non-circularly account even for the logical operator '&'.

idea is that just like a token ‘I’ picks out a particular person because of singular causal relations with this person, so a token ‘C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t)’[104] picks out a configuration involving a particular visual system because of singular causal relations with this visual system. In this way, the system that is tokening ‘x’ and ‘C_j’ is implicitly specified (as opposed to explicitly, through ‘I’). The benefit of adopting this suggestion is that ‘I’ is no longer needed in the definite description, which is a plus when we are concerned with early representation—a stage of representation when it is dubious that the representing system has the capacity to token first-person reflexive terms.

Explicit representation of the visual system is needed for the capacity to represent objects in allocentric space, since such capacity allows one to distinguish between spatial relations that involve one from those that do not. Take, for example, a cat observing two mice separated by 5 m with the cat being 5 m away from one of the mice. Suppose that there is a nomic relation between the spatial relation [being a configuration involving two objects separated by 5 m] and a symbol ‘5 m away’. For the cat to be able to distinguish the spatial relation [being separated by 5 m] that involves it from the spatial relation [being separated by 5 m] that does not, then ‘5 m away’ must be related to a visual ‘I’ when the spatial relation involves the cat, and not related with the visual ‘I’ otherwise. However, for visual systems that represent objects only in egocentric space, the visual ‘I’ is not needed. Such systems do not represent spatial relations that do not involve the systems, and as a result they do not

[104] Note that ‘I’ has been removed.

need to distinguish between spatial relations that involve the systems from those that do not. Since now I am concerned only with egocentric space, I do not need a visual 'I' to be explicitly represented in the definite description that refers to particular objects.

One might think that we can also get rid of the 't': isn't the time of the 'C_j''s tokening enough to specify a time implicitly? However, getting rid of 't' would be plausible only if we were concerned with representing objects *at a moment*. But we are not—our visual systems are capable of distinguishing past from present instantiations of the same object. Likewise, explicit representations of time are needed to process the velocities of moving objects. This implies that objects' temporal coordinates are represented explicitly and this is a reason to keep 't' in the definite description. Furthermore, as I argued above, in order to explain why 'x' refers to an object and not merely to a trope (or state of affairs), one has to appeal to two different definite descriptions—one that contains 't_1' and one that contains 't_2'. That is, without the explicit representation of time in the definite descriptions, we would not be able to explain how particular *objects* are represented.

The upshot is that we are left with the following definite description:

$$(\boldsymbol{E}x)((y)((C(y, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{ry}, d_{ly}, d_{ny}, t) \,\&\, Sy) \leftrightarrow (x = y)))$$

This proposal is simpler than the first one in Section 3.D.II.1 and relies almost as much on conceptual role semantics: the only difference is that here we do not have 'I' and so do not need to appeal to conceptual role semantics for its meaning. But we still need conceptual role semantics to account for predicate logic, representations of numbers, and 't'.[105]

3.D.III. Tracking of particulars:

If we are concerned with representing particular objects *at a time* and if the above is on the right track, then we have a general sketch of how particular objects are represented. However, we are interested not only in representing particulars at a time but *through time*, as well. In tracking a moving object we obviously represent it as being the same one through time. That is to say that even though at time t_2 the object may have changed its spatial and non-spatial properties, we still represent it as

[105] Prior to Fodor (2004), Fodor assumed that ADA goes hand in hand with conceptual role semantics, where the latter *only* concerned itself with accounting for the logico-syntactic roles of symbols. An important implication of my proposals is that conceptual role semantics is needed to account for *more* than the logico-syntactic roles of symbols. The reason I am stressing this is to point out that my proposal does not treat representations of particulars as complex representations composed by constituents bound by logical operators where the contents of those constituents are accounted for by ADA. If that was my proposal, then I would have held that ADA, prior to Fodor (2004), would have been able to account for representations of particulars. In my proposals, however, the definite descriptions incorporate 'now' and representations of numbers, and those do not get their content merely through nomic relations with external properties. That is why I believe that ADA, prior or after Fodor (2004), would *not* be able to account for representations of particulars.

being identical to the one it was at time t_1. The "tracking problem" is to explain how this is achieved—by appealing to either implicit or explicit re-identification.

An explicit re-identification account is one where in virtue of an object representation 'x' at t_1 being related to an object representation 'y' at t_2', the two are treated *as* representing the same object. An example is where '=' is used to identify the referent of 'x' at t_1 with that of 'y' at t_2'. An implicit re-identification account is one where no relation between the two representations is taken to constitute them being treated as representing the same object. The most plausible example of such an account goes as follows[106]: if tokens of 'x' are causally connected with the object x from t_1 to t_2, then 'x' tracks x from t_1 to t_2.[107] Here we can say that 'x' tracks x not in virtue of a token 'x' at t_1 being connected with a token 'x' at t_2, but in virtue of the tokens being continuously causally related with x. However, the problem with such an account is that it cannot explain visual tracking through occlusion (a case where the visually tracked object is briefly occluded by an intervening obstacle). The reason is that while the visually tracked object is occluded, there is no causal chain connecting tokens of 'x' with x. This suggests that if we want an account of visual

---

[106] One could try to define implicit re-identification as "continuously referring to an object". However, I can continuously refer to Bob from t_1 to t_2 (by continuously thinking about him) without being able to recognize Bob at t_2, due to him drastically changing his appearance, for example. Thus, if we define implicit re-identification as above, it is not clear how we are re-identifying anymore.

[107] The need to use a *token* of 'x' in the antecedent of the conditional, as opposed to just the symbol 'x', is because I am assuming (see Chapter 1) that objects enter only in singular causal relations, where the latter have as relata particular objects.

tracking through occlusion, then we probably need to include some explicit representation of re-identification. Below I give an account of such explicit re-identification that is based on my account of how particular objects are represented.

First note that the problem of explicit re-identification is not solved merely by appealing to '=' to connect two representations of particular objects. The reason is that one still would have to explain why they are *treated as* identical in content with each other. Here one cannot simply say that they are *treated as* identical in content with each other, if they refer to the same object. The reason is that one can have two concepts that refer to the same object under different descriptions, *without* treating the concepts as identical in content with each other. We need a mechanism to explain how the system treats the content of 'x' at t_1 as identical to that of 'x' at t_2. I provide such a mechanism below:

To solve the tracking problem I suggest appealing to the definite descriptions in the two proposals above:

$$(\boldsymbol{E}x)((z)((C(z, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rz}, d_{lz}, d_{nz}, t) \& Sz) \leftrightarrow (x = z)))$$

Where 'y' and 'x' are representations of particular objects, the idea is to associate the referent of 'y' at time t_2 (y) with the broad content of 'x' at time t_1 (x), *if* at time t_2 y is closest to x at time t_1.[108] The metric for closeness piggy-backs on the predicates 'C(z, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rz}, d_{lz}, d_{nz}, t)' ('C_j'). For exposition's sake, let's call the variables in the definite description of 'x': {'d_{rx}', 'd_{lx}', and 'd_{nx}'}, and likewise for 'y'.[109] Then, how close to each other x and y are represented as being is reflected by the values of the variables in the following two sets: {'d_{rx}', 'd_{lx}', 'd_{nx}'} and {'d_{ry}', 'd_{ly}', 'd_{ny}'}. The distance between x and y can be computed based on these values. This distance then serves to ground the closeness relation. Then x is associated with y if the distance between them is the smallest when compared to the distances between x and other represented particular objects.

More specifically: at t_1 'x' is tokened together with 'C_j'. At t_1 'x' refers to the object denoted by the definite description that contains 'C_j' (let's call it: D-'C_j'). D-'C_j' is stored in the object file associated with 'x'. At t_2, we have 'C_k' being tokened as a result of a perceived super-determinate pyramidal configuration (where 'C_j' ≠ 'C_k'). 'C_j' together with 'C_k' are used to compute a distance in virtual space between the point defined by {'d_{rx}', 'd_{lx}', 'd_{nx}'} and the point defined by {'d_{ry}', 'd_{ly}', 'd_{ny}'}. If

[108] Some further condition is needed to specify a *small* time period within which such association is to take place. Particular objects represented at times whose temporal difference exceeds this time period, will not be associated with each other. The purpose of this condition is to eliminate most of the cases where two *distinct* particular objects get associated with each other: if the time period is sufficiently small, then, unless the particulars are moving *really* fast, the particular object at time t_2 that will be closest to the particular object at time t_1 will be the same particular object as that at time t_1.

[109] 'd_{rl}', 'd_{rn}', and 'd_{ln}' stay the same, being representations of the distances between the eyes and the nose.

this distance is the smallest when compared to the distances produced by other predicates referring to super-determinate pyramidal configurations, then D-‘C_j’ is updated with D-‘C_k’ in ‘x’’s object file. Now (at t_2) ‘x’ refers to the object denoted by D-‘C_k’. Importantly, when D-‘C_j’ is updated, it is not deleted, but stored in the object file. This fact means that ‘x’ keeps its inferential relation to it (on top of its relation to D-‘C_k’). It is ‘x’ having both of these relations that explicitly represents that ‘x’ refers to the *same* particular at t_2 that it referred to at t_1.

So, given this account, what happens in cases of occlusion during visual tracking? The answer may go as follows: at t_1 ‘x’ represents an object denoted by D-‘C_j’—x. D-‘C_j’ is stored in the object file associated with ‘x’. At t_2, when x is occluded, no ‘C_k’ is tokened and there is no D-‘C_k’. However, we still have D-‘C_j’ in ‘x’’s object file. As a result, ‘x’ still refers to x, in virtue of D-‘C_j’. At t_3 we have D-‘C_k’ which updates D-‘C_j’, where the latter remains stored. ‘x’ at t_1 is identified with x at t_3 not because ‘x’ keeps referring to x from t_1 to t_2, but because of ‘x’’s relation with both D-‘C_j’ and D-‘C_k’.[110]

3.E. Conclusion:

[110] Importantly, the above is an account of tracking that can deal with occlusions. I grant that we do not need explicit representation of re-identification for sensory systems that cannot deal with occlusions.

In this chapter I have argued that to explain how particular objects are represented a) the logico-syntactic roles of representations of particular objects need to be taken into account, and for this ADA needs to be complemented by conceptual role semantics; and b) an appeal to the temporal and first-person indexicals must be made, and for this ADA must be supplemented with singular causal relations. In particular, I have argued that a representation of an object 'x' gets its content *via* a definite description that denotes 'x''s referent without constituting 'x''s content. The definite description contains a predicate 'C_j' that refers to a super-determinate pyramidal configuration instantiated at an observer and a time—a trope. 'C_j' refers to this trope in virtue of a) a nomic relation *á la* ADA with the pyramidal configuration; and b) functional relation with the indexicals 'I' and 't'. The definite description containing 'C_j' reads as follows: right now x is the one and only particular that has the relational spatial property of being at the non-base vertex of the super-determinate configuration, where this configuration has its base positioned on my eyes and nose. This view makes use of conceptual role semantics for more than the logical apparatus needed for the definite description. It is also used to define the narrow content of 'I' and 't'. In this chapter I have also argued that particulars are tracked in virtue of an updating and storing of the definite descriptions that determine the content of representations of particular objects.

In these proposals I have appealed to conceptual role semantics to account for predicate logic, representations of numbers, 'I', and 't'. But conceptual role semantics has been held by many to be deeply problematic. So problematic that Fodor (2004) doesn't even think it can account for its paradigmatic examples—logical operators. I turn to a defense of conceptual role semantics in the next chapter.

Chapter 4: In Defense of Conceptual Role Semantics

4.A. Introduction:

In Chapter 1 I argued that a) the singular causal strategy for accounting for the conent of representations of particular objects faces the "which link" problem; b) the nomic strategy pursued in this dissertation is at worst as complex as the singular causal strategy. The upshot was that the nomic strategy also deserves consideration. In chapter 2 I presented an example of the nomic strategy—Fodor's ADA—and pointed out that it faces the problem of representing individuals. In the previous chapter I argued that ADA, as a nomic theory, when complemented with conceptual role semantics and singular causal relations for indexicals can account for how particular objects are represented and tracked. In this chapter I will discuss my use of conceptual role semantics: why I have chosen it instead of the other accounts of narrow content, and will argue that it has the potential to deal with the most serious objections thrown in its way.

Conceptual role semantics is not the only conception of narrow content that might be able to help nomic theories account for how representations of particular objects get their content. In Section 4.B I briefly overview several accounts of narrow content (the descriptivist, the mapping, and Devitt's one), and explain that I picked conceptual role semantics over them because it is able to help with the problem of representing particular objects better than the other accounts. However, there are two serious problems against conceptual role semantics: the holism and the analytic/synthetic distinction problem. Following Rey, I suggest that these problems can be dealt with if one appeals to asymmetric dependencies between symbol relations. This is presented in Section 4.C. Another problem raised against conceptual role semantics that I discuss in this section is the circularity problem.

4.B. Notions of narrow content:

According to narrow content theories, a representation can have a content that is not individuated merely by the representation's relation with external entities. This is in contrast with broad theories of content where a representation's content is defined only through causal relations with the external world. On narrow content accounts, the meaning of a representation is not to be equated with the representation's referent, whereas on broad theories of content the meaning of a representation *is* the representation's referent. For example, according to ADA, if 'pet fish' enters in asymmetrically dependent nomic relations with pet fish, then the content of 'pet fish' is simply its referent—the property of being a pet fish. Whereas, according to a (pure) conceptual role semantics, the content of 'pet fish' is not the referent, but the inferential/functional relations that ‘pet fish’ enters into with percepts in the sensory systems. In this section I overview some of the different notions of narrow content and some of the problems they face.

4.B.I. The mapping notion of narrow content:

An example of the mapping notion of narrow content can be found in Fodor (1987): the narrow content of a term is a function from contexts to contents[111]. Thus, for Fodor, a term’s narrow content is not equated with the term’s contents or truth-conditions. A term’s narrow content is what is left once the truth-conditions of a term have been abstracted away and what remains is the function only. The narrow content of a term is a “set of ordered pairs” (Fodor in Loewer & Rey (1991, p. 299)) of worlds and properties where in each world (context) a term’s (narrow) content picks the term’s extension (broad content). Thus, to specify the narrow content of a term one must specify all the ordered pairs. This means that a) a particular term’s narrow content is a “construct out of broad content” (Fodor in Loewer & Rey (1991, p. 301))—the broad contents of the term in each possible world define the term’s narrow content; and b) the individuation condition of a term’s narrow content is just the set of ordered pairs of contexts and contents—thoughts with identical sets have the same (narrow content).[112]

[111] As such, the mapping notion is essentially Kaplanian in character.

[112] *Cf*. Fodor (1987, p. 48)): “Your thought is content-identical to mine only if in every context in which your thought has truth condition *T*, mine has truth condition *T* and vice versa.”

So how helpful is Fodor's mapping notion in the project of accounting for the content of representations of particular objects? This question should be split into two parts: a) does Fodor's notion of narrow content help by itself; and b) can Fodor's notion of narrow content replace conceptual role semantics in the account given in Chapter 3? The answer to the first question is, "No". The problem of accounting for the content of representations of particulars is to explain *why* it is the case that in a given context a token symbol refers to a given object. Fodor's mapping notion can at most tell us that given a context a symbol will have a given referent, but it does not explain *why* it is the case that in this context the symbol has this referent.[113]

Can Fodor's mapping notion replace conceptual role semantics in the account given in Chapter 3? The answer is, "No". In my account of how particular objects are represented I argued that we need to appeal to logical operators. However, Fodor's mapping notion is of no help when it comes to defining the meanings of logical operators because logical operators have no referents (that is why Fodor, prior to his

---

[113] As Curtis Brown describes this problem in "Narrow Mental Content" in SEP: "although the mapping conception gives us an abstract, formal conception of narrow content, it does not give us an algorithm for *finding* the narrow content of a particular state. Although apparently any function from contexts to contents would count as a "narrow content" in Fodor's sense, some of these functions could not really be the content of a mental state. To use a computational analogy, we are really interested only in "computable" functions from context to content, functions that can be implemented somehow in a human mind, and this suggests that it is not the function itself that is of interest but rather the algorithm by means of which it is computed."

(2004) appeals to conceptual role semantics to account for the meanings of logical operators). This means that even if Fodor's mapping notion can account for the contents of indexicals and numbers, Fodor's notion would not completely suffice for all that I need narrow content to do for me in Chapter 3. Instead of adopting multiple notions of narrow content, considerations of simplicity suggest to stick with conceptual role semantics if it can help me with predicate logic, numbers, and indexicals, as I argued that it can in the previous chapter.[114]

4.B.II. The descriptive notion of narrow content:

This notion is most notable in the work of Frege. For Frege the sense of a term is a definite description which determines the referent of the term. Thus, the referent of 'water' is determined by its sense: [the clear, odorless, liquid that freezes at 0 degrees Celsius] and the referent of 'Barack Obama' is determined by its sense: [the current president of the USA].[115]

---

[114] A substantial problem for Fodor's notion of narrow content is that it, unlike the other notions to be discussed below, does not seem promising in solving the necessarily co-instantiated properties problem (the Gavagai problem). This, as Rey (1997) points out, generalizes to necessarily un-instantiated properties: an account, like Fodor's, would have to treat terms like 'free will' and 'magic' as identical in narrow content, since in every world they will refer to the empty set. The other notions of narrow content, however, do offer hope in dealing with this problem, which can be a reason for preferring them to Fodor's notion.

[115] Note that this notion of narrow content offers to deal with the problem of distinguishing between necessarily co-instantiated properties (e.g. rabbit-hood and being an undetached rabbit part) by

Does the descriptive notion of narrow content help by itself to account for how representations of particular objects get their content? The answer is, “No”, because of well-known problems which were reviewed in the last chapter. For example, Strawson (1959, p. 20) points out that Fregean definite descriptions would not be able to distinguish between massively qualitatively identical (real world) duplications. DK can be as rich a description as one wants, but the possibility would remain that the universe contains another individual that fits it just as well as I do.

Can the descriptive notion replace conceptual role semantics in the account given in Chapter 3? The answer is, again, “No”. In my account of how particular objects are represented I argued that we need to appeal to indexicals like ‘I’ and ‘t’. However, definite descriptions are unable to account for the content of such indexicals either by themselves or by being helped out by singular causal strategies. Definite descriptions cannot account for the content of ‘I’ and ‘t’ by themselves because of Strawson’s duplication argument. In the previous chapter I appealed to singular causal relations to account for the content of indexicals: because a token ‘I’ can affect only the states of the token system within which it plays its role, the *broad* content (referent) of ‘I’ would be the token system within which the token indexical plays its role. ‘I’ plays its role in virtue of entering in singular causal relations with other tokens in a token system. However, the same move is unavailable for the

---

specifying through descriptions whether the referent of ‘rabbit’ is the property of being a rabbit or the property of being an undetached rabbit part.

descriptive notion theorists because they abstract away from a token term's inferential/functional relations: a definite description specifies a term's, as opposed to a *token* term's, relations with other terms. Details about the *causal* interactions of the token terms are omitted. As a result, if one appeals only to definite descriptions, then one cannot also appeal to singular causal relations involving the *token* terms to determine the content of indexicals, the way I did in the previous chapter.

4.B.III. Devitt's notion of narrow content:

According to Devitt,

> narrow meaning (or content) is very rich. Not only does it include all the functional roles that determine the syntactic structure of sentences, but also the inner functional roles that partly determine the reference of words. The latter roles are what is left of wide word meaning when the extra-cranial links are subtracted. The roles constitute narrow word meanings. Those meanings are functions taking external causes of peripheral stimuli as arguments to yield wide (referential) meanings as values. *Narrow word meaning is (mostly) not a matter of syntax*. (Devitt (1990, p. 380))

Devitt provides the qualification "mostly" because almost all terms enter in more than syntactical relations with other terms—they enter in functional relations with proximal stimuli and states of the organism. The syntactic and functional relations are included in a term's narrow content only if they are instrumental in determining the term's reference given a context.[116] Likewise, descriptions (that are instantiated *via* inferential/functional relations) can also be part of a term's narrow content—as long as they play a part in determining the term's wide content. So, Devitt & Sterelny's idea is that the narrow content of a term is specified by the relations in which the term enters, including those relations holding between the term and proximal stimuli. Once the context is provided—the distal causes causing the proximal stimuli—the narrow content determines a referent: the cause of the proximal stimuli. If a description is used to categorize the referent, then the description is also part of the term's narrow content.

Devitt divides terms into three kinds: "pure-causal terms", "descriptive-causal terms", and "descriptive terms".[117] The first are terms whose narrow content involve only "the functional-role connections between peripheral stimuli and the term". For pure-causal terms, these connections, given a context, suffice to determine a referent. The descriptive-causal terms require a richer narrow content to determine, given a context, the referent: their narrow content involves a) links between the terms and stimuli; and b) associations with other terms. The narrow content of descriptive terms involves only associations with other terms and its broad content is determined through the broad contents of these other terms.[118]

---

[116] *Cf*. Devitt & Sterelny (1999, p. 210): "Where then do we differ from two-factor theorists? The only functional roles that go into our meanings are ones that determine reference."

[117] *Cf*. Devitt (1990, p. 380).

[118] Devitt's view is attractive because it avoids the holism charge levied at most functionalist accounts of not being able to define in a principled manner the meaning-constitutive roles of a symbol: on

Does Devitt's notion of narrow content help by itself to account for how representations of particular objects get their content? The answer is, "Almost". Devitt's descriptive-causal terms and descriptive terms allow for a rich narrow content. This means that, if we ignore for a moment the logical operators to which I appeal in my definite descriptions, my definite descriptions can be treated as the narrow content of one of Devitt's terms that determines the term's reference—much like it is on my own account. However, there is an important difference between Devitt's notion of narrow content and the conceptual role semantics notion that I appeal to: Devitt's notion of narrow content is tied to a referent—no referent, no narrow content—whereas the conceptual role semantics notion is not so constrained. It is this difference that explains the answer "almost" above and explains why I chose conceptual role semantics instead of Devitt's notion of narrow content:

My definite descriptions appeal to logical operators. Logical operators do not seem to have referents and thus none of the roles of logical operators determine any. Obviously then, Devitt's notion of narrow content, which determines, given a context,

---

Devitt's account the syntactical/inferential/functional roles which comprise the narrow content of a symbol are those instrumental in determining the term's referent. His view also promises to deal with the Gavagai problem by appealing to descriptions which categorize the referents of terms but do not identify them (*Cf.* Devitt & Sterelny (1999, p. 80): "a name is associated, consciously or unconsciously, with a description in a grounding…the associated general categorial term does not *identify* the object"). For example, 'rabbit' refers to rabbit-hood in virtue of being associated with a description which categorizes 'rabbit' as referring to wholes, as opposed to parts.

the term's *referent*, does not apply to logical operators. This means that if one appeals only to Devitt's notion of narrow content, then one cannot use my proposal to account for how particular objects are represented. It also means that Devitt's notion of narrow content cannot replace conceptual role semantics in the account given in Chapter 3 because it cannot account for the content of logical operators. But if it weren't for the inability to account for the content of logical operators, then my account of how particular objects get represented would have been easily incorporated into Devitt's account of narrow content and this would have meant that his account, by itself, would have been a promising strategy for accounting for the content of representations of particular objects. Hence, the "almost" above.

Still, the above only shows that my account cannot be incorporated into Devitt's account of narrow content. It does not show that Devitt's account of narrow content cannot explain by itself how particular objects are represented. So can it? The answer is, "No". If one appeals only to Devitt's notion of narrow content, then one cannot appeal to *definite* descriptions to pick out the referents of representations of particular objects (because this entails an appeal to logical operators like '***E***'). If one cannot appeal to definite descriptions, then one can try appealing to singular causal strategies. However, as I argued in Chapter 1, this approach runs into the "which link" problem. One can try appealing to singular causal strategies *and* to the inferential/functional relations between a representation of a particular object ('x') and other terms. Here, the hope is that by relating 'x' with terms referring to shapes, surface reflectances, and Spelke properties (being bounded, unified, and persisting in time), one would specify a particular object. However, this also runs into problems: a) one cannot necessarily individuate a particular entity; and b) one cannot distinguish between a region and an object:

With respect to (a): since I am concerned with explaining how particular objects are represented in the visual system, I am assuming that predicates such as 'being the only even prime' are unavailable. Granted, such predicates can individuate a particular (in this case—the number 2), but they are too sophisticated to be found in the visual system. What about predicates like 'T_j': 'being at the tip of a super-determinate pyramidal configuration set on my eyes and nose right now'? Those aren't too sophisticated. Wouldn't a term concatenated with 'T_j' refer to a particular object? The answer is, "No", because it is logically possible that there may be more than one object at the same space-time location. It is here that definite descriptions would help, were they available—they could define the referent of 'x' as the *only one* that has the property represented by 'T_j'. In cases where there is more than one object at the same space-time location, the definite description would denote the complex object composed of all the entities at the location. But, granted, this may be too "sketchy" a possibility. Furthermore, what about predicates like 'R_j': 'being the reddest entity in my visual field, now'? Purportedly, no two entities can be the reddest ones in my visual field, now. If there are two equally red ones, then 'R_j' would simply not get tokened. So, wouldn't a term concatenated with 'R_j' refer to a particular object? To answer this, let us move onto (b):

With respect to (b): suppose that 'x' is concatenated with 'R_j'. Doesn't it refer to a particular object? The answer is, "No", because the entity in question can be either an object (the reddest object in my current visual field) or a region (the space-time region at which we have the property of being the reddest in my current visual field).[119] What if the predicate 'object' were present? Wouldn't 'x' refer to an object if it were also concatenated with 'object'? First of all, being concerned with the visual system, I am assuming that predicates like 'object' are not available. But even if they were, the answer is again, "No", because we still have a region available as a candidate—the space-time region at which we have the property of being the reddest object in my current visual field. Now, on my account of how particular objects are represented, I argued that it is in virtue of definite descriptions being stored and updated that 'x' represents an object, as opposed to a trope. Can't we make the same move, but without appeal to definite descriptions? Here is how it may go:

We store predicates like 'R_i'. When observing the reddest entity in our visual field through time, we have a succession of 'R_i' predicates being stored in memory. Let 'x' be concatenated with one of them: 'R_ix'. This reads as "x has the property of being the reddest object in my visual field, at t_i". Let 'x' be concatenated with the others, as well. So we have 'R_ix', 'R_jx', and 'R_kx'. In virtue of being concatenated

[119] One can object that the concatenation 'Px' means that x *has* P and that a region cannot *have* the property of being the red, for example. Only objects can have such properties. But it strikes me that the *syntax* of predicate logic does not make ontological assumptions about regions having properties. It makes assumptions only about there being particulars, properties, and properties being related to particulars. Whether the relation is one of having, being instantiated at, or whatever, is left open.

with all of them, 'x' represents an entity that has different temporal locations (at t_i, t_j, and t_k). Given that a space-time point cannot be at two space-time places at once, 'x' can only refer to an object. This would work if we are concerned only with space-time points. However, notice that points (being infinitesimal) can be problematic as the locations of properties—one wonders whether something square can be at an infinitely small point. Instead, let's treat properties as being instantiated at regions. Similarly, instead of treating tropes as properties at space-time points, we can treat them as properties at space-time *regions*. Second, note that when observing the reddest object in my visual field through time, there is a space-time region occupied by the reddest object. At this region we have instantiated the property of being the reddest object in my visual field at t_i, t_j, and t_k. That is to say that at this region we have instantiated the properties of being R_i, R_j, and R_k. So, let 'x' refer to the space-time region at which we have instantiated the property of being reddest at t_i, t_j, and t_k. We can represent the fact that at this region we have instantiated the properties above by having 'x' be concatenated with all three predicates at once: 'R_ix' and 'R_jx' and 'R_kx'. Thus, the fact that 'x' is concatenated with all three predicates does not serve to define 'x''s content as that of an object, as opposed to a region.[120]

4.B.IV. The conceptual role notion of narrow content:

[120] On my approach this problem is avoided because a definite description like $(\boldsymbol{E}x)((y)(R_jy \leftrightarrow (x = y)))$ specifies the content of 'x' as *wholly* present at a space-time region. This means that the content of 'x' at t_j and t_k must be an object since a space-time region cannot be wholly present at two different space-time locations.

The notion of narrow content that I want to defend is the conceptual role notion. Block (1986), Peacocke (1992), Rey (1997, 2007, 2008), and Botterill & Carruthers (1999) are partial to this view. The conceptual role notion of narrow content defines the narrow content of a symbol as the symbol's relations with other symbols. These relations are grounded in the casual and counterfactual relations between tokens of symbols. The counterfactual relations are specified by the syntactical rules that govern the formations and transformations of symbols. The syntactic transformations of symbols correlate with semantic relations between the symbols (e.g. implication, entailment, validity, etc.). In this sense, the *inferential* roles of symbols also constitute part of the symbols' relations. Functional relations can also be part of the symbol's relations. These include a symbol's relation with states of the system (as opposed to other symbols). Thus, on the conceptual role notion of narrow content, the narrow content is the syntactic/inferential/functional roles of a symbol.

Paradigmatic examples of how the conceptual role notion of narrow content is applied arise in explaining the meanings of logical operators. Take for example the logical operator '&': the inferential role of '&' is such that there are dispositions to infer 'a' from 'a & b', 'b' from 'a & b', and 'a & b' from 'a' and 'b'.[121] Because of its inferential role, we can say that the meaning of '&' is that of conjunction. Thus, we can speak about a content of '&' that is not broad, but narrow, and this content is its inferential role. Note that a symbol's inferential role need not involve relations to specific symbols (or types of symbols). Instead, a symbol's inferential role can be

[121] *Cf.* Peacocke (1992, p. 6).

defined through the *structure* of relations the symbol enters in. A symbol's relations with specific symbols (or types of symbols) are the non-logico-syntactic inferential roles of the symbol. In contrast, the structure of relations that a symbol enters in is the symbol's logico-syntactic inferential roles. For example, the inferential roles of '&' above are logico-syntactic in that they specify a structure of relations. The inferential roles of 'bachelor' that tie it to 'unmarried' and 'male' are non-logico-syntactic in that they tie it to specific symbols.[122]

Another example of how the conceptual role notion of narrow content is applied arises in explaining the narrow contents of indexicals. I presented the following example in the last chapter, but for purposes of illustration I present it here again: Rey (1997) defines the narrow content of 'I' (what he calls "first-person reflexive term" or "FPRT") as follows:

α is a FPRT for agent x iff:

(1-1) Whenever an input φ is received, x stores 'φα';

---

[122] Logico-syntactic and non-logico-syntactic roles are perhaps not completely independent: as Peacocke (1995, p. 231) points out, without some semantic assignment to the symbols in a syntactically defined operation, the operation cannot be said to perform a *specific* logical operation. For example, the same table of '1' and '0' symbols (imagine the truth table for '&') can be used to describe an operation as conjunction (when '1' is assigned to "true" and '0' is assigned to "false") or a disjunction (when '0' is assigned to "true" and '1' is assigned to "false"). This serves to show that conceptual role semantics should be considered as a part of a 2-factor account, and not as *the* account of content.

(1-2) Whenever x is in a mental state M, x is prepared to comp-judge a predication ‘φα’ that (ordinarily) gets released only when x is in M;

(1-3) All preference states, and all basic action descriptions in x’s decision system that lead up to action in a standard decision theoretic way are states and descriptions whose subject is α. (Rey (1997, p. 291))

On Rey’s view of FPRTs, ‘α’ means “receiver of present inputs, the instigator of outputs, and the subject of intervening mental states”. Thus, ‘α’’s narrow content is ME. Note that one need not be suspicious of the above use of intentional terms like “φ”, “M”, and “decision system”: as Rey points out, conceptual role semantics *can* appeal to intentional properties in its explanations, *as long as* it avoids those intentional properties it is trying to explain through inferential roles.

As I argued in the previous sections, Devitt’s account of narrow content and the mapping notion account cannot explain how logical operators get their content because the accounts specify a term’s narrow content *via* the term’s referent, while logical operators have no referents. Also, the descriptivist account cannot explain how indexicals get their content because the account abstracts away from a token term’s causal relations, while the latter are needed to pick out a particular individual/time/place. In contrast, as argued in the previous chapter, the conceptual role notion of narrow content *can* account for the content of logical operators, indexicals, and perhaps numbers. That is why I chose it over the others. However, there are serious problems raised for the conceptual role account:

4.C. Problems for the conceptual role notion of narrow content:

In this section I overview some of the strongest objections against conceptual role semantics and present replies to them.[123] Importantly, the replies are meant merely to show that a conceptual role semantics account *sufficient for my purposes* can be defended against the objections thrown at it. I am *not* going to defend a conceptual role semantics account as *the* account of content or defend that it applies to *all* representations. I care about conceptual role semantics to the extent that it allows me to account for logical operators, indexicals, and numbers. Also, the section is more of an overview than an exhaustive analysis of all the possible moves. It should serve merely to show that conceptual role semantics does have coherent and plausible ways to defend itself. In Section 4.C.I I present Rey's claim that the holism problem can be avoided if one appeals to asymmetric relations between symbols' relations. The circularity problem is discussed in Section 4.C.II where I argue that it can be avoided on a 2-factor strategy of content.

4.C.I. The holism problem:

The holism problem for conceptual role semantics arises when one tries to define the meaning-constitutive roles of a symbol. One can argue that a) *all* of the

[123] See for discussion Rey (1997, 2007), Fodor (2004), Carruthers (1996).

symbol's relations are meaning-constitutive; or b) only some of the symbol's relations are meaning-constitutive. The former option is the holistic option. The problem with holism arises when one tries to reconcile it with what Rey and Fodor call a "crucial condition on concepts": "their being *stable* over a wide variety of cognitive systems sharable by different people at different times" (Rey (2009, p. 2)). If holism is right, then if two people share a symbol with the same narrow content, then the symbol has identical causal and counterfactual relations in both persons' minds. However, humans do not have the same beliefs, desires, goals, etc. This means that it is practically impossible for people to share symbols with the same narrow content. This is counterintuitive since it means that if some of the content of some of our thoughts is narrow, then we cannot share those thoughts with others. Some other startling consequences about thoughts with narrow content are pointed out by Rey (1997): if beliefs have narrow content, then

> no normal person gaining new beliefs as she peers at the world around her could ever *remember anything* (after all, those new beliefs would change the contents of all old states)….One couldn't hope to find new evidence for or against a belief, since the discovery of the evidence would change its liaisons and thus its content. (Rey (1997, p. 240))

The upshot is that psychology would be unable to make generalizations over people's mental states with narrow content—surely not a positive result.

But perhaps, as Block (1986) suggests, the holism problem can be avoided by assuming that the thoughts with narrow content are sufficiently similar in narrow content. However, he admits that "there are no specific suggestions as to what the dimensions of similarity of meaning are or how they relate to one another" (forthcoming). Until such are provided the holism problem remains a threat—it is what has chased Fodor (1995) away from conceptual role semantics.

We have seen that taking all of the symbol's roles to be meaning-constitutive runs into serious problems. So why isn't the obvious alternative to take only some of the symbol's roles as meaning-constitutive? The reason is that doing so is also deeply problematic: as Fodor (1990, 1995, 2004) has pointed out, there doesn't seem to be a principled manner of individuating a subset of roles as a symbol's meaning-constitutive roles in the mind. A symbol's role is defined in terms of the symbol's causal and counterfactual relations with other symbols. However, there is a multitude of such relations. Which one are the meaning-constitutive ones? For example: '&' might have an inferential role specified as above, but it might, on top of that, have a role such that, whenever it is tokened, it causes 'conjugal' to be tokened, as well (because 'conjunction' and 'conjugal' sound similar, for example). Which of the two roles is meaning-constitutive, then?

To define in a principled manner some of a symbol's relations as *the* meaning-constitutive ones is in essence to provide an analytic/synthetic distinction with regards to a symbol's roles. Quine (1956) argued that there is no principled manner of distinguishing analytic from synthetic truths[124]. Fodor, convinced by Quine, concludes that there is no principled manner of distinguishing the meaning-constitutive relations of a symbol from those that are not such. Since taking *all* of the symbol's relations as meaning-constitutive is also implausible, and since these two options are the only ones for conceptual role semantics, Fodor concludes that conceptual role semantics cannot be the right account for narrow content.

But not all hope is lost. Horwich (1992, 1998), contra Quine, argues that we can have an analytic/synthetic distinction. According to Horwich, the language faculty contains "meaning postulates" that define the meanings of their constituent linguistic concepts.[125] A sentence gets to count as a "meaning postulate" of a concept if accepting this sentence explains why other sentences containing the concept are also accepted. In other words, all other meanings of the concept are explained through the meaning assigned to it in the meaning postulate. An example of a meaning postulate for the concept 'bachelor' is 'a bachelor is an unmarried male'. All other meanings of 'bachelor' are defined through the meaning assigned to 'bachelor' in the meaning postulate. Devitt (2002), while disagreeing with Horwich's proposal, points out that Horwich's idea can be extended to apply to non-linguistic concepts. Thus, a sentence gets to count as a meaning postulate of a non-linguistic concept if the

[124] But see Rey (2007, 2008) for a plausible objection to Quine.

[125] Another option is offered by Peacocke (1992): the meaning-constitutive relations are "primitively compelling". The problem then becomes to give sufficient conditions for an inference being "primitively compelling" in a way that would not presuppose that certain roles are meaning-constitutive. I discuss this option in the next section.

presence of such a sentence in the belief box explains why other sentences containing the concept are also accepted. So, the suggestion is that Horwich's proposal could provide the principled manner of defining the meaning-constitutive roles of a symbol that is needed to solve the holism problem.

Fodor (1998), however, argues against meaning postulates because he does not see what work they do:

> Imagine two minds that differ in that 'whale → mammal' is a meaning postulate for one but it is 'general knowledge' for the other. Are any further differences between these minds entailed? If so, which ones? Is this wheel attached to anything at all? (Fodor (1998, p. 111-112))

But Horwich's proposal does predict *counterfactual* differences between two such minds. In a mind where 'bachelor → unmarried male' is a meaning postulate, all the roles of 'bachelor' will be explained by the meaning postulate, but not *vice versa*. In a mind where 'bachelor → unmarried male' is not a meaning postulate, then this dependency relation will not hold. In other words, whether a sentence is a meaning postulate or not determines the counterfactual relations involving a symbol whose meaning is defined through the postulate. This leads to a common objection that Horwich, Rey, and Carruthers have all made against Fodor in different ways: that appealing to counterfactual relations allows the charge of holism to lose some of its menace. The idea is that the beliefs of an agent can change, without this changing a belief's dispositions to enter in relations with other beliefs—a belief's counterfactual relations with other beliefs:

> [D]ifferent thinkers, despite their differences of belief, and despite the differences in the epistemic liaisons of any given belief, may nevertheless entertain many of the same beliefs. For the same *conditionals* can be true of them. For example, it can be true of me *both* before *and* after my formative experience with the reporter from the *Guardian,* that I believe Mrs Thatcher to be the longest serving UK prime minister of the twentieth century. For although the epistemic liaisons of that belief have changed, it will be true of me on both occasions that *if* I believe the *Guardian* to be reliable, *then* I shall regard, 'the *Guardian* says so', as a reason for holding my belief about Mrs Thatcher. (Carruthers (1996, p. 111))

As Rey has pointed out (p. c.), it is striking that Fodor has not seen the potential of counterfactuals to deal with the holism problem, given that he uses them to define the broad content of symbols in his ADA. According to Fodor, the way to distinguish between the meaning constitutive nomic relations involving 'x' and the other ones is to look at the counterfactuals: the meaning constitutive nomic relations involving 'x' are the ones, but for which the others wouldn't exist. Rey (2007, 2008) noticed the common appeal to counterfactuals to define a symbol's content in Fodor and Horwich and pointed out that Horwich's idea is "simply a redeployment of the asymmetric dependency proposal, only this time *inside* the agent" (Rey (2009, p. 6)). Rey goes on to suggest an account of which both Fodor's and Horwich's proposal can be special cases:

(BAS-LOT) The content of an internal symbol in an agent's LOT is determined by the property of a meaningful tokening of a term that is explanatory basic: the one on which all other tokens with that meaning asymmetrically/explanatorily depend by virtue of that property. (Rey (2009, p. 11))

Rey acknowledges that this account is not a reduction of meaning because it appeals to "meaningful tokenings"[126]. Nevertheless, assuming one can individuate which are a symbol's meaningful tokenings, one can then distinguish between the symbol's meaning-constitutive and non-meaning-constitutive inferential roles. In this way one can define the meaning of a concept that does not enter into the language faculty or the meaning of a concept in creatures that do not posses language faculties. This makes Rey's proposal a more general solution to the holism problem than Horwich's.[127]

To sum up: the way to distinguish between the meaning constitutive inferential relations involving 'x' and the other ones is to look at the counterfactuals:

[126] It is important to distinguish "meaningful" from "meaning-constitutive". If a tokening of a symbol is the latter, then it is the former, but not *vice versa*.

[127] Rey (p. c.) points out that there is the following problem with BAS-LOT: suppose that I name my cat 'Beauty' and after her death, I use 'Beauty' to refer to one of her kittens that has taken her place in my home. According to BAS-LOT, since the reference of 'Beauty' to the kitten asymmetrically depends on the reference of 'Beauty' to my former cat, then 'Beauty' refers to my dead cat. But this seems wrong. One way this worry can be avoided is to use, á la Fodor, a condition that stipulates that the asymmetric dependency is synchronic.

the meaning-constitutive inferential relations involving ‘x’ are the ones, but for which the others wouldn’t exist. In other words, one can appeal to asymmetric dependencies between inferential relations to specify which of the symbol’s inferential relations are meaning-constitutive. What follows is an example. Just like ADA, it is a special case of BAS-LOT. Unlike ADA, however, it is about a symbol’s inferential roles. It is modeled on ADA and for ease of reference I will call it AIR (for Asymmetric Inferential Roles):

The inferential role of a symbol ‘x’ is defined by its relation with representations ‘ψ’$_1$, ‘ψ’$_2$, … , ‘ψ’$_n$ if:

i) There is an inferential[128] relation between ‘x’ and ‘ψ’$_1$, ‘ψ’$_2$, … , ‘ψ’$_n$.

ii) Any inferential relation between ‘x’ and symbols other than ‘ψ’$_1$, ‘ψ’$_2$, … , ‘ψ’$_n$ depends on the one between ‘x’ and ‘ψ’$_1$, ‘ψ’$_2$, … , ‘ψ’$_n$, but not *vice versa*.

iii) There are inferential relations between ‘x’ and representations other than ‘ψ’$_1$, ‘ψ’$_2$, … , ‘ψ’$_n$.

iv) The dependence in (ii) is synchronic.

[128] Instead of “inferential”, I could have used “nomic”, just as in ADA. The reason is that inferential relations are a type of nomic relations between symbols.

v) The asymmetric relation in (ii) is one on which other asymmetric relations involving a law between 'x' and 'ψ'$_1$, 'ψ'$_2$, … , 'ψ'$_n$ depend, but not *vice versa*.[129]

This proposed solution could solve the holism problem because it gives us a principled way of specifying which of the symbol's inferential relations are meaning-constitutive. One doesn't need to define the inferential role of a symbol by specifying *all* of the roles the symbol has—just the ones that are meaning-constitutive.

4.C.II. The circularity problem:

Fodor (2004) argues that conceptual role semantics is unable to account even for the meaning of logical operators without running into circularity. If he is correct, then appealing to conceptual role semantics is a broken tool and cannot be of use in defining a conceptual role notion of narrow content.

Fodor argues that the possession conditions of a concept cannot be deconstructed into inferential roles. The reason Fodor gives is that "*understanding* S is prior in the order of analysis to *grasping its role in inference*" (Fodor (2004, p. 43)). In other words, only after one possesses and grasps a concept can one grasp the concept's inferential relations with other concepts. Prior to grasping 'bachelor', I cannot grasp 'bachelors are unmarried men'. 'Bachelors are unmarried men' can fix the meaning of 'bachelor' only if 'bachelor' is already possessed and its meaning

[129] This is Rey's addendum condition. See Chapter 2, Section 2.C.III.

provided. Any attempt to fix the meaning of a concept through inferential roles ends up presuming that the concept is already possessed and such an attempt is therefore circular:

Fodor illustrates this point with Peacocke's account of conjunction:

> Conjunction is that concept C to possess which a thinker must find transitions of the … forms [F] primitively compelling, and must do so because they are of these forms. (Peacocke (1992, p. 6))

Here "[F]" stands for two elimination forms (pCq → p; pCq → q) and one introduction form (p, q → pCq). *Prima facie*, Peacocke's account is not circular—it doesn't appeal at all to conjunction to specify the inferential roles that supposedly define the possession conditions of '&'. Peacocke argues that the inferential transitions of the forms are primitively compelling *because of their form*. This is in order to avoid cases where the transitions are found compelling because of accidental or pathologic causes. It is in analyzing what this form is that Fodor finds the hidden circularity in Peacocke's account of conjunction.

Fodor points out that the reason for accepting these inferences cannot be because they have a certain *logical* form, for fear of circularity. To see the circularity, first note that the logical form of an inference is defined in terms of logical constants. So, if one tries to explain the meaning of logical constants *via* primitively compelling transitions, proceeds to account for the latter *via* logical form, and finally accounts for logical form *via* appeal to logical constants, then one would have tried to account for logical constants *via* logical constants.

The reason for accepting these inferences cannot be *syntactic* form, either. Syntactic form is distinct from validity/acceptability because the latter are semantic notions. As a result, syntactic form cannot explain why a certain inference is found valid or acceptable:

> ‘p and q → p’ is valid (not in virtue of its syntax, but) in virtue of the meaning of ‘and’. It is, to repeat, (only because you accept the argument (because you know the meaning of ‘and’)) that your accepting it manifests your grasp of CONJUNCTION. (Fodor (2004, p. 45))

Fodor concludes that Peacocke cannot give an account of ‘&’. Since Peacocke’s story of ‘&’ is the paradigmatic account of conjunction in conceptual role semantics, then supposedly conceptual role semantics cannot account for the possession conditions of ‘&’.

A related problem is raised by Peacocke[130]. Consider the inferences appealed to in defining ‘&’: the inferential/functional role of ‘&’ is such that there are

[130] *Cf.* Peacocke (1995, p. 231): “An “and”-gate may be characterized simply relative to certain assignments of 0 and 1 to the nodes to which it is connected. That is not yet a semantical characterization—indeed, if falsity and not truth were the semantic significance of the assignment of 1 to a node, such an “and”-gate would function semantically like alternation (the output node indicating

dispositions to infer/token 'a' from 'a & b', 'b' from 'a & b', and 'a & b' from 'a' and 'b'. If we take the inferring/tokening of 'a' and 'b' to imply being *true*, then we can have the following *truth* table:

---

falsity iff both the input nodes indicate falsity)." Peacocke makes this observation to show that if we *completely* abstract away semantic characterizations, then we are not going to be able to treat a system as performing any computations—even ones that are as simple as addition. For Peacocke this means that there is "no computation without representation" (Fodor's words (*cf.* Fodor (1981) and Pylyshyn (1984))): every computational system operates on symbols with semantic properties and preserves the semantic properties of strings of symbols. For Peacocke the upshot is the one that Rey (2004) (see below) makes—conceptual role semantics need not be averse to appealing to a symbol's intentional properties in its explanations, as long as these are not the very properties it is trying to explain.

In arguing against "no computation without representation" Piccinini (2006) argues that if we take into account inputs and outputs, then a system's function can be defined *non-semantically* in a mechanistic fashion. His point is that, contrary to Peacocke and Fodor, a computation need not be defined over symbols: for instance, instead of the computation involving *symbol* α, we can speak of the computation involving inscriptions on a tape of type "α". While he is right that we can define a system's function in such a mechanistic fashion, such an account seems problematic because a) it would miss generalizations among differently constituted mechanisms where the inscriptions are different (instead of "α" we have "a"); and b) the mechanistically specified computation that does not appeal to symbols would be of no use to conceptual role semanticists who are concerned with relations between symbols that can explain difference in content (if such a deflationary notion is used by conceptual role semanticists, then, as Fodor (2004a, p. 101) points out, the notion of having a concept would be equally deflationary, which would jeopardize the project of conceptual role semanticists to explain difference in content through difference in computational roles).

a	b	a * b
1	1	1
1	0	0
0	1	0
0	0	0

This, then, is the truth table of '&'. But if we take the tokening of 'a' and 'b' to imply being *false*, then the truth table above is the truth table of 'v' ('or'). This suggests that we cannot define the meaning of '*' through appeal to inferential/functional relations alone. Appeal to independently assigned semantic meaning seems necessary.

Essentially Fodor's (and Peacocke's) point is that what is lacking in conceptual role semantics is an account of why a symbol's inferential roles should be identified with its meaning. With respect to '&', the point is that even if '&''s syntactical role mimics its semantic one, this does not explain why the syntactical role should account for the meaning of '&'. Pure conceptual role accounts that attempt to explain that end up being circular. However, as Rey (2004) points out, Fodor makes a mistake in assuming that conceptual role semantics cannot appeal to *any* intentional properties in specifying the content of '&'. According to Fodor conceptual role semantics "wants an account of content that doesn't presuppose semantic or intentional properties (Fodor (2004a, p. 110))". But even if some conceptual role accounts strive to be pure, the strategy that I have adopted in this dissertation is of mixed blood: my appeal to BAS-LOT entails that the content of a symbol is determined by the property of a meaningful tokening of a term that is explanatory basic—this property can be one that specifies the symbol's nomic relations with external properties (á la ADA), the symbol's relations with other symbols (á la AIR), or both at the same time. This means that we have the following way of addressing the circularity objection: one can appeal to nomic relations with external properties for some symbols and then use those to define á la conceptual role semantics the content of *other* symbols. As long as one avoids presupposing the symbol's intentional properties that one is trying to explain, one can appeal to *other* symbols' intentional properties while free of the charge of circularity.

Rey (2004) gives an example of such an approach in accounting how the (narrow) content of '&' is specified through conceptual role semantics: Fodor claimed that syntax is of no help in defining what counts as valid or acceptable. But its help is not needed to define what counts as acceptable. Rey (2004, p. 76) argues that one way of defining which inferences count as acceptable is to have a "Belief-box"—if one has sentence 'p' in one's 'Belief-box' then one believes/accepts that "p".[131] Slightly simplifying Rey's story, we can say that what accounts for why 'p' means p is ADA. What accounts for 'Perceptual module' being a perceptual module is that it tokens symbols (that have content á la ADA about certain properties) when

[131] Fodor (2004a) objects to Rey's proposal by arguing that a conceptual role semantics cannot appeal to Rey's solution, since Rey appeals to beliefs, and conceptual role semantics cannot presuppose any intentional properties in its account of '&'. However, as is pointed out above, conceptual role semantics doesn't have to avoid any intentional properties. It just has to avoid those intentional properties it is trying to explain through inferential roles.

there are objects having the specified properties causing the token symbols. What accounts for a certain sub-system being a 'Decision system' is that its outputs cause the machine to move in various ways. What accounts for a certain register being a 'Belief-box' is that it takes as input outputs from 'Perceptual module' and stores them and also outputs to 'Decision system'. What accounts for 'p' being a belief is that it is stored in 'Belief-box'. Then, if a sentence is granted entry in 'Belief-box', it is accepted and believed. If not, then it isn't. One can then define '&' as follows: '&' means "and" iff ['p & q' would be granted entry to the Belief-box, only if 'p' and 'q' would be given entry to the Belief-box]. Then, one can think "p and q" iff ['p & q' would be granted entry in the Belief-box and 'p' would mean "p", 'q' would mean "q", and '&' would mean "and"]. Note that in specifying the content of '&' an appeal was made to *other* symbols with intentional properties (*via* the appeal to 'Belief-box', which was specified through 'Perceptual module', which in turn was specified through appeal to symbols that get their content á la ADA). Since the content of '&' was not accounted for *via* an appeal to the intentional properties of '&', circularity has been avoided.

What this means is that, contrary to Fodor, one can fix the narrow content of a concept C through inferential roles *without* presuming that the concept is already possessed and grasped. That some intentional properties (e.g. broad content) of some symbols need to be assumed is granted, but these symbols would not be what conceptual role semantics gives an account of in the case of the concept C. This avoids the problem of circularity that Fodor raises. While Fodor is completely right that conceptual role semantics cannot account by itself for *all* the intentional properties of the mind, it need not do so. It can avail itself of the fruits of externalist accounts as long as it doesn't presuppose the intentional properties it is trying to explain.[132]

4.D. Conclusion:

If what I have said above is on the right track, then conceptual role semantics has available and plausible moves to make in defending itself against the holism and the circularity problems. These two problems that are of prime concern to me, because I appeal to conceptual role semantics to account for how indexicals, logical operators, and numbers get their content. If the holism objection runs, then I would not be able to provide the meaning-constitutive roles of any of the above. If the circularity objection runs, then I would not be able to provide the content of logical operators in a non-circular manner. Granted, there are other objections against conceptual role semantics (e.g. the compositionality of thought problem (*cf.* Fodor (1998, 2004, 2004a)). However, such problems do not concern me, since they either

[132] Granted, there are other objections against conceptual role semantics (e.g. the compositionality of thought problem (*cf.* Fodor (1998, 2004, 2004a)). However, such problems do not concern me, since they either attack the capacity of conceptual role semantics to be *the* account of content—something that I am happy to deny—or they attack its capacity to give an account of the content of concepts for whose content individuation conditions I appeal to ADA (e.g. recognitional concepts).

attack the capacity of conceptual role semantics to be *the* account of content—something that I am happy to deny—or they attack its capacity to give an account of the content of concepts for whose content individuation conditions I appeal to ADA (e.g. recognitional concepts).

Chapter 5: Spatial Representations and Narrow Content for FINSTs

5.A. Introduction:

In the previous chapters, I have argued that definite descriptions are necessary for the ability to represent particular objects. If the ability to form definite descriptions presupposes concepts, then a) FINSTs refer through conceptual spatio-temporal representations; and b) FINSTs are conceptual representations themselves. In Section 5.B I will present Pylyshyn's position (Section 5.B.I), expose the problems with his account that arise if it is interpreted to depend on a singular causal strategy of content (Section 5.B.II), and expose the problems with his account that arise if it is interpreted to depend on a nomic strategy of content (Section 5.B.III). In Section 5.C I will present problems with Pylyshyn's account of how FINSTs track (Section 5.C.I) and then I will criticize Pylyshyn's empirical argument that tracking of and reference to visual objects cannot be done through representations of location (Section 5.C.II). Finally, in Section 5.D, I present Clark's view of sentience and point out how my view fits Clark's.

5.B. Pylyshyn's account and representations of particulars:

In this section I present Pylyshyn's account (Section 5.B.I) and then examine his claim that FINSTs refer through an information link with their referents (Section 5.B.II). I point out that *if* this information link is read as a singular causal one, then, for reasons given in Chapter 1, Pylyshyn's account runs into the "which link" problem, which makes it questionable. In Section 5.B.III I argue that *if* the information link between FINSTs and their referents is read as a nomic one, then, then only one of the following options is the case: a) it is implausible to think that FINSTs refer to particular objects; or b) FINSTs refer to particular objects *via* definite descriptions. If the latter option is the case, and if definite descriptions presuppose concepts, then FINSTs are conceptual representations referring *via* conceptual representations of spatio-temporal relations. I finish Section 5.B.III by criticizing Pylyshyn's theoretical arguments for FINSTs being unmediated by concepts.

5.B.I. Pylyshyn's account:

Pylyshyn postulates his visual indexes (FINSTs) to account for his Multiple Object Tracking (MOT) experiments[133]. In these experiments the test subjects are presented with some number (usually more than five) of what looks like objects on a

[133] *Cf.* Pylyshyn (2001b, 2007).

computer screen. Following a cue that draws the subjects' attention to four or five of them, the test subjects are tasked to track them for a short period of time. The objects move randomly but, nevertheless, people are extremely good in tracking four to five objects. Here is an illustration of what the MOT experiment looks like[134]:

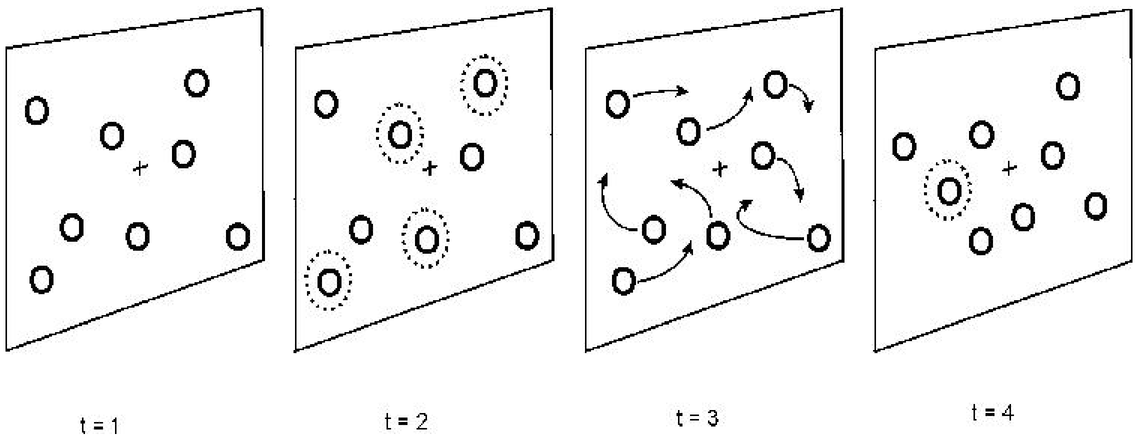

Figure 5.1.: Eight of what looks like identical objects are shown (at t = 1). A proper subset is flashed (at t=2), after which the objects randomly move. Then the observer has to pick out the objects initially flashed at t=4.

According to Pylyshyn's FINST model[135] the data can be explained by supposing that about 4-5 objects are preattentively indexed in parallel by FINSTs. The primary function of FINSTs is "to individuate a small number of objects so that they may be directly accessed and subjected to focused attentional processing." A FINSTed object need not be found by attentional scanning. According to the FINST

[134] *Cf.* Pylyshyn (2001b, p. 142).

[135] *Cf.* Pylyshyn (2000, p. 2).

model, FINSTs are tokened in a bottom-up manner. FINSTs are "object-centered" and refer to objects despite changes in the location of the latter: in other words, FINSTs not only refer, but track objects.

Pylyshyn and Storm (1988) did an experiment where the participants had to track a flashed proper subset of identical, randomly moving objects. Since the objects had identical properties (apart from location) the only way the participants could track them was *via* their historical continuity. Pylyshyn and Storm (1988) found when up to 5 objects were tracked the success rate was about 90%. Using a computer simulation of tracking, where focal attention moved serially among the tracked objects and updated a record of their locations, a success rate of only 50% was produced. The speed of the focal attention was as high as it has been recorded in the psychological literature and it was even assumed that participants stored "the predicted locations based on the direction and speed of the targets' motion, and that they used a guessing strategy when they were uncertain"[136]. Pylyshyn concludes on the basis of these results that tracking is done in parallel without the use of focal attention and without encoding the locations of the tracked objects.

In further experiments Pylyshyn has tested whether the subjects are tracking the objects by using representations of their properties[137]. Pylyshyn takes the results

---

[136] *Cf.* Pylyshyn (2000, p. 3).

[137] *Cf.* Pylyshyn and Storm (1988), Scholl, Pylyshyn, and Franconeri (1999), and Dennis and Pylyshyn (2002).

to suggest that our visual system employs four to five representations that refer to and track, not properties, but what typically turn out to be particular objects, and do so without using concepts to encode the object's properties:

> [FINSTs] provide a reference to some sensory individual…without thereby encoding any property of the individual that is indexed. (Pylyshyn (2007, p. 67))

These supposed representations of objects Pylyshyn calls "FINSTs" and he takes them to explain how the visual system binds properties together, tracks objects, and allocates attention.

What FINSTs refer to, Pylyshyn calls "FINGs". FINGs are *not* to be identified as objects in the ontological sense. Instead, they are *intentional* objects (what Pylyshyn calls "sensory individuals" and "visual objects")—things we perceive as objects.[138] Another way to think of FINGs is as *instantiations* of a complex property that nomically co-varies with FINSTs. The complex property could be something like the property of being bounded, unified, and persisting through time (the property of being a Spelke object). A *token* FINSTs would refer, according to Pylyshyn, to a particular instantiation of that property—a FING.

---

[138] *Cf*. Pylyshyn (2007, p. v): FINTSs "attach to what in our sort of world typically turn out to be individual visible physical objects."

While Pylyshyn does not believe that any concepts play a role in explaining how FINSTs refer, he believes that the FING's properties (e.g. having a location, being bounded, being unified, etc.) are causally responsible for the tokening of FINSTs—it is just that they need not be encoded as concepts[139] for the FINST to refer:

[139] A brief clarification: Pylyshyn says below that the properties are "not encoded", not that they are "not encoded *as concepts*". I take him to mean the latter, though, given that prior to the tokening of FINSTs there are states that carry information about the FING's properties. On some accounts of reference this is sufficient to treat those states as representations. Pylyshyn can grant that, but deny that they are conceptual ones. *Cf*. Pylyshyn (2007, p. 73): "My concern at this point is primarily with the question of whether certain spatial properties such as location and distance are encoded (represented) early in vision and whether they can serve as the basis for primitive selection. For this purpose we should recall that there is a substantial difference between the claim that certain properties, such as the locations and distances between objects in the world, play a causal role in a visual process, and the claim that these properties play a role by virtue of being represented. Objects are always at some location or other, and the effect they have on a perceiver may depend on where they are, but the locations they are at (even their locations relative to one another) may or may not be represented. The same is true of the location of objects on the proximal stimulus (e.g., on the retina) or further up in the nervous system, such as patterns of activity on the retinotopically-organized fibers leading from the eye, or in primary visual cortex, which is largely retinotopically mapped. Since these locations are past the sensors, are they necessarily representations? If so what is the essential difference between the way that distance in the world affects perception and the way that the corresponding distance on a neighborhood-preserving (i.e., homeomorphic) anatomical mapping affects perception (for ease of reference I will refer to the result of such mappings as "neural layouts" or NLs)? We can say that such neural layouts register (rather than represent) spatial properties. Such neural layouts help to illustrate the general theme that there are many types of representations, ranging from conceptual, through

> [there] are specific properties that cause a FINST index to be assigned and that enable it to keep track of the indexed individuals – but those properties *just are not encoded and a representation of those properties is not used in carrying out those functions*. (*Cf.* Pylyshyn (2007, p. 139, *emphasis* mine))

According to Pylyshyn, FINSTs represent FINGs the way demonstratives do—simply as "that":

> what is selected is merely indexed or demonstrated and not conceptualized in any way. (Pylyshyn (2007, p. 62))

Note that FINSTs being demonstratives does not necessitate that conceptual representations do not play a role in how FINSTs refer. If one treats FINSTs as a Kaplanian 'dthat', then one can treat FINSTs as demonstratives that refer *via* conceptual representations in the definite description that defines the content of a FINST—one will only deny that the definite description (and the concepts within) is constitutive of the content of a FINST. However, as evidenced by the previous quote, Pylyshyn holds that conceptual representations of a FING's properties are not used in the case of referring and tracking. In what follows I will attack this position. If Pylyshyn adopts a Kaplanian position, however, then there is no disagreement between us.

---

subpersonal to informational states that are better referred to as registrations of spatial information rather than representations."

Since Pylyshyn's experiments purportedly show that conceptual representations of properties play no role in MOT, Pylyshyn concludes that FINGs

> are not represented [by FINSTs] *as* objects or *as Xs* for *any* possible category *X*. They are just picked out transparently by a causal or informational process without being conceptualized as something or other. (Pylyshyn (2007, p. 56))

The notion of information appealed to here is a Dretskian one:

> The minimum function needed for an object to have the right kind of causal or informational link with a [FINST] token is that there be some causal or nomologically-supported dependency between the object and its associated symbol token (this is similar to the informational view of reference, as developed by Dretske, 1981). What kind of dependency? Any simple causal connection will do. (Pylyshyn (2007, p. 56-57))

A symbol is taken to carry Dretskian information about another if there is a *nomic* relation between the two. But, given the discussion in Chapter 1, I will explore both options with respect to FINSTs—that they get their reference in virtue of singular causal relations to FINGs (Section 5.B.II) and that they get their reference in virtue of nomic relations to FINGs (Section 5.B.III):

5.B.II. FINSTs referring through singular causal relations:

Given the discussion in Chapter 1, we have reasons to be dubious that FINSTs get their reference in virtue of singular causal relations with FINGs: we have the "which link" problem—the problem of distinguishing which of the links in the causal chain leading to a FINST is its referent. I argued that Fodor's attempt at a solution—his triangulation account does not work. Pylyshyn in fact is sympathetic to this approach:

> Another possibility [for solving the "which link" problem (recently suggested by Fodor [2008, Chapter 9] is that counterfactuals may rule out all but the correct link in the causal chain. This proposal works because the reference for which we are trying to give an account is a visual reference, so only currently visible things are relevant and only links in a causal chain to the FINST from some initial but currently visible cause have to be considered (which excludes the big bang, and the switching on of a light earlier, among other things; but it allows the light source if it is visible). Such a chain must pass through some property of the referent. Which property? The answer cannot be determined solely from that one chain—it needs another parameter. According to Fodor's proposal (which he calls a triangulation), if we consider counterfactual causal chains that end with the same FINST but have a slightly different perspective (a slightly different viewer location) then if the chains intersect they will intersect at the link that is the referent of that FINST. (Pylyshyn (2007, p. 97, ff. 14))
>
> .

However, as I argued in Chapter 1, there are problems with Fodor's triangulation approach. I briefly recapitulate them here:

Interference effects:

The first problem for Fodor’s triangulation account is that the first intersection need not be the one we would intuitively label as the “right” one. We get such cases when the causal chains are instantiated by light beams. The reason we get them is because of a *fundamental* property of light—it interferes with itself. The upshot, as I explain below, is that the first intersection of the actual and the counterfactual causal chains is the link where the interference happens and that is not the intuitively “right” referent. In more detail:

Suppose that a laser is directed at a transducer and that all the photons from the laser end up at the transducer. Suppose that there is no other light present and finally suppose that the laser is instantiating a causal chain that *via* the transducer leads up to a token of a symbol (‘x’). A problem emerges once we consider lasers and interference effects:

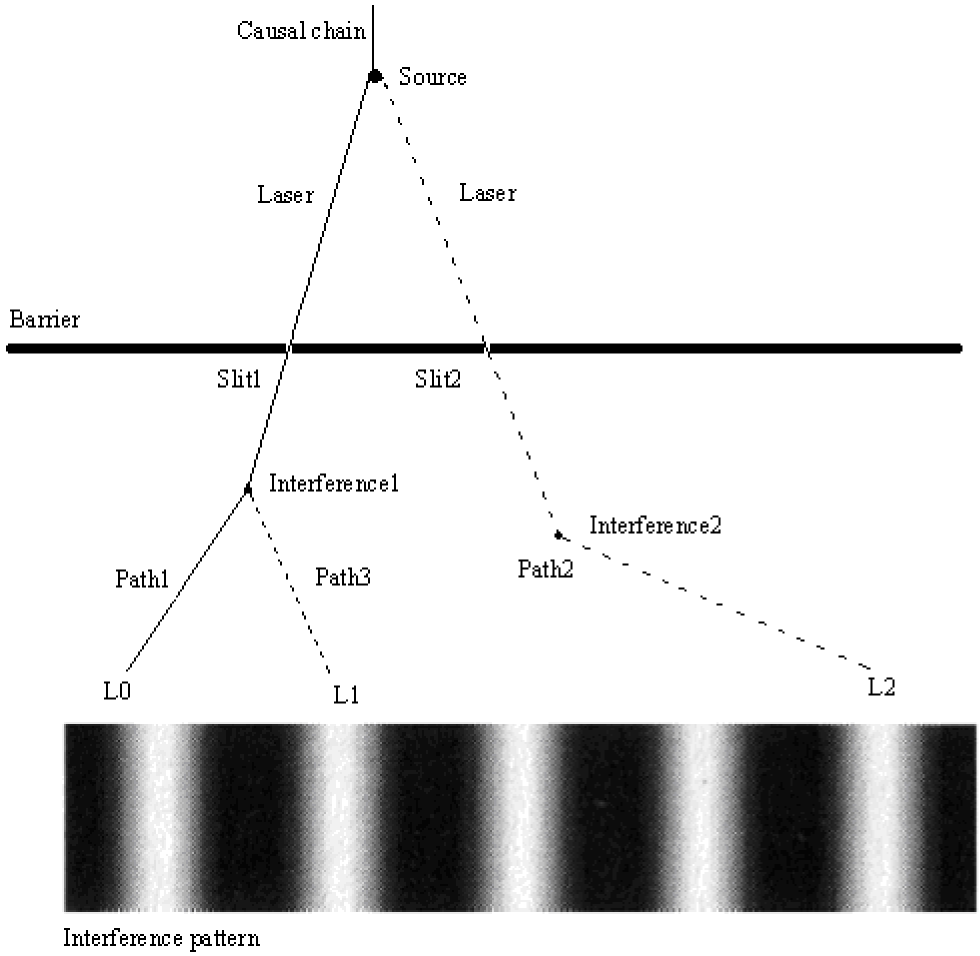

Figure 5.2.: Schemata of a source emitting a laser, photons of which pass (one per day) through one of the slits in a barrier. Since a second slit is open, there are interferences, which explain the interference pattern. The striped lines denote the counterfactual paths of photons. L0 denotes the actual location of a transducer, while L1 and L2 denote two counterfactual locations of the transducer.

Suppose that a transducer is placed at L0—in the path of a photon that has been interfered with. Let's apply the triangulation account here: if the transducer were to be placed at L2 and if a photon were to travel down that path, then we would have an actual and a counterfactual causal chains (they would be bent a little, but that matters not) that intersect at Source, which is the answer that we intuitively want. However, suppose the transducer were to be moved in location L1. Suppose that in this counterfactual scenario the photon that moved along Path1 now moves along Path3. Then, the two causal chains would intersect at the point of interference (Interference1) and so Interference1 would be the referent of ‘x’. First of all this seems wrong. Furthermore, another problem is that depending on where we choose to place the transducer, we get two different intersections of the causal and the counterfactual chains: Source and Interference1. This means that the triangulation account does not determine a unique referent of ‘x’.

Single-slit diffraction:

A second problem with the triangulation account can be seen once we consider single-slit diffraction. When light passes through a single slit it can be observed to spread out (or “fray”):

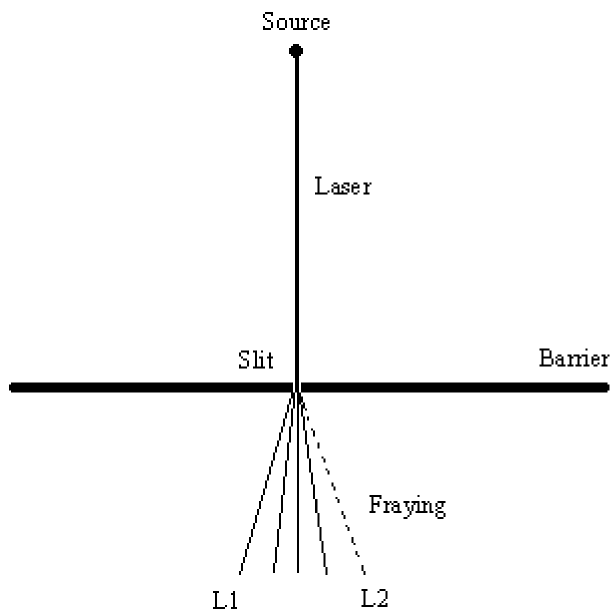

Figure 5.3.: Schemata of a source emitting a laser, photons of which pass through a slit in a barrier. The striped line denotes a counterfactual path of a photon. L1 is the location of the actual observer. L2 is the counterfactual location of the observer.

Now, consider an observer at L1 and let's apply the triangulation account: were she to move, say to L2, then she would be in the path of a fraying light beam that intersects the actual light beam at the location of the slit. The actual and the counterfactual light beams intersect and converge at Slit. By the triangulation account then, it is the photons at the location of Slit which are the referent of a token symbol 'x'. However, as Cole points out[140], "we can refer to things further back on the causal

[140] *Cf.* Cole (2009, p. 442). In his thought experiment he uses sound waves passing through a key hole, but the point is the same: "[W]e can refer to things further back on the causal chain than any perceptual

chain than any perceptual intersection". That is, if I am standing at L1 and observing the light beam, I am tokening 'x' and using it to refer not to the location of Slit, but to a point behind it. Thus, the triangulation account gives the wrong result.

Objects, tropes, and states of affairs:

The triangulation account is an account of how particular objects get represented. It takes for granted that only objects enter in singular causal relations. As a result, when it individuates a singular causal link, the link can only be an object. Now, objects enter in singular causal relations *in virtue* of the properties they have. That is the same as saying that nomic relations between properties "cover" the singular causal relations that objects enter in. Yet another way to put the point is that properties play a causal role in singular causal relations between objects. So when we are individuating *via* triangulation *a* causal link in a singular causal chain, why are we specifying only an object and not also the properties that the object has? One answer is that we are individuating a particular *at* a space-time location. Objects can have space-time locations. Properties don't. But now, it seems that a lot hinges on what metaphysical view we are committed to:

Suppose that one denies that properties are universals and appeals instead to tropes—super-determinate properties instantiated at space-time locations. On this

---

intersection. Suppose e.g. you and I listen to a conversation through a keyhole. We can be thinking about the speakers, even though our perceptual causal chains intersect at the keyhole."

view, instead of a property being causally relevant to a singular causal relation that an object enters in, we have a trope, instantiated at the space-time location of the object, being causally relevant. Triangulation, then, triangulates two particular entities at the same space-time location—an object and a trope at the space-time location of the object. Both play a causal role in the causal chain leading to a token symbol ‘x’. So, which one is the referent of ‘x’? One can argue that since both are playing a causal role in the singular causal chain, both, together, should be the referent of ‘x’. This is treating them as an event or, alternatively, a state of affairs (a particular object instantiating a property at a space-time location). But if one goes this way, then one has given an account of how states of affairs (or events) are represented, not how objects are. Note that the two are different, since objects, unlike states of affairs (and events), can move through space-time, while the latter can’t. This means that *if* one makes the metaphysical moves above, then triangulation ceases to be the account of representation of particular objects that we are looking for. Granted, the metaphysical moves may be deemed “sketchy” due to the appeal to tropes. But such moves *are* defended in the literature by nominalists and, in any case, the point is merely that if the triangulation account is an account of how objects are represented, then it needs to be elaborated to prevent “sketchy” moves like the ones above being made.

Beyond pointing to Fodor’s triangulation account, Pylyshyn does not address the issues above. With respect to the “which link” problem he says that

> [i]t is one of the “big questions” about how reference is naturalized and is beyond the scope of this monograph. (Pylyshyn (2007, p. 97))

But addressing the "big questions" is essential if Pylyshyn's claims—that FINSTs represent without appeal to concepts—are to be held as plausible. If one tries to account for the content of FINSTs through singular causal relations, then one runs into the "which link" problem. But if one abandons singular causal strategies for nomic ones, due to the difficulties with the "which link" problem, then, as I argued in Chapter 3, there are good reasons to believe that representations of particulars get their content *via* concepts, contrary to Pylyshyn:

5.B.III. FINSTs referring through nomic relations:

Suppose that the relation between a FINST and a FING is a nomic one. As I discussed in Chapters 2 and 3, nomic relations, by themselves, cannot explain how representations of particulars get their content: Adams & Aizawa (1997, p. 274-5) point that ADA cannot account for how representations of particulars get their content. ADA is based on nomic relations and Fodor himself assumes that there are no nomic relations between particular objects. Nomic relations are assumed to be between properties. This leaves ADA as incapable of addressing the question of how representations of particular objects get their content. One way to help ADA here is to appeal to properties such as the property of [being an instance of Aristotle] and to nomic relations between the property of being an instance of Aristotle and the property of being a token of 'Aristotle'.[141] But, as Adams (2003, p. 156) objects, this

[141] *Cf*. Fodor (1995, Index B, p. 118).

makes ‘Aristotle’ mean a property that can be shared by clones of Aristotle—entities that are numerically different from (the original) Aristotle, but share with (the original) Aristotle the property of being an instance of Aristotle. What makes it the case that ‘Aristotle’ means the original, as opposed to the clones? Fodor could insist on treating the property [being an instance of Aristotle] as unique for an individual. But this is dubious: as Adams (2003, p. 156) points out, “why would anyone ever have thought that individuals do not feature in laws?” Allowing such properties and laws between such properties entails a vast profligacy of laws—surely not the result we want if we are after a lean ontology.

If nomic relations, by themselves, cannot explain how representations of particulars get their content, then one could appeal to my account of how representations of particulars get their content. I briefly recapitulate it below:

In Chapter 3 I argued that ADA, when supplemented with conceptual role semantics, can account for how representations of particular objects refer. The idea is to use conceptual role semantics to account for the logical and non-logical roles of representations of particular objects in the visual system. The relations of particular interest are those between the representations of particular objects and representations of spatio-temporal relations. The representations of super-determinate spatio-temporal relations that I appeal to—$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$—co-vary, in the way Fodor suggests, with the property of being a spatio-temporal relation at time t such that: a) it involves the objects x, l, n, and r forming a pyramid with sides of length d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx}; and b) r, l, and n form the base of the pyramid and r and l are my eyes, while n is my nose:

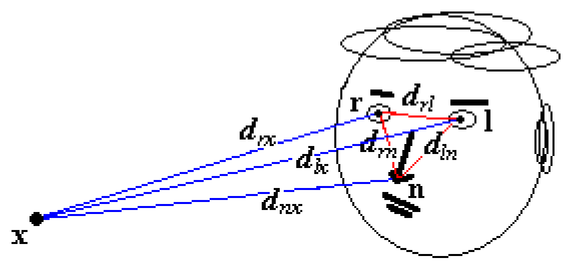

Figure 5.4.: This is an instance of a configuration involving four objects x, l, n, and r forming a pyramid with sides of length d_{rx}, d_{lx}, and d_{nx} and a base of a triangle with sides d_{rl}, d_{rn}, and d_{ln}. The configuration instance is instantiated so that l, r, and n are the left, right eye, and the nose of "Pinocchio" here.

From now on, for brevity, I will use "'C_j" for the predicate "'C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)'".

Importantly, the content of the predicate 'C_j' is not defined entirely *á la* ADA, but is inferentially and functionally related to a first-person reflexive term ('I') and an indexical for time ('t'). It is this characteristic that complements and helps ADA here: in virtue of the predicate being related to 'I' and 't' it doesn't pick out just any pyramidal super-determinate configuration of type C, instead it picks out a particular instance of it: *the* one that is set on me at t. This makes the referent of the predicate 'C_j' a trope. 'C_j' is nomically related to the property of being a pyramidal super-determinate configuration but is *not* nomically related to a trope. Instead, the trope is *picked out* in virtue of the ‘C_j’’s nomic relations with the property of being the super-determinate spatial configuration C_j *and* in virtue of its inferential and functional relations with ‘I’ and ‘t’. The picking out of a trope is cashed out as the satisfaction relation that obtains between Fregean definite descriptions and their referents. In short: when ‘I’ and ‘t’ in the predicate ‘$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, O)$’ are saturated, then the referent of this predicate is a trope.

I suggest that we use the predicate ‘$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, O)$’ in the following definite description in virtue of which the symbol ‘x’ gets to refer to a particular object:

$$(\boldsymbol{E}x)((y)(C(y, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{ry}, d_{ly}, d_{ny}, t, O) \leftrightarrow (x = y)))$$

Put crudely, this reads: right now x is the one and only particular that has the relational spatial property of being at the specified distances with respect to my nose and eyes.

This approach makes use of conceptual role semantics for more than the logical apparatus needed for the definite description. Conceptual role semantics is also used to define the content of ‘I’ (the first-person reflexive term) and is perhaps needed also for the content of the temporal predicate ‘t’.

According to this view representations of particular objects get their content *via* definite descriptions *without* the definite descriptions being constitutive of the content of the representations of particular objects. *If* the capacity to express definite descriptions presupposes concepts, then my view implies that FINSTs are conceptual representations and that they get their content *via* conceptual representations of spatio-temporal representations. Now, *if* the capacity to express definite descriptions presupposes concepts, then my account is barred for Pylyshyn. He claims that FINSTs are non-conceptual and that they refer and track independently of conceptual representations of properties.[142]

Now, the debate on what exactly concepts are and whether there are such things as non-conceptual representations is messy and there is no consensus as to what concepts are and whether there are such things as non-conceptual representations. Some, like Fodor (2008), would grant that if one appeals to definite descriptions with predicates, argument variables, and quantifier operators, then one is appealing to conceptual representations.[143] My claim is that *if* being able to compose

---

[142] *Cf*. Pylyshyn (2007, p. x): "I have proposed that the capacity to individuate and track several independently-moving things is accomplished by a mechanism in the early vision module that I have called FINSTs…. This primitive nonconceptual mechanism functions to identify, reidentify and track distal objects."

[143] Fodor would agree that being able to compose definite descriptions is a sufficient condition for being a concept. Fodor (2008) argues that being a predicate entails having the property of being a conceptual, as opposed to non-conceptual representation. Fodor (2008, Ch. 6) assumes that conceptual and non-conceptual representations are compositional (p. 171)—i.e. their syntactic and semantic

definite descriptions is a sufficient condition for being a concept, then, obviously, my account of how particular objects get represented appeals to concepts and the representation of an object itself is conceptual. However, one can deny that being able to compose definite descriptions is a sufficient condition for being a concept. For example, as Rey (p. c.) has pointed out, one may argue that what distinguishes conceptual from non-conceptual representations is that the former can enter into indefinitely many novel combinations, while the latter, being produced by modules that are encapsulated from and inaccessible to the rest of cognition, enter in severely restricted number of combinations. Thus, it is possible for one to argue that non-conceptual representations can compose definite descriptions. Or alternatively, one can insist, as Fodor and Pylyshyn do, that that a) conceptual representations are those that represent an entity *as* being X; and b) conceptual representations are to be contrasted with non-conceptual representations that represent an entity without conceptualizing it *as* being X. For both Fodor and Pylyshyn non-conceptual

---

properties are determined by the syntactic and semantic properties of their parts. One of the distinctions between the two types of representations is that they compose differently. Conceptual representations have canonical decompositions: they have a constituent structure (logical form) in virtue of which they have canonical parts that have semantic and syntactic properties. Their non-canonical parts (obtained by decomposing the conceptual representations in a manner different from that defined by their constituent structure) have no semantic and syntactic properties. Non-conceptual representations don't have canonical decompositions as a result of which all of their parts have semantic and syntactic properties (p. 173). Because non-conceptual representations don't have canonical decompositions, all of their parts are "homogeneous" (p. 174-175)—they are of the same type, as opposed to of different types as are the parts in conceptual representations: predicates, arguments, and logical constants.

representations carry Dretskian information about X without representing it as X, or as anything else. Now, as long as a definite description only picks out the content of a symbol without also constituting the content of the symbol[144], it seems possible that definite description can be built from logical constants and symbols that do not represent *as*.

Defining what concepts and non-conceptual representations are is beyond the scope of this dissertation. As a result I am prepared to grant that the mere ability to compose definite descriptions is not a sufficient condition for conceptual representations and either the ability to enter into indefinitely many novel combinations or the ability to represent as[145], or both, or some other condition is such

[144] If the definite description were to constitute the content of a symbol, then it would count as a mode of presentation and modes of presentation are standardly taken to be a mark of representations that represent *as*.

[145] I think, however, that there is a sense in which non-conceptual representations can also represent as. The argument is as follows: Information-carrying non-conceptual representations come in two types: demonstrative and non-demonstrative (robust). The content of demonstrative representations depends on the particular circumstances of their tokening. That is to say that in the case of demonstrative representations we are concerned with a particular cause causing a token symbol. However, as I have argued in Chapter 1, if such demonstrative representations are unmediated, then we get the "which link" problem. This suggests that we need demonstrative representations to be mediated by robust ones. However, we have the disjunction problem for robust representations and to avoid it, one has to appeal to ADA. In appealing to ADA one commits to treating a representation as having veridicality conditions: the nomic relations on which all the other nomic relations depend, but not *vice versa*, pick out a symbol's referent that is veridical, whereas all other symbol's causes are misrepresented *as* the

a condition. In that case, the argument in this dissertation would have to be scaled back to just establishing that to represent particular objects, we need logical form rich enough to express definite descriptions. This does not permit me to criticize Pylyshyn's position that FINSTs are non-conceptual and that they are not mediated through conceptual representations. However, the more important point that I am making in this dissertation is about the logical structure needed to represent particular objects. If this structure turns out to be sufficient for concepts, then I can also point out that, contrary to Pylyshyn, representations of objects are conceptual and are mediated through conceptual spatio-temporal representations.

Obviously, Pylyshyn can deny that definite descriptions are needed to determine the content of a FINST. If he is right, then it wouldn't matter if the capacity to form definite descriptions presupposes concepts. I would no longer be able to argue that, since FINSTs get their content *via* definite descriptions and such descriptions presuppose concepts, FINSTs are conceptual and conceptually mediated. Below I discuss this objection in light of what has been said in the previous chapters:

5.B.III.1. Logical form:

---

referent and as such are non-veridical. This means that there are two modes of presentation of an entity—a veridical (*via* representation 'P') and a non-veridical one (*via* representation 'Q'). Modes of presentation are standardly taken to be a mark of representations that represent *as*. Since robust non-conceptual representations have modes of presentation, then they represent *as*.

I am now going to argue that abandoning the logical form of definite descriptions would make one unable to represent particular objects.

Suppose that Pylyshyn goes with Fodor (2008) and claims that heterogenous syntactic role is a sufficient condition for being a concept. Essentially, the idea is that if representations have at least two types of syntactic roles (e.g. of predicates and of arguments), then they count as concepts.[146] Defending the view that FINSTs are non-conceptual would then mean that in giving an account of FINSTs, Pylyshyn cannot appeal to predicates and arguments. This entails that appealing to definite descriptions is out of the question.

Despite this, Pylyshyn can argue that he can make use of my account by substituting my conceptual representations (predicates and arguments) with registers that carry information (non-conceptual representations).[147] Thus, just like tokens of the predicates in my definite descriptions enter in singular causal relations with the visual system tokening them and in virtue of this they specify it implicitly, so would the registers. Also, just like the temporal indexical in my account, Pylyshyn's registers would carry information about the time they are tokened in virtue of causal relations with a register and a pacemaker. If this is on the right track, then my account is impotent in trying to show that a) the logical form of definite descriptions is needed to represent particular objects; and b) FINSTs are conceptually mediated

[146] *Cf.* Fodor (2008, p. 171-175).

[147] I am using here "register" to mean a representation that is not conceptual

representations of particular objects. In essence, Pylyshyn can argue that I haven't provided a reason why predicates (conceptual representations) are needed, as opposed to registers (non-conceptual representations).

Let's suppose that my account is modified in the way proposed above. One of the results of the above modification is that the registers are not going to be concatenated with FINSTs as predicates. If they were, then they would be predicates and the FINST would be an argument. That would make both be conceptual representations—contrary to what Pylyshyn wants—*if* we go with Fodor's way of defining the conceptual/non-conceptual distinction (which for the moment we are assuming). So, instead of logical concatenation, the registers are merely going to be functionally related to FINSTs. Of course, one cannot argue that a definite description specifies the content of a FINST, because that would be appealing to the predicates and arguments that form the definite description. Instead, one must give a causal story where a FINST gets its content through representationally unmediated or mediated relations with its referent. If a FINST were to get its content unmediated by registers (something that Pylyshyn does not believe is the case[148]) then the "which link" problem would obviously arise. This suggests that a FINST gets its content mediatedly through registers: the content of a FINST would be the information provided by the registers. However, this information would not be specific enough to specify a particular object because of the "which link" problem. It can only specify a

[148] *Cf*. Pylyshyn (2001b, p. 145): "In assigning indexes, some cluster of visual features must first be segregated from the background or picked out as a unit".

property. However, then the FINST would end up being a representation of a property, not of a particular object, and what we are after is representations of particular objects.

Pylyshyn can try to argue that 'x' (a token FINST) a) carries information about the system tokening it and the time of its tokening in virtue of being related to a system *via* singular causal relations; and b) carries information (*via* its causal relation with a register) about the property of being at the tip of a spatio-temporal configuration such that: i) it involves the objects x, l, n, and r forming a pyramid with sides of length d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx}; and ii) r, l, and n form the base of the pyramid and r and l are a system's eyes, while n is a system's nose. While we are at it, let's also suppose that the FINST carries information (*via* registers) about shapes, surface reflectances, and Spelke properties (being bounded, unified, and persisting in time). Haven't we, in virtue of having the FINST carry so much diverse information, specified a particular already? The answer is that we haven't. Nothing guarantees that all of these properties are about the same particular object—'x' could just as well carry information about independent particular objects. That is to say that the system cannot use 'x' to individuate a particular object.

Also, note that what is not explained is how representations of the properties of object p get bound with each other and not with the representations of the properties of object q (this is known as the "binding problem"). Clark (2004) argues, as I present in Section 5.D, that such binding requires the use of representations with different syntactic roles—arguments and predicates. Since we are currently assuming Fodor’s account of concepts, we can now argue that such binding requires conceptual representations. Since such binding is required for the FINST to get its content, then it follows that FINSTs are conceptually mediated.

Let us move away from Fodor’s account of the conceptual/non-conceptual distinction. So, *suppose* now that predicates and arguments can count as non-conceptual. This means that the above objections can be discarded: now Pylyshyn can argue that in virtue of being predicates of the same argument variable (at the same time), the *same* entity is represented as having different properties. The idea is that we have the following concatenations conjoined with each other: ‘$L_i x$ ^ Ux ^ $S_k x$’. Here ‘L_j’ refers to the property of being at a particular space-time location with respect to me, ‘U’ refers to the property of being unified, and ‘S_j’ refers to the property of being of a particular shape. ‘$L_i x$ ^ Ux ^ $S_k x$’ reads as “x has the property of being the unified entity with a certain shape at L_i”. In virtue of ‘x’ being concatenated with all of these predicates at the same time, ‘x’ refers to an entity that has all of these properties. Furthermore, suppose that we can store ‘L_i’, ‘U’, and ‘S_k’. When observing an object in our visual field through time, we have a succession of these predicates being stored in memory. Let ‘x’ be concatenated with the predicates stored in memory, as well. For example, let ‘$L_i x$ ^ $L_i x$ ^ $L_k x$’. In virtue of being concatenated with all of them, ‘x’ represents an entity that has different temporal locations (at t_i, t_j, and t_k). Given that a trope or a space-time point cannot be at two space-time places at once, ‘x’ cannot refer to a trope or a space-time point. So, isn’t that enough to treat ‘x’ as a representation of a particular object?

The answer is, “No”. I am willing to grant that the referent of ‘x’ is a particular (after all, I argued that ‘$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$’ refers to a particular—a trope—merely in virtue of ‘C_j’ being related to ‘t’ and ‘I’)[149]. Also, I am willing to grant that ‘x’ does not refer to a trope or a space-time point. However, what I am not willing to grant is that a particular *object* has been represented. To see this note that one reason why I appeal to definite descriptions in my proposal is so that I can say that ‘x’’s referent is the *one and only* entity that has a certain set of properties. Without definite descriptions (and we are currently assuming that they are unavailable to Pylyshyn), one has to allow that more than one object can be at a certain space-time location (which is possible in some ontologies). This would mean that no particular object has been individuated. The upshot is that if Pylyshyn adopts the current proposal, then he would not be able to account how particular objects get represented.

The point above, however, can be objected to, due to the implausibility of there being two objects (wholly) located at the same space-time location. So here is a second argument for the point that without definite descriptions we cannot determine if the referent of ‘x’ is a particular object: First, notice that points (being infinitesimal) can be problematic as the locations of properties—one wonders

[149] I am using ‘C_j’ as an abbreviation of ‘$C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)$’.

whether something square can be at an infinitely small point. Instead, let's treat properties as being instantiated *at regions*. Similarly, instead of treating tropes as properties at space-time points, we can treat them as properties at space-time *regions*. Second, note that when observing an object in my visual field through time, there is a space-time region occupied by the object. At this space-time region we have instantiated the properties of being of a certain shape (S_k) and of being unified (U). We also have instantiated at different spatio-temporal parts of this space-time region the properties of being at a certain location with respect to me (L_j, L_i, and L_k)—since the region is partially at these locations. So, when we have '$L_ix \wedge L_ix \wedge L_kx$', we can treat this as representing the fact that a region is partially at L_j, L_i, and L_k (and wholly at all of them). Alternatively, we can treat this as representing the fact that L_j, L_i, and L_k are instantiated at the region. So, if '$L_ix \wedge L_ix \wedge L_kx$' can be treated as representing the fact that a region is partially at L_j, L_i, and L_k, then the space-time region in question is an available candidate for the referent of 'x'. The upshot is that the current proposal cannot determine whether 'x' refers to a region or an object.

One can object here that the concatenation 'Px' means that x *has* P. Thus, 'L_jx' cannot mean that L_j is instantiated at x. Instead, it must mean that x has the property of being at L_j. But, the objection continues, a region that is composed of the sub-regions specified by L_j, L_i, and L_k cannot wholly be at any one of them. Alternatively, if 'Px' means that x has P, then one can object that a region cannot *have* the property of being unified, for example. Only objects can have such properties. However, it strikes me that the *syntax* of predicate logic does not make ontological assumptions about regions having properties. It makes assumptions only about there being particulars, properties, and properties being related to particulars. Whether the relation is one of having, being instantiated at, or whatever, is left open. As a result, 'Px' need not mean that x *has* P. It can mean that P is instantiated at x.

In my account, the problem of determining whether 'x' refers to a region or an object is avoided because a) a definite description like $(\boldsymbol{E}x)((y)(L_jy \leftrightarrow (x = y)))$ specifies the content of 'x' as *wholly* present at a space-time region; and b) there are multiple definite description like this one stored in memory (e.g. $(\boldsymbol{E}x)((y)(L_ky \leftrightarrow (x = y)))$). This means that the content of 'x' at t_j and t_k must be an object since a space-time region cannot be wholly present at two different space-time locations.

The upshot from the points above is that unless the ability to form definite descriptions is present, a sensory system would not be able to represent particular objects. If this ability is sufficient for being a concept, then the further conclusion can be derived that representations of objects are a) conceptual; and b) mediated through conceptual spatial representations. But Pylyshyn gives three theoretical arguments as to why FINSTs cannot be mediated through conceptual spatial representations.[150] I criticize them below:

5.B.III.2. The circularity argument:

[150] Pylyshyn treats predicates as concepts. Thus, in what follows I will use "predicate" and "concept" interchangeably.

Pylyshyn's circularity argument seems to be based on Kripke's requirement that a successful theory of reference must not include in the conditions of reference the notion of reference itself for fear of being circular.[151] Likewise, one would be right to argue against theories that include in the conditions for conceptualizing, the notion of conceptualization. Pylyshyn seems to suggest that a similar circularity is committed by a theory where the reference of an argument is explained by predicates:

> there is no such thing as a purely top-down process, or rather, a process cannot be top-down all the way out to the world. If representations are to have a content that is about the world, then the world must impose itself upon the perceptual system – which is to say it must act bottom up at some stage. What I am proposing here is that what is bottom up is what will be needed to produce the predicate-argument pairs that constitute a conceptual encoding of the world (encoding that something has the property P). In order to prevent circularity the arguments of such predicates must be identified (or as I say "picked out") by a process which itself is not conceptual (does not use other predicates or properties in order to identify the referents of the arguments). This desideratum also entails that things that are bearers of properties must be selected and referred to in a bottom-up or data-driven manner. (Pylyshyn (2007, p. 12))

I don't see why a theory where the reference of arguments is provided by predicates needs to be circular. As long as the reference of predicates is explained in a way that does not appeal to the notion of predicate reference, there is nothing circular

[151] *Cf*. Kripke (1980, p. 68-70).

with using predicates to determine the reference of an argument. That is, circularity is avoided if one appeals, for example, to ADA to explain the reference of predicates. Then, using descriptions with predicates and indexicals would be enough to provide the reference of the predicated arguments, as I have argued in Chapter 3.

5.B.III.3. The infinite regress argument:

Pylyshyn argues that if the reference of a term is explained through the reference of another term, whose reference is explained through another term, and so on *ad infinitum*, then we have an infinite regress and we never explain how a term actually gets to be about something in the external world. From this Pylyshyn infers that a) the reference of terms must eventually be explained through an unmediated (direct reference) relation between a term and an entity in the world; and b) it is visual indexes that provide the grounding:

> [w]hile it is clear that you cannot individuate objects in the full-blooded sense without a conceptual apparatus, it is also clear that you cannot individuate them with only a conceptual apparatus. Sooner or later concepts must be grounded in a primitive causal connection between thoughts and things. The project of grounding concepts … in perception remains an essential requirement if we are to avoid an infinite regress. Visual indexes provide a putative grounding for basic objects – the individuals to which perceptual predicates apply, and hence about which cognitive judgments and plans of action are made…Without such a preconceptual grounding, our percepts and our thoughts would be disconnected from causal links to the real-world objects of those thoughts. With

indexes we can think about things…without having any concepts of them: one might say that we can have demonstrative thoughts. (Pylyshyn (2001b, p. 154))

For the purposes of this dissertation I agree with (a): the *reference*[152] of terms is eventually grounded in representationally unmediated directly referential relations with external entities. ADA is an example of a theory where i) predicates are unmediated in the sense that the content of other predicates is not constitutive of their content; and ii) some predicates are unmediated in the sense that no other representations are needed for their tokening. In either case, according to ADA, predicates directly refer to properties. However, I disagree with (b): the infinite regress argument does not show that predicates cannot also provide the grounding. ADA, for example, shows how this can be the case. Then, we could have predicates and visual indexes that provide the grounding.[153]

5.B.III.4. The evaluated predicates argument:

Pylyshyn seems to argue that since only those predicates which are predicated of arguments can play a role in the sentences of the language of thought, then it must be the case that predicates can play a role in the language of thought only after an

[152] Reference is, of course, to be distinguished from content. The content of indexicals, logical operators, and numbers is not their referent (well, in the case of indexicals, not *only* their referent).

[153] I, as evidenced by Chapter 3, also believe that visual indexicals are *causally* mediated by representations (but still directly referential). That is, their content is not mediated *via* the content of other representations, but their tokening is.

argument is available for them to be predicated of. This makes FINSTs—the arguments—logically prior to the predicates:

> some properties do get encoded in the form of predicates, since predicates are properties of indexed things, so FINSTs are logically antecedent to predicates…I will argue that it is a general property of conceptualizations of the perceptual world that only indexed objects can serve as arguments of predicates and consequently only properties of selected objects are conceptually encoded. (Pylyshyn (2007, p. 43))

However, we can agree with Pylyshyn that for predicates to play a role in the language of thought they must be concatenated with an argument, without agreeing that the argument must be a FINST—a visual index that *already* refers to a visual object in the external world. As I have argued in Chapter 3, the argument may get its content *via* its concatenation with predicates (which, of course, does not entail that the content of the predicates is constitutive of the content of the argument). Thus, even if one grants that the tokening of an argument is logically prior to (i.e. independent of) the tokening of a predicate[154], one can still hold that the reference of the argument is logically posterior to (i.e. dependent on) that of the predicate. And it is the latter point that is of interest here.

5.C. FINSTs, FINGs, and the tracking problem:

[154] I don't see why one should grant that, though—the arguments and predicates can be tokened together, so neither tokening would be logically prior.

If the above is on the right track, then Pylyshyn's theoretical arguments do not show that concepts cannot play a role in grounding thought. But Pylyshyn's most powerful argument for his claim that conceptual representations of properties play no role in MOT is an empirical one: a simulation by Pylyshyn purportedly shows that tracking by conceptual location representations is unviable. The simulation and the solution to the tracking problem that I presented in Chapter 3—a solution based on conceptual representations of locations—are practically identical. In what follows I will first present and criticize Pylyshyn's FINST based solution to the tracking problem (Section 5.C.I) and then I will present and rebuff his empirical argument (Section 5.C.II).

5.C.I. Pylyshyn's solution to the tracking problem:

According to Pylyshyn, the tracking problem—the problem of associating an object at time t_1 with an object at time t_2[155]—is solved by the FINST mechanism in virtue of the fact that FINSTs normally (i.e. barring mistakes) keep being connected to the same FING through time through a causal chain linking the two. Pylyshyn hypothesizes two possible mechanisms that play a role in that:

[155] I will leave the problem of associating an object at time t_1 with multiple objects at time t_2 for another time.

> Any early vision system will contain sensors and a way of clustering features (e.g. Marr, 1982). In order to maintain the identity of moving clusters (i.e. to implement a 'sticky' binding) all one needs is a mechanism that treats time-slices of clusters that move continuously over the retina as the same cluster. It could do so, for example, by following the rule that if the majority of the elements in a cluster (represented, for example, in a 'list of contributing points') continue to be present in a succeeding cluster then consider both clusters to be the same. Or alternatively, one could simply spread the activation arising from a cluster of elements to neighboring elements, thereby favoring the activation of nearby regions and so favoring continuously moving clusters. This is essentially the technique suggested by Koch and Ullman (1985) in their proposal for a neural implementation of attentional scanning. (Pylyshyn (2001b, p. 147, ff. 7))

The neural implementation in question is a winner-take-all network:

> The network I have in mind uses a slightly modified version of what is known as a fully connected Winner-Take-All (WTA) network. Such a network has the property that when provided with an array of inputs that vary in their activation level (i.e., that have different magnitudes of inputs) the network settles into a state in which the unit with the most highly active input retains its value while the activity level of all the others is reduced to zero. This is, in fact, a maximum-finding circuit. (Pylyshyn (2006, p. 167))

Pylyshyn's account of tracking is (more or less) as follows: a FING causes a token of a FINST. The direction (in retinal coordinates) in which the FING is lying is registered, but not conceptually represented, *via* sensors on the retina. The distance from the eyes to the FING is likewise registered, but not conceptually represented, through sensors that register convergence of the eyes[156]. The causal chain leading from the FING to the token FINST passes through a neural array where neurons (buffer units) are registering information about the object. These neurons are connected with other neurons so that they have neighboring receptive fields. The neurons currently being active stimulate the neurons they are connected with. As a result the latter neurons will get even more active, once they are stimulated by a FING, than a neuron that is stimulated by a FING without being already "primed". This would make the token FINST be more likely to be activated by the neurons close by the neurons that had been activated. In this way a causal chain between the FING and the token FINST is maintained and tracking is performed. Note how for Pylyshyn, the tracking of FINGs, just as referring to FINGs, is accomplished without the help of conceptual representations of properties. While the object's properties may play a role in causing the FINST, they are not (conceptually) represented.

This is an intuitive proposal that seems to work: since the same token FINST is continuously caused by the same FING, in virtue of a causal chain linking the two,

[156] Also, the distance is registered through sensors that register accommodation: *cf*. Palmer (1999, p. 203-4): "Accommodation is the process through which the ciliary muscles in the eye control the optical focus of the lens by temporarily changing its shape. It is a monocular depth cue because it is available from a single eye, even though it is also present when both eyes are used…[T]he lens of the human eye has a variable focusing capability, becoming thin to focus light from faraway objects on the retina and thick to focus light from nearby ones…. If the visual system has information about the tension of the muscles that control the lens's shape, then it has information about the distance to the focused object."

the token FINST tracks the FING. No appeal is made to any conceptual representations but only to registrations. Despite its intuitive appeal, this proposal is problematic. First, it cannot explain visual tracking through occlusion (a case where the visually tracked object is briefly occluded by an intervening obstacle). The reason is that while the visually tracked object is occluded, there is no causal chain connecting tokens of FINSTs with a FING.[157] Furthermore, the proposal suffers from the "which link" problem just as much as Pylyshyn's account of how FINSTs refer to FINGs. Namely, it is not clear why the FINST should be tracking a distal object, as opposed to a proximal one, e.g. a locus of activation on the retina or on the neural array. In fact, it is not clear why there should be any tracking at all—the notion of tracking implies that the *same* object is being tracked. However, because of the "which link" problem, it is indeterminate what a FINST refers to at a moment and therefore from one moment to the next. At time t_1 it could refer to object x (a link in a causal chain leading to the token FINST). At time t_2 it could refer to object y (also a link in a causal chain leading to the token FINST). However, if $x \neq y$, then, since the FINST does not represent the same object at different times, it does not track. Lastly, if Pylyshyn has a hard time explaining why a FINST refers to a particular object, as opposed to a property or a space-time region, as argued above, then it is not clear why a FINST should track at all. Instead it can be treated as merely representing that a

[157] To avoid this problem Pylyshyn allows that during occlusion tracking is performed *via* conceptual representations of location: *cf*. Pylyshyn (2007, p. 40): "[it] seems at least that when tracked targets disappear there is a record of where they were when they disappeared" and *cf*. Pylyshyn (2007, p. 80): "our assumption is that the disappearance itself causes locations to be conceptualized and stored in memory".

property is present. Alternatively, it can be treated as merely individuating a space-time region.

The point is that there are theoretical problems for Pylyshyn's account of tracking by FINSTs. But Pylyshyn thinks that the problems for alternative accounts for tracking, like the one I presented in Chapter 3, are much worse:

5.C.II. Pylyshyn's empirical argument against conceptual location representations:

Pylyshyn takes his experiments[158] to show that conceptual representations of non-spatial properties are not used in tracking. Let's grant this for the moment. Pylyshyn also takes his experiments to show that conceptual representations of *spatial* properties are not used in referring and tracking. If he is right, then my account of how particular objects get represented—according to which particular objects are represented through conceptual spatio-temporal representations—and the solution to the tracking problem presented in Chapter 3—a solution based on conceptual spatio-temporal representations—are mistaken. What I will now argue for is that Pylyshyn's argument is not conclusive.

[158] *Cf.* Pylyshyn and Storm (1988), Scholl, Pylyshyn, and Franconeri (1999), and Dennis and Pylyshyn (2002).

Pylyshyn considers a mechanism for tracking through conceptual location representations[159]: the locations of four targets are represented (encoded) and recorded (perhaps in something like Treisman's object files[160]). Then, these conceptual representations of location are updated continuously by moving focal attention to each target in turn: if target x has moved, the representation of location in the object file corresponding to x is updated with a representation of the location of the nearest object to x's previously represented location.[161] However, Pylyshyn dismisses this solution because the simulation of this strategy when using data about the movements of the dots from an actual MOT experiment yielded

> a predicted performance of only about 30 percent under the most conservative conditions—that is, using the highest estimates of attentional speed reported in the literature and even considering the possibility that not just location but also speed and direction of each target are also encoded to enable some degree of prediction of the targets' location. This is far from the 87 percent performance we actually observed with our volunteer subjects. (Pylyshyn (2007, p. 36-37))

---

[159] *Cf.* Pylyshyn (2007, p. 36-37).

[160] *Cf.* Kahneman, D., Treisman, A. and Gibbs, B. J. (1992).

[161] More specifically the process is something like the following: 1) Create object file α with a location representation at time t_1. 2) For any object file β such that it is a) created after t_1 but within time t; and b) its location representation specifies a location whose distance to the location specified by α's location representation is the shortest, but also shorter that x, then delete the location representation in α, copy β's location representation in α, and delete β.

Hill (2008), however, points out that this argument does not work against all accounts where representations of location are used in MOT. The accounts immune to Pylyshyn's argument are those where FINSTs are object files (*á la* Treisman) and where each FINST tracks in virtue of a location representation in the object file, which gets continuously and automatically (i.e. in parallel and without the use of focal attention) updated:

> [Pylyshyn's] only truly decisive argument against [the hypothesis that MOT requires representations of location], stated on p. 37, does not rule out all versions of it. In particular it has no force against the version which claims (i) that subjects track moving objects by deploying object files, (ii) that object files contain representations of locations, and (iii) that these representations are continuously and automatically updated. (Hill (2008))

The accounts which Pylyshyn dismisses are ones where the locations of the targets are represented and recorded and where these representations of location are updated *by moving focal attention to each target in turn*. Focal attention is a *serial* process and it is because of that that Pylyshyn gets the abysmal predicted performance of only 30 percent (compared to the 87 percent performance observed with actual subjects)—it takes time for focal attention to move from one object to another and the longer that time is, the greater the chance is that a non-tracked object can be the one closest to the previous location of a tracked one (and so be mistaken for the tracked object). However, accounts where representations of locations are represented, recorded, and updated in *parallel* are immune to Pylyshyn's criticism.

Pylyshyn accepts that one cannot show with a great degree of certainty that such alternative explanations are wrong:

> one can't exclude all logically possible alternative processes for achieving these results. For example, we cannot exclude the possibility that location encoding occurs in parallel at each tracked object and then serially allocated focal attention is used for tracking, or that four parallel "beams of attention" independently track the four targets. (Pylyshyn (2001b, p. 143, ff. 3))

But what then are Pylyshyn's grounds for preferring his own view? Hill conjectures that Pylyshyn insists on his hypothesis on methodological grounds: it is simpler since it doesn't require representing locations. But Hill argues that Pylyshyn's proposal is not in fact simpler, as Pylyshyn intends: Pylyshyn allows for representations of locations to play a role in tracking:

> [it] seems at least that when tracked targets disappear there is a record of where they were when they disappeared. (Pylyshyn (2007, p. 40))

Pylyshyn's idea seems to be that normally tracking does not need to involve representations of locations, but in some cases—occlusion—they are involved. Pylyshyn claims that representations of locations are created only when objects disappear:

> our assumption is that the disappearance itself causes locations to be conceptualized and stored in memory. (Pylyshyn (2007, p. 80))

Hill agrees that this is a possible interpretation of the results, but he points out that since the FINST mechanism also appeals in some cases to representations of location, then it may not be simpler than mechanisms that use representations of locations to track in all cases:

> Pylyshyn seems not to realize that a mechanism that detects disappearances and then creates representations of last known locations (perhaps by drawing on iconic memories of the locations) may not be simpler than a mechanism that creates representations of locations at the outset of a tracking venture and then automatically updates them. In other words, his interpretation of tracking across disappearances adds significantly to the complexity of his initial FINST-based hypothesis about tracking, and may thereby undercut the methodological argument for its correctness. (Hill (2008))

The upshot is that Pylyshyn's argument on methodological grounds is not persuasive since it at most shows that representations of locations need not play a *key* role in tracking. Furthermore, the argument based on Pylyshyn's empirical findings does not exclude all variants of accounts where spatial representations are used for tracking. Thus, what Hill (2008) calls Pylyshyn's "only truly decisive argument" against the use of representations of location in tracking leaves open the possibility that representations of locations are used in tracking.

If the above is on the right track, then Pylyshyn has not presented a conclusive argument against my account of how particulars get represented and the tracking mechanism proposed in Chapter 3.

5.D. Connections with Clark's theory of sentience:

Throughout this dissertation I have appealed to representations of spatio-temporal relations to account for how particular objects are represented. Before I conclude, I will present Clark's argument as to why such representations are needed for representing objects. The gist is that to represent objects we first need to segregate some cluster of visual features from the background. To do so, however, requires solving the binding problem, which in turn requires concatenating predicates to arguments, where the arguments refer to regions. I give an account of how regions get represented and this account appeals to representations of spatio-temporal relations. The upshot is that Clark's argument can be read to fit with and support my account of how particular objects get represented.

According to Clark, our visual system must solve the binding problem if it is to distinguish between a glossy red object next to a matte green one and a matte red object next to a glossy green one. Clark argues that to solve the binding problem a system needs at least two syntactically different types of symbols. Suppose that a system had only one type—predicates. The problem is that simply conjoining predicates together would not work: 'glossy' and 'red' and 'matte' and 'green' does not settle which representations are bound together—do we have ('glossy' and 'red') and ('matte' and 'green'), or ('glossy' and 'green') and ('matte' and 'red')?[162]

One can attempt to appeal only to predicates by appealing to *primitive* representations of properties like 'glossy-red', 'matte-green', 'glossy-green', and 'matte-red'. In other words, one can try to solve the binding problem not by figuring out conditions for conjoining 'glossy' and 'red' but by appealing to non-complex representations that refer to complexes of properties like glossy redness. However, this would entail that for any combination of properties we would need to have a distinct representation. This would lead to an indefinite number of representations, which is problematic. What makes this implausible is not only the profundity of symbols, but the fact that minds are assumed to be instantiated in computational systems with *finite* alphabets of primitive symbols. Having a mind with an indefinite number of primitive symbols threatens this presumption.[163]

---

[162] *Cf.* Clark (2000, p. 72): "we need two different kinds of place-holders in any schema describing the contents of sensory experience. It cannot be collapsed to a univariate form. We cannot capture those contents by substituting different qualities *Q* in a schema of the form 'appearance of qualities *Q*.' Instead we need two place-holders: 'appearance *of* qualities *Q at* region *R*.'

[163] I say "threatens" and not "disproves" because the fact that a set has indefinite number of members does not entail that the set is infinite. Thus, the problem of indefinite number of demonstratives is a problem of implausibility as opposed to a problem of incoherency between assumptions.

One can also attempt to appeal only to predicates by appealing to Quine's (1960). In Chapter 2 I mentioned Quine's (1960, p. 343-347) argument that we can transform a sentence with singular terms/arguments/variables into a sentence without one *via* six combinators. However, as Robin Clark (p. 3-4) points out, what Quine actually shows is that we can eliminate only the *syntactic* contribution of variables. The *semantic* contribution remains intact. That is to say that symbols corresponding to variables in a language do not need to be explicitly tokened/written. Whatever their syntactic contribution to a sentence is, it can be substituted *via* the use of Quine's combinators. But this does not entail that we can eliminate the *semantic* contribution of variables. The semantic contribution is needed by Quine's own admission to define the arity of predicates. So we still need variables.

Clark argues that what is needed to solve the binding problem is to *identify* the thing that is glossy with the thing that is red. To do so, we need the identity operator and a term that is syntactically and semantically distinct from representations of properties—a singular term:

> [t]o get to identity statements we need to add a new *kind* of term, with a distinct function. These are singular terms, names or terms like names, that are used to identify. (Clark (2004, p. 8))

The singular term that Clark introduces to solve the binding problem is a term that refers to places (space-time regions):

> Terms for features and terms for places must play fundamentally distinct and non-interchangeable roles, because otherwise one could not solve the binding problem. (Clark (2004, p. 6))

The reason Clark appeals to singular terms that refer to places is to account for the processing that goes on in early vision. Essentially, different properties are registered in different "feature maps". Such maps register a given property and its spatial relation to the observer. We can think of their outputs as composed of a representation of a property concatenated with a singular term—a representation of a place. Treisman and Gelade (1980) suggest that there is also a "master map"—a map that detects when two feature-maps token singular terms that refer to the same region. It is *via* the "master map" that two singular terms are identified with each other.

Clark (2004, p. 13) points out that the binding problem must be solved prior to FINSTs being tokened, according to Pylyshyn's own admission:

> In assigning indexes, some cluster of visual features must first be segregated from the background or picked out as a unit. (Pylyshyn (2001b, p. 145))

As a result Clark (2004, p. 13) takes himself to be giving an account of an "earlier and simpler kind of direct reference" than the visual indexing of FINSTs:

> Proto-objects will be the values of variables in one layer of visual representation. That layer is above feature-placing but below the apparatus of individuation and sortal concepts found in a natural language. (Clark (2004, p. 1))

Given my account of how particulars are represented, I am sympathetic to Clark's account. I agree that we need representations prior to FINSTs and more specifically that the representations needed must include spatial ones. I also agree with his argument that we need both predicates and arguments in order to solve the binding problem and that the arguments refer to regions. Now I will point out how his view can be read to fit with and support mine:

In my account of how particular objects get represented I appealed to representations of super-determinate pyramidal configurations—'C(x, l, n, r, d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, d_{nx}, t, I)' (I will use 'C_j' for brevity). The configuration is such that: a) it involves the objects x, l, n, and r forming a pyramid with sides of length d_{rl}, d_{rn}, d_{ln}, d_{rx}, d_{lx}, and d_{nx}; and b) r, l, and n form the base of the pyramid and r and l are a system's eyes, while n is a system's nose. Now, the concatenation 'C_jx' represents that x is at the tip of the pyramid. This means that in virtue of being concatenated with 'C_j', 'x' can be used a representation of a space-time location—the one at the tip of the pyramidal configuration C_j set on me, now. 'x' can also represent a region when it is concatenated with a multitude of such predicates: 'C_jx', 'C_ix', and 'C_kx'. Suppose that 'x' is also concatenated with a non-spatio-temporal property, as in 'Rx', where 'R' refers to redness. Then, 'Rx' will denote a feature at a region. Finally, note that according to my account of how particular objects are represented, as long as 'x' doesn't enter into definite descriptions that are stored in memory, 'x' would not count as a representation of an object. Thus, it is the fact that 'x' doesn't enter into definite descriptions that are stored in memory that makes 'x' refer to a region, as opposed to an object.

This shows that my account can provide an explanation of how representations of features at regions get their content and it does so in virtue of two syntactically different types of symbols—much as on Clark's account. Also, given that there are no definite descriptions needed for representing features at regions, I can hold, as Clark does, that the level of representation of features at regions is below that required for representing objects. More importantly, however, above I have given an account of how representations of regions get their content—something which Clark does not do. According to this account, representations of regions get their content *via* representations of spatio-temporal relations ('C_j'). Thus, Clark's view can be developed to appeal to such spatio-temporal representations to explain how regions are represented and, *via* the terms that represent regions, to explain how the binding problem is solved. Since solving the binding problem is a necessary requirement for representing objects in the sensory systems, Clark's view can be read as supporting my position that representations of spatio-temporal relations are necessary to represent objects.

5.E. Conclusion:

In this chapter I have argued that if Pylyshyn's view is read as an account where FINSTs get their content *via* singular causal relations with FINGs, then Pylyshyn's view runs into the "which link" problem. If Pylyshyn's view is read as an account where FINSTs get their content *via* registers that enter into nomic relations with properties of FINGs, then, unless it appeals also to definite descriptions, it has difficulty explaining why a FINST is a representation of a particular object, as opposed to a representation of a region. If one appeals to definite descriptions, however, then on some views of the conceptual/non-conceptual distinction, one has appealed to concepts. This means that if to represent particular objects requires definite descriptions, and if an appeal to the latter entails an appeal to concepts, then FINSTs are conceptual representations and are conceptually mediated (in the sense that their content is specified through other concepts, but without the content of these concepts being constitutive of the content of FINSTs). Pylyshyn's empirical argument—the argument that tracking by conceptual representations of location is unviable—is also shown to be lacking because it does not apply against accounts that appeal to the use of conceptual representations of location in parallel. Pylyshyn's own solution of how tracking takes place is shown to suffer from the "which link" problem as much as his solution of how FINSTs represent FINGs. The overarching conclusion is that Pylyshyn's view that FINSTs are not conceptual and are not mediated by conceptual representations of spatial properties is questionable.

Chapter 6: Conclusion

6.A. Main claims:

I have argued that singular causal strategies of content face a significant problem in the "which link" problem. The latter cannot be resolved by an appeal to triangulation (whether actual or counterfactual). The troubles of singular causal strategies with the "which link" problem constitute a reason to pursue other strategies, which is what I have done in this dissertation.

I have pointed out that ADA faces the problem of representing individuals. I have argued that to explain how particular objects are represented a) the logico-syntactic roles of representations of particular objects need to be taken into account, and for this ADA needs to be complemented by conceptual role semantics; and b) an appeal to the temporal and first-person indexicals must be made, and for this ADA must be supplemented with singular causal relations. In particular, I have argued that a representation of an object 'x' gets its content *via* a definite description that denotes 'x''s referent without constituting 'x''s content. This view makes use of conceptual role semantics for more than the logical apparatus needed for the definite description. It is also used to define the narrow content of 'I' and 't'. I have also argued that particulars are tracked in virtue of an updating and storing of the definite descriptions that determine the content of representations of particular objects.

In my proposal for accounting for the content of representations of particular objects I have appealed to conceptual role semantics to account for predicate logic, representations of numbers, ‘I’, and ‘t’. But conceptual role semantics has been held by many to be deeply problematic because of the holism and the circularity problems. I have argued that by looking at the roles of symbols that are explanatorily basic, the holism problem can be solved. I have also argued that conceptual role semantics need not be circular if one doesn’t use it as the sole account of content.

Finally, I have argued that if Pylyshyn’s view is read as an account where FINSTs get their content *via* singular causal relations with FINGs, then Pylyshyn’s view runs into the “which link” problem. If Pylyshyn’s view is read as an account where FINSTs get their content *via* registers that enter into nomic relations with properties of FINGs, then, unless it appeals also to definite descriptions, it has difficulty explaining why a FINST is a representation of a particular object, as opposed to a representation of a region. If one appeals to definite descriptions, however, then on some views of the conceptual/non-conceptual distinction, one has appealed to concepts. This means that if to represent particular objects requires definite descriptions, and if an appeal to the latter entails an appeal to concepts, then FINSTs are conceptual representations and are conceptually mediated. The overarching conclusion is that Pylyshyn's view that FINSTs are not conceptual and are not mediated by conceptual representations of spatial properties is questionable.

6.B. Directions for future research:

I have suggested how spatio-temporal relations such as configurations can be represented. But this leaves open the issue whether these spatio-temporal relations are organized within a space (e.g. Euclidean or Hyperbolic) or whether the mechanisms for spatial perception are fragmented in a way that does not allow one to pick a specific or any geometry. In examining this issue I plan to start with Wagner (2006) and Masrour (2010). So, one direction for future research is to see how space-time is represented.

In this dissertation I have assumed Fodor's metaphysical assumptions. I want to explore how alternate metaphysical assumptions would change the conclusions of this dissertation. Obviously, if one were to adopt a nominalist proposal, then the singular causal theories of content would become much more plausible and this would boost Pylyshyn's case. But adopting regularity theories of nomic relations might also have effects on the conclusion I have drawn—for example, the assumption on which ADA relies is that nomic relations cover causal relations (i.e. they are ontologically fundamental). Questioning this assumption might weaken ADA.

In my proposal of how particulars are represented in nomic accounts I appealed to temporal representations. I would like to explore further this issue beginning with the work of Roberts (1998) and Gallistel (2009).

Some of the implications of my account—as to whether FINSTs are conceptual or not—obviously depend on what "concept" means and whether there are non-conceptual representations. Thus, one of the future research projects would be to relate the views on the conceptual/non-conceptual distinction to my project in a much more specific manner than I have done in this dissertation.

Finally, I have appealed to memory in giving an account of how representations of particular objects get their content. My hunch is that Treisman-style object files can serve as memory files that encode types of features through time (e.g. shape S_i at t_i, S_j at t_j…). I plan to examine this hypothesis starting with the work of Gallistel and King (2010).

Bibliography

Adams, F. (2003). Thoughts and Their Contents: Naturalized Semantics. *The Blackwell Guide to Philosophy of Mind*. Eds. S. Stich and T. Warfield. Blackwell Publishing Ltd.

Adams, F. and Aizawa, K. (1992). ‘X’ Means X: Semantics Fodor-Style. *Minds and Machines*. 2, 175-183.

Adams, F. and Aizawa, K. (1997). Rock Beats Scissors: Historicalism Fights Back. *Analysis*, 57, 273-281.

Armstrong, D. (1997). *A World of States of Affairs*. New York, NY: Cambridge University Press.

Aydede, M. and Robins, P. (2001). Are Frege Cases Exceptions to Intentional Generalizations? *Canadian Journal of Philosophy*, 31(1): 1–22, March 2001.

Arjo, D. (1996). Sticking Up for Oedipus: Fodor on Intentional Generalizations and Broad Content. *Mind & Language*, Vol. 11. No. 3 September 1996, pp. 231-245.

Block, N. (1986). Advertisement for a Semantics for Psychology. *Midwest Studies in Philosophy*. Eds. P.A. French, et. al.. Volume X, pp. 615-678.

Block, N. (forthcoming). Holism, Mental and Semantic. Forthcoming in *The Routledge Encyclopedia of Philosophy*.

Botterill, G. and Carruthers, P. (1999). *The Philosophy of Psychology*. New York, NY: Cambridge University Press.

Burge, T. (1986). Individualism and Psychology. *Philosophical Review*, Vol. 95, No. 1, 1986.

Burge, T. (2009). Perceptual Objectivity. *Philosophical Review*, Vol. 118, No. 3, 2009.

Carruthers, P. (1996). *Language, Thought, and Consciousness*. Cambridge University Press.

Carruthers, P. (2006). *The Architecture of the Mind, Massive Modularity and the Flexibility of Thought*. New York, NY: Oxford University Press Inc.

Chomsky, N. (1995). Language and Nature. *Mind*, 104, p. 1-61.

Clark, A. (2000). *A Theory of Sentience*. Oxford University Press.

Clark, A. (2004). Feature-Placing and Proto-Objects. *Philosophical Psychology*, 17(4), December 2004, 451-477.

Clark, A. (2005). Location, Location, Location. Zenon Pylyshyn Conference, University of Guelph.

Clark, R. (1999). Variables, interpretations, and Quine-like combinators. Presented at the Sixth Meeting on Mathematics of Language, University of Central Florida, July 23-25, 1999

Cole, D. (2009). Review of LOT2: The Language of Thought Revisited. *Minds & Machines* (2009) 19:439–443.

Crick F. & Koch, C. (1990). Towards a neurobiological theory of consciousness. *Seminars in the Neurosciences*, Vol. 2, pp. 263-275.

Davidson, D. (1970). Mental Events. *Experience and Theory*. Eds. Foster and Swanson. London: Duckworth.

Deutsch, D. (1997). *The Fabric of Reality*. Penguin Books.

Devitt, M. (1990). A Narrow Representational Theory of the Mind. *Mind and Cognition*. Ed. W.G. Lycan. Oxford, UK: Basil Blackwell, 1990.

Devitt, M. and Sterelny, K. (1999). *Language and Reality* (2nd edition). Cambridge MA: MIT Press.

Donner, K. (1989). The Absolute Sensitivity of Vision: Can a Frog Become a Perfect Detector of Light-Induced and Dark Rod Events? *Physica Scripta*, Vol. 39, 133-140, 1989.

Dretske, F. (1982). *Knowledge and the Flow of Information*. Cambridge MA: MIT Press.

Egan, F. (2003). Naturalistic Inquiry: Where does Mental Representation Fit in? In *Chomsky and his Critics*. Eds. L. Antony and N. Hornstein. Blackwells, 2003, pp. 89-104.

Evans, G. (1979). Reference and Contingency. *The Monist* 62 (2): 178-213.

Fodor, J. (1975). *The Language of Thought*. Harvester Press.

Fodor, J. (1981). The Mind-Body Problem. *Scientific American*, 1981.

Fodor, J. (1983). *The Modularity of Mind*. MIT Press.

Fodor, J. (1987). *Psychosemantics*. Cambridge, MA: MIT/Bradford.

Fodor, J. (1990). *A Theory of Content and Other Essays*. Cambridge, MA: MIT/Bradford.

Fodor, J. (1991). Replies. *Meaning in Mind*. Eds. B. Loewer and G. Rey. Basil Blackwell.

Fodor, J. (1995). *The Elm and the Expert*. Bradford Books/MIT Press.

Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.

Fodor, J. (2004). Having Concepts: a Brief Refutation of the Twentieth Century. *Mind & Language*, Vol. 19 No. 1 February 2004, pp. 29–47.

Fodor, J. (2004a). Replies to Commentators. *Mind & Language*, Vol. 19 No. 1 February 2004, pp. 29–47.

Fodor, J. (2008). *LOT 2: The Language of Thought Revisited*. Oxford, England: Oxford University Press, 2008.

Franconeri, S., Jonathan, S., and Scimeca, J. (2010). Tracking Multiple Objects Is Limited Only by Object Spacing, Not by Speed, Time, or Capacity. *Psychological Science*, XX(X) 1-6.

Funkhouser, E. (2006). The Determinable-Determinate Relation. NOUS, 2006, VOL 40; # 3, pages 548-569.

Gallistel, C. R. (2009). The Foundational Abstractions. *Of Minds and Language: A Dialogue with Noam Chomsky in the Basque Country*. Eds. M. Piattelli-Palmirini, J. Uriagereka, and P. Salaburu. Oxford University Press.

Gallistel, C. R. and King, A. P. (2010). *Memory and the computational brain: why cognitive science will transform neuroscience*. Blackwell Publishing.

Green, B. (1961). Figure coherence in the kinetic depth effect. *Journal of Experimental Psychology*. 62, 272-282 (1961).

Hill, C. (2008). Review of *Things and Places*. *Philosophy Reviews*, Notre Dame Philosophical Reviews.

Horsey, R. (2000). Meaning Postulates and Deference. *UCL Working Papers in Linguistics 12, 45-64*.

Horwich, P. (1998). *Meaning*. New York, NY: Oxford University Press Inc.

Jeshion, R. (forthcoming). Descriptivism and the Representation of Spatial Location. *Reference and Referring, Topics in Contemporary Philosophy*. Eds. M. O'Rourke and J. Campbell. Volume 10, MIT Press.

Julesz, B. (1991). Early vision and focal attention. *Reviews of Modern Physics*. 63, 735–772.

Kahneman, D. and Treisman, A. (1992). The Reviewing of Object Files: Object Specific Integration of Information. *Cognitive Psychology* 24, 175-219 (1992).

Kaplan, D. (1989). *Themes from Kaplan*. Eds. J. Almog, J. Perry, and H. Wettstein. New York, NY: Oxford University Press.

Kaplan, D. (1989a). Demonstratives. *Themes From Kaplan*. Eds. Almog, Perry, and Wettstein. Oxford University Press.

Kaplan, D. (1989b). Afterthoughts. *Themes From Kaplan*. Almog, Perry, and Wettstein. Oxford University Press.

Kripke, S. (1972). Naming and Necessity. *Semantics of Natural Language*. Dordrecht: Reidel, pp. 253-355.

Levine, J. (2007). Demonstrative Thought. *Mind and Language*, 25 (2): 169-195.

Loewer, B. and Rey, G. (1991). *Meaning in Mind*. Basil Blackwell.

Mackie, P. (2006). Transworld Identity. *Stanford Encyclopedia of Philosophy*. Retrieved 05/08/12, from http://plato.stanford.edu/entries/identity-transworld/

Mansson, J. (1998). A Neural Model of Stereopsis. Lund University Cognitive Studies—LUCS 64. ISSN 1101—8453.

Marr, D. (1982). *Vision.* San Francisco: W.H. Freeman.

Marr, D. and Poggio, T. (1976). Cooperative Computation of Stereo Disparity. *Science, New Series*, Vol. 194, No. 4262. (Oct. 15, 1976), pp. 283-287.

Masrour, F. (2010). The Geometry of Visual Space and the Nature of Visual Experience. Retrieved 05/08/12, from http://web.me.com/fmasrour/My_webpage/Papers_files/Masrour-geometry-2011-11-15.pdf

Neander, K. (1995). Misrepresenting and Malfunctioning. *Philosophical Studies*, 79: 109-41.

Nothdurft, H. C. (1990). Texton Segregation by Associated Differences in Global and Local Luminance Distribution. *Proceedings of the Royal Society of London*, Series B, Biological Sciences, Vol. 239, No. 1296 (Apr. 23, 1990), pp. 295-320.

Palmer, S. (1999). *Vision Science: Photons to Phenomenology*. MIT

Peacocke, C. (1992). *A Study of Concepts*. MIT Press: Cambridge, MA.

Peacocke, C. (1995). Content, Computation, and Externalism. *Philosophical Issues*, Vol. 6, Content (1995), pp. 227-264. Ridgeview Publishing Company.

Perry, J. (1977). Frege on Demonstratives. *Philosophical Review*, 86: 474-97.

Piccinini, G. (2006). Computation Without Representation. *Philosophical Studies* (2008) 137: 205-241.

Pietroski, P. (1992). Intentionality and Teleological Error. *Pacific Philosophical Quarterly* 73 (1992), pp. 267–82.

Pietroski, P. (2010). Concepts, Meanings and Truth: First Nature, Second Nature and Hard Work. *Mind & Language*, Volume 25, Issue 3, pages 247–278, June 2010.

Putnam, H. (1973). Meaning and Reference. *Journal of Philosophy* 70, 699-711.

Pylyshyn, Z. (1984). *Computation and Cognition*. Cambridge, MA: MIT Press.

Pylyshyn, Z. (1999). Is Vision Continuous with Cognition? The case for Cognitive Impenetrability of Visual Perception. *Behavioral and Brain Sciences*, 22 (1999), 341-365.

Pylyshyn, Z. (2001a). *Seeing and Visualizing: It's Not What You Think*. Bradford Books/MIT Press.

Pylyshyn, Z. (2001b). Visual indexes, Preconceptual Objects, and Situated Vision. *Cognition*, 80 (1/2), 127-158.

Pylyshyn, Z. (2007). *Things and Places: How the Mind Connects with the World*. Bradford Books/MIT Press.

Pylyshyn, Z. and Sears, C. (2000). Multiple Object Tracking and Attentional Processing. *Canadian Journal of Experimental Psychology*, 2000, 54 (1), 1-14.

Pylyshyn, Z. and Storm, R. W. (1988). Tracking Multiple Independent Targets: Evidence for a Parallel Tracking Mechanism. *Spatial Vision*, 3(3), 1-19.

Quine, W. (1951). Two Dogmas of Empiricism. *The Philosophical Review*, 60: 20-43.

Quine, W. (1956/76). Carnap and Logical Truth. *Ways of Paradox and Other Essays*. 2nd ed., Cambridge, MA: Harvard University Press.

Quine, W. (1960a). *Word and Object*. MIT Press

Quine, W. (1960b). Variables Explained Away. *Proceedings of the American Philosophical Society*, Vol. 104, No. 3, (Jun. 15, 1960), pp. 343-347.

Recanati, F. (1990). Direct Reference, Meaning, and Thought. *Nous*, Vol. 24, No. 5 (Dec., 1990), pp. 697-722.

Recanati, F. (2003). Indexical Concepts and Compositionality. *Two-Dimensionalism*. Ed. J. Macia. Oxford University Press.

Rey, G. (1995). Keeping Meaning More in Mind. *Intellectica*, 1995/2, 21, pp. 65-80.

Rey, G. (1997). *Contemporary Philosophy of Mind*. Cambridge, MA: Blackwell Publishers Inc.

Rey, G. (1998). A Narrow Representationalist Account of Qualitative Experience. *Nous*, Vol. 32, Supplement: Philosophical Perspectives, 12, Language, Mind and Ontology (1998), pp. 435-457.

Rey, G. (2003). Just Where is the Explanatory Gap? Dubrovnik conference paper.

Rey, G. (2004). Fodor's Ingratitude and Change of Heart? *Mind & Language*, Vol. 19 No. 1 February 2004, pp. 29–47.

Rey, G. (2005). Mind, Intentionality, and Inexistence. *Croatian Journal of Philosophy*, vol. 5, no15, pp. 389-415.

Rey, G. (2006). Conventions, Intuitions and Linguistic Inexistents: A Reply to Devitt. *Croatian Journal of Philosophy* 2006, vol. 6, no 18 (p. 230).

Rey, G. (2007). The Analytic/Synthetic Distinction. *Stanford Encyclopedia of Philosophy*. Retrieved 05/08/12, from http://plato.stanford.edu/entries/analytic-synthetic/

Rey, G. (2009). Concepts, Defaults and Internal Asymmetric Dependencies: Distillations of Fodor and Horwich. *The* A Priori *and Its Role in Philosophy*. Eds. N. Kompa, C. Nimtz, and C. Suhm. Paderborn: Mentis (2009), pp185-204.

Roberts, S. (1998). The Mental Representation of Time: Uncovering a Biological Clock. *An Invitation to Cognitive Science* – 2nd ed. The MIT Press.

Russell, B. (1917). Knowledge by Acquaintance and Knowledge by Description. *Mysticism and Logic,* paperback edition. Garden City, NY: Doubleday, 1957.

Saarinen, J. (1996). Localization and Discrimination of "Pop-out" Targets. *Vision Research*, 36(2), 313-316.

Sagi, D. and Julesz, B. (1985). 'Where' and 'what' in vision. *Science*, 228, 1217-1219.

Schiffer, S. (1978). The Basis of Reference. *Erkenntnis*, 13: 171-206.

Schiffer, S. (1981). Indexicals and the Theory of Reference. *Synthese*, 49: 43-100.

Scholl, B. J., Pylyshyn, Z. W., and Franconeri, S. L. (1999). When Are Featural and Spatiotemporal Properties Encoded as a Result of Attentional Allocation? *Investigative Ophthalmology & Visual Science*, 40(4), 4195.

Searle, J. (1983). *Intentionality*. Cambridge. Cambridge University Press.

Segal, G. (2000). *Slim Book about Narrow Content*. MIT.

Sober, E. (1982). Why Logically Equivalent Properties may pick out different Properties. *American Philosophical Quarterly*, 19, 183-9.

Spelke, E. (1990). Principles of Object Perception. *Cognitive Science*, 14, 29-56 (1990).

Stich, S. (1983). *From Folk Psychology to Cognitive Science*. Cambridge, Mass.: MIT Press.

Stich, S. and Warfield, T. (2003). *The Blackwell Guide to Philosophy of Mind*. Blackwell Publishers Inc.

Strawson, P. (1959). *Individuals: An Essay in Descriptive Metaphysics*. Routledge, Taylor and Francis Group.

Strawson, P. (1974). *Subject and Predicate in Logic and Grammar*. Ashgate Publishing Limited.

Tooley, M. (1977). The Nature of Laws. *Canadian Journal of Philosophy*, Vol. 7, No. 4 (Dec., 1977), pp. 667-698.

Treisman, A. (1996). The binding problem. *Current Opinion in Neurobiology*, 6: 171-78.

Treisman, A. (1998). The perception of features and objects. *Visual Attention*. Ed. Richard Wright, p. 26-55. Oxford University Press, Inc., 198 Madison Avenue, New York, NY 10016.

Treisman, A. and Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology*, 12, 97–136.

Tye, M. (1989). *The Metaphysics of Mind.* Cambridge: Cambridge University Press.

Wagner, M. (2006). *The Geometries of Visual Space*. London: Erlbaum

Wakefield, J. (2002). Broad Versus Narrow Content in the Explanation of Action: Fodor on Frege Cases. *Philosophical Psychology*, 15: 2, 119-133.

Wakefield, J. (2003). Fodor on Inscrutability. *Mind & Language* 18 (5), 524-537.

Wolfe, J. M. and Bennett, S. C. (1997). Preattentive object files: shapeless bundles of basic features. *Vision Research*, 37 (1), 25-43.