

Abstract

Non-Relational Intentionality

Justin D'Ambrosio

2017

This dissertation lays the foundation for a new theory of non-relational intentionality. The thesis is divided into an introduction and three main chapters, each of which serves as an essential part of an overarching argument. The argument yields, as its conclusion, a new account of how language and thought can exhibit intentionality intrinsically, so that representation can occur in the absence of some thing that is represented. The overarching argument has two components: first, that intentionality can be profitably studied through examination of the semantics of intensional transitive verbs (ITVs), and second, that providing intensional transitive verbs with a nonrelational semantics will serve to provide us with (at least the beginnings of) a non-relational theory of intentionality. This approach is a generalization of Anscombe's views on perception. Anscombe held that perceptual verbs such as "see" and "perceive" were ITVs, and that understanding the semantics of their object positions could help us to solve the problems of hallucination and illusion, and provide a theory of perception more generally. I propose to apply this strategy to intentional states and the puzzles of intentionality more generally, and so Anscombe's influence will be felt all through the dissertation.

In the first chapter, titled "Semantic Verbs are Intensional Transitives", I argue that semantic verbs such as "refers (to)", "applies (to)", and "is true (of)" have all of the features of intensional transitive verbs, and discuss the consequences of this claim for semantic theory and the philosophy of language. One theoretically enriching consequence of this view is that it allows us to perspicuously express, and partially reconcile two opposing views on the nature and subject-matter of semantics: the Chomskian view, on which semantics is an internalistic enterprise concerning speakers' psychologies, and the Lewisian view, on which semantics is a fully externalistic enterprise issuing in theorems about how the world must look for our natural language sentences to be true. Intensional Transitive Verbs have two

readings: a *de dicto* reading and a *de re* reading; the *de dicto* reading of ITVs is plausibly a nonrelational reading, and the intensional features peculiar to this reading make it suitable for expressing a Chomskian, internalist semantic program. On the other hand, the *de re* reading is fully relational, and make it suitable for expressing the kinds of word-world relations essential to the Lewisian conception of semantics. And since the *de dicto* and *de re* readings are plausibly related as two distinct scopal readings of the very same semantic postulates, we can see these two conceptions of semantics as related by two scopal readings of the very same semantic postulates.

In chapter two, titled “Hallucination and the New Problem of Empty Names”, I argue that the problem of hallucination and the problem of empty names are, at bottom, the same problem. I argue for this by reconstructing the problem of empty names in way that is novel, but implicit in much of the discussion on empty names. I then show how, once recast in this light, the two problems are structurally identical down to an extremely fine level of granularity, and also substantially overlap in terms of their content. If the problems are identical in the way I propose, then we should expect that their spaces of solutions are also identical, and there is significant support for this conclusion. However, there are some proposed solutions to the problem of hallucination that have been overlooked as potential solutions to the problem of empty names, and this realization opens new non-relational approaches to the problem of empty names, and to the nature of meaning more generally.

In chapter three, titled “Intensionality is Additional Phrasal Unity”, I argue for a novel approach to the semantics of intensional contexts. At the heart of my proposal is the Quinean view that intensional contexts should, from the perspective of the semantics, be treated as units, with the material in them contributing to the formation of a single predicate. However, this proposal is subject to a number of objections, including the criticism that taken at face value, this would render intensional contexts, which seem to be fully productive, non-compositional. I begin by discussing the concept of the unity of the phrase, and pointing to various ways that phrases can gain additional unity. I then proposes that the intensionality of intensional transitive verbs is best construed as a form of semantic incorporation; ITVs, on their intensional readings, meet all of the criteria for qualifying as incorporating the nominals in their object positions. I then give a semantics for ITVs

that builds on existing views of the semantics of incorporation structures, and gesture at how this can be extended to intensional clausal verbs, including the so-called propositional attitude verbs.

Non-Relational Intentionality

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Justin D'Ambrosio

Dissertation Director: Zoltán Gendler Szabó

December, 2017

Copyright © 2017 by Justin D'Ambrosio
All rights reserved.

Contents

Acknowledgements	viii
0 Introduction	1
0.1 Unpotting the History of Intentionality	1
0.2 The Features and Puzzles of Intentionality	3
0.3 Intensionality as the Mark of Intentionality	9
0.4 Semantics and Metaphysics: Anscombe’s Project	12
0.5 Metaphysical Views and Semantic Views	15
0.6 Locating the Project	24
0.6.1 Semanic Verbs are Intensional Transitives	24
0.6.2 Hallucination and the New Problem of Empty Names	25
0.6.3 Intensionality is Additional Phrasal Unity	26
1 Semantic Verbs are Intensional Transitives	28
1.1 Introduction	28
1.2 Intensional Transitive Verbs	31
1.3 Representational Verbs	35
1.4 Speaker’s Reference and Application are Intensional	37
1.4.1 Experiment 1: Nonexistence	38
1.4.2 Experiment 2: Nonspecificity	40
1.4.3 Experiment 3: Opacity	44
1.4.4 Conclusions	47
1.5 Semantic Reference and Application are Intensional	50

1.6	Technical Terms and Ordinary English	56
1.7	Consequences	57
2	Hallucination and the New Problem of Empty Names	61
2.1	Introduction	61
2.2	Intentionality	64
2.2.1	The Problem of Non-Existence	64
2.2.2	The Anscombian Approach	68
2.3	Reconstructing the Arguments	71
2.3.1	The Argument from Hallucination	71
2.3.2	Empty Singular Reference	81
2.4	Identity of the Arguments	90
2.5	The Space of Possible Responses	94
2.6	Semantic Adverbialism and Intensional Transitive Verbs	98
2.7	Conclusion	104
3	Intensionality and Phrasal Unity	105
3.1	Introduction	105
3.2	Intensionality and Hyphenation	107
3.2.1	Intensional Contexts	109
3.2.2	Hyphenation	113
3.3	What is Heightened Phrasal Unity?	115
3.4	Incorporation and Phrasal Unity	122
3.4.1	Syntax and Morphology	123
3.4.2	Semantic Incorporation	124
3.5	Intensional Constructions are Incorporated	127
3.5.1	Semantic Features	127
3.5.2	Nonexistence	131
3.5.3	Opacity and Restriction	133
3.5.4	Syntactic Features	136
3.6	A Semantic Proposal: Lexical Modifiers	137

3.7	Intensional Clausal Verbs	141
3.8	Conclusion	145

List of Figures

- 1.1 Mean ratings by condition in Experiment 1. Error bars show standard error. 39
- 1.2 Mean ratings by condition in Experiment 2. Error bars show standard error. 42
- 1.3 Mean ratings by condition in Experiment 3. Error bars show standard error. 45

List of Tables

1	Views on Intentionality, Categorized by Content-Type and Relationality . .	23
---	--	----

Acknowledgements

To Alex and Zoltán.

Chapter 0

Introduction

0.1 Unpotting the History of Intentionality

Minds are really the ultimate multi-purpose tools. With them, we manage to think about all sorts of things: existent, nonexistent, true, false, useful, and fanciful. But we don't only think about things; we also desire, hope, fear, need, notice, plan, ponder, love, and suspect things, and even these are just a fraction of the things our minds help us do. But despite how varied our mental activities and lives are, there is something that holds these various mental states and activities together: they are all about, or directed toward things. Some would put this by saying that all of our various mental states have objects, or are directed toward objects. But regardless of the phrasing, this fact seems to give us a clue as to how we should understand mental phenomena generally. Mental states are about, or directed out toward the world, and this directedness or aboutness seems to be at least part of what makes mental states distinctive.

The term “intentionality”, as it is used in modern philosophy, is a technical term that philosophers use to talk about this kind of directedness. Even *qua* technical term, “intentionality” has a long and complex history.¹ The term comes from the Latin word “intentio”, which roughly means “direction toward”, or “striving toward”. The use of “intentio” as a technical philosophical term can be traced back at least to Augustine, who used it in both a practical sense, roughly matching our “intend”, and in a related cognitive sense [Caston,

¹See [Caston, 2001, Perler, 2001b], among other articles in [Perler, 2001a].

2001]. In its cognitive sense, Augustine used “intentio” as a form of striving or will that “directs the faculty of sense toward its object and keeps it fixed on it”. Roughly speaking, Augustine used “intentio” like we use “attention”. However, through the middle ages, use of “intentionality” moved more toward a purely cognitive use, one similar to our ordinary notion of “aboutness” [Pasnau, 1997, Burnyeat, 2001, Perler, 2001a]. Intentionality played a significant role in many scholastic theories of the mind, and was often characterized as the feature of of minds in virtue of which they can come to be about objects.

After the scholastic period, use of the term “intentionality” declined markedly, and didn’t figure prominently in the theoretical vocabulary of philosophers again until being rehabilitated in the late 19th century by Franz Brentano [Brentano, 1973, Caston, 2001, Crane, 1998, 2014b]. It was with Brentano’s rehabilitation that “intentionality” came to have its more-or-less modern meaning, which is closely related to the notion of representation. Roughly speaking, the things that exhibit intentionality are just those things that are representational. So for instance, pictures, paintings, photos, some words, and plausibly all mental states are representational, and so exhibit intentionality. If we are attracted to a view of intentionality on which representations are about their intentional objects, then we can say that all of the representational entities above have intentional objects.

Brentano’s rehabilitation of intentionality had deep ramifications for the philosophy of mind. Brentano held that intentionality was the characteristic feature of the mental: the feature that distinguishes mental phenomena from all other phenomena [Brentano, 1973]. Let’s call this thesis “Brentano’s Thesis” [Crane, 2014c]. In modern terminology, Brentano’s thesis says that all and only mental phenomena are representational. Of course, this formulation can’t be quite right, because there seem to be many non-mental representations. However, there are refinements in the area that make the view much more plausible. My central concern is not with Brentano’s thesis, but with the concept of intentionality more generally. However, Brentano’s thesis plays an important motivating role in any investigation of intentionality, because even the possibility that intentionality can serve as the distinguishing feature of the mental gives intentionality a central theoretical role in investigations of human cognition.

0.2 The Features and Puzzles of Intentionality

Brentano thought that intentionality was characteristic of mentality, but throughout its history, there have been features that were themselves seen as characteristic of intentionality. If some form of Brentano's thesis turns out to be true, then it will likewise turn out that these features are the defining features of the mental. So what are those features? The most traditional one is what we can call Nonexistence.² Roughly speaking, Nonexistence is just the slightly puzzling fact that we are able to think about and represent things that do not exist. Alternatively, we might say that many representations have nonexistent intentional objects. Brentano had a version of this feature in mind when he claimed that the defining feature of intentionality was what he called "intentional inexistence" [Brentano, 1973, Crane, 2014c]. His idea was that, in many cases, the intentional object of a mental state does not exist in the world, but rather exists *immanently*, and only immanently, to the mind. That is to say, many intentional objects are *mere* intentional objects, and do not have any existence outside of the intentional acts or states to which they are immanent. Thus, on Brentano's view, the ability to think of merely intentional objects—which do not exist, at least in the same sense that ordinary objects exist—is a central feature of intentionality, and in turn is potentially a defining feature of the mental.

The features of intentionality are often not just considered features: often they come to be seen as puzzles, and many of these puzzles have a distinguished philosophical lineage. For instance, we often hear about the puzzle or problem of Nonexistence. However, it is not in itself a puzzle that we are able to think about the non-existent. This is just a fact, acknowledged by basically everyone, and one that no ordinary person would even think twice about. However, the fact becomes a puzzle when we try to work out a theory of how we manage to do it, i.e. when we try to provide a general theory of intentionality. The hope is to be able to provide a theory of intentionality that does not leave us with a puzzling philosophical residue. Many theories of intentionality—for example, the intentional object view to which I have several times alluded—make the claim that we think about objects that do not exist, and for some, this leads to residual questions. What are non-existent

²See Caston [2001], Crane [2012].

objects? What are merely intentional objects? How do they differ from ordinary objects? Do we need to give an account of their metaphysics? Of course, there are theories that address just these residual questions, but the hope, and particularly *my* hope, is to be able to give a theory that makes thought about the nonexistent exactly as unremarkable as it is to the folk. As will become clear, I think views of intentionality that posit intentional objects often times fail to explain exactly what needs to be explained, and this is still often the case after providing a view of their metaphysics.

Nonexistence, like, as we will see, the other features of intentionality, can be detected by noting the failure of certain class of inferences. We can start at the most basic. Sentences involving verbs such as “think (about)” are not existence-entailing within their complements. For instance:

(1) John is thinking about a unicorn

does not entail:

(2) A unicorn exists.

The same is true for intentional verbs that take clausal as opposed to phrasal complements. Consider “believes”:

(3) John believes that unicorns have horns.

(4) Unicorns exist.

Clearly, (3) does not entail (4). Similar things can be said for a large class of verbs that express intentional notions, some of which have already been mentioned: “search”, “need”, “want”, “hallucinate”, “sense”, “experience”, “desire”, “suspect”. None of these notions are existence-entailing, which provides some evidence that Nonexistence is a good test for intentionality in addition to the long history; it is present in most states that are obviously representational.

Another feature of intentionality that has a distinguished philosophical lineage is what

we can call Nonspecificity. It will be easiest to illustrate Nonspecificity with an example (due to Victor Caston) [Caston, 2001]. Suppose that I greet someone. It is not possible to greet someone who is not of a particular height. Every person is of a particular height. But it is certainly possible for me to think of someone without thinking of a person of any particular height. In this sense, my thinking can be Nonspecific. The same holds for a host of other attitudes. I can perfectly well search for a bear, without searching for a bear whose fur is a particular shade (of brown, or any other color), even though every bear has fur of a particular shade. Even more generally, I might think of a bear, but not think of any particular bear, while of course, if I get attacked by a bear, it must have been a particular bear that did the attacking. In other words, just as a writer may not fix how many total cells are in the body of one of his fictional characters, I may not have fully specific thoughts. Moreover, for beings like us, it may even be *impossible* to have fully specific thoughts about things like ordinary objects; the world may well be too complex to represent specifically in thought. How could I, for instance, have a thought that was as detailed as every contour of a face? Such complexity does not seem capturable in thought; our representations of the world are lossy.

Similarly to the case of Nonexistence, Nonspecificity can be detected inferentially, by seeing what is entailed by instances of representation. Importantly, such inferential criteria are not yet anything specific to language; at this point we are just considering the nature of certain intentional states, and seeing what they entail, and what they fail to entail. We then express these entailments, or lack thereof, in language. Consider the following sentence:

(5) I am searching for a dog.

Clearly, I need not be searching for a particular dog, so (6) does not follow:

(6) I am searching for a particular dog.

It may well be the case that I am searching for a Vizsla, but not a particular one, and yet (5) can still aptly describe my search. Thus we have our first inferential test for Nonspecificity. But there are also slight variations on this inference that fail to hold for intentional locutions.

Consider the following inference:

(7) I imagined a man.

Every man has a specific number of hairs on his head.

Therefore, I imagined a man with a specific number of hairs on his head.

Obviously, this inference is no good. Why? Because “imagine” is an intentional verb, and what I imagine will typically not be one of the many concrete, flesh-and-blood men in the world, each of whom has a particular number of hairs on his head. Instead, when I imagine a man, I might imagine him incompletely, conjuring only some properties and not on others. It is obvious that we don’t imagine things down to every detail. Our imaginings can be incomplete, and allow for further determination, and it is this lack of full specificity that is unique to, and characteristic of, intentional states.

But not only can intentional states fail to be fully specific, but they can also be more and less specific. I can represent things in greater or lesser amounts of detail, and different states may represent in greater and lesser degrees of specificity. For instance, I might think of a dog, without thinking any particular one, and I certainly need not think of a dog with any particular color of fur, or of any particular size. In this sense, representation can occur at certain levels of generality or specificity (or perhaps “abstractness”), and we can think about this generality or nonspecificity in terms of representing determinables without their determinates. I can represent a dog, qua determinable, without representing any of its determinates, or I can represent it with all of its determinates, which would yield (at least something close to) a fully specific thought.

What this points us to is that with respect to specificity, some kind of hierarchy, or algebra, is needed to adequately account for the structure of intentionality. This hierarchy may be a hierarchy of properties, as developed in property theory, or perhaps something even higher-order: a hierarchy of properties of properties.³ Intentionality allows us to represent parts of the hierarchy without representing others. In the case of ordinary objects, it may

³For different developments of property theory, see Chierchia and Turner [1988], Bealer and Mönlich [2003]. For an account of the hierarchy of higher-order properties, or generalized quantifiers, see Peters and Westerstahl [2008].

be the case that possessing a certain property entails the possession of many others, this is rarely the case when we are considering representation. Representation lifts the strictures of metaphysical necessity: just because everything is some particular size does not mean that we must represent them as such, nor does it mean that just because a grizzly bear is of the species *Ursus Arctos* that I must represent it as one. The structure of intentionality need not fully mirror metaphysical structure.

But how exactly to interpret this hierarchy is still an open question. We might interpret the hierarchy as a collection of properties that are the objects of our intentional states, or we might interpret it as a conceptual hierarchy in which the elements of the hierarchy are various ways that we can represent the world. This latter option points us toward the idea that intentionality generally, and many intentional states more specifically, are not representations of properties, but rather *guide us toward the world* in certain ways with varying degrees of generality. On such a view, intentionality is direction toward the world, and the direction is provided by the content of certain representations. But importantly, the direction and specificity of our representations is not given by an intentional object. Rather, we use representations to point ourselves toward aspects of the world. Representations are kind of like internal maps.

The last feature of intentionality, which is perhaps its most well-known, is related to (Non)specificity, in that it concerns our ability to focus on particular aspects of objects and not others. In allowing us to focus on particular aspects of objects, intentional states are “finer grained” than objects themselves. Since the linguistic turn, this has often been captured by saying that we can think of—or search for, want, fear, or hope for—things under some descriptions but not others. That is to say, even if two descriptions or names pick out the same object, we need not have the same attitudes toward the object characterized in the two different ways. The attitudes may only hold toward the object characterized in a certain way—the object *qua* described in a certain way, or *qua* having a certain name.

This last feature is often associated with the property of contexts within sentences called “referential opacity”. Roughly speaking, a context within a sentence is referentially opaque, or just “opaque”, if substitution of coextensive expressions within that context can change the truth value of the sentence. Alternatively, we can give an inferential characterization of

opacity, by noting the failure of inferences such as that from (8) to (9), even if the shortest spy is Ortcutt:

(8) John is looking for Ortcutt.

(9) John is looking for the shortest spy.

John may be looking for Ortcutt without looking for the shortest spy, so while both “Ortcutt” and “the shortest spy” may have the same extension, substitution of one for the other need not preserve truth-value. On the assumption that “Ortcutt” and “the shortest spy” are constituents of (8) and (9) respectively, and assuming that the truth-value of a sentence is a function of the extensions of its constituents together with its syntax, this failure to preserve truth-value shows that the truth of the above sentences is sensitive to something over and above the ordinary extensions of the expressions. This sensitivity to more than mere extension is what is meant by calling a contexts within a sentence “referentially opaque”.

These three features form the core of a cluster of features that are characteristic of intentionality. But the status of these features at the center of the cluster of features raises two important points. First, these are not the only features that arise in discussions of intentionality. Historically there have been others as well. For instance, it is often seen as characteristic of intentional states that they come in two forms: a *de dicto* form and a *de re* form. Some authors have characterized this distinction as one between nonspecific and specific or relational and notional forms of intentionality, which correspond to *de re* and *de dicto* readings of intentional reports. These two readings are often associated with two scopal readings of intentional reports, and the non-equivalence of the readings on which a quantifier is interpreted inside and outside of the verb phrase. However, this is not a universally recognized distinction, and it is closely associated with the semantics of intentional reports, so for the moment it will suffice to note it and delay further discussion until our discussion of the linguistic hallmarks of intentionality in the next section.

However, these properties do not provide a hard and fast criteria for identifying intentional states. Instead, these features are merely *symptomatic* of intentionality. For it may well be the case that there are some intentional phenomena that lack all of these features,

and there may also be non-intentional phenomena that have some of these features. Once a theory of intentionality is developed, the theory may make some of these features criterial for intentionality, but that conclusion is a consequence of substantive theorizing, rather than a starting point for the investigation. I myself am inclined toward treating these features, with appropriate qualifications, as definitive of intentionality, and so definitive of the mental, but very little in what follows will turn on this view. Rather, as we proceed, I will treat these features as indicative of intentionality, and so I will treat an account of them as an important aspect of a theory of intentionality. However, I will allow that there may be intentional phenomena that do not exhibit these features, and so may admit of a different sort of explanation. However, I believe that these features can be used to help construct a novel and nuanced account of our representational lives.

0.3 Intensionality as the Mark of Intentionality

Thus far we have discussed three properties that are traditionally associated with intentional states: Nonexistence, Nonspecificity, and Opacity. Importantly, we've talked about these three features as features of the metaphysics of intentionality: they are data concerning intentionality that our theory of intentionality needs to explain. But we also gave these features inferential characterizations, and these inferential characterizations show us that intentionality manifests itself in language in distinctive ways. It turns out that the three inferential features of intentionality are the definitive features of the class of Intensional Transitive Verbs (ITVs), and perhaps the class of intensional verbs more generally. In one sense, this is a surprising confluence. When linguists discuss the features of intensionality, they rarely acknowledge that these features have been discussed for a millenia-and-a-half as the features characteristic of intentionality. Rather, as a linguistic phenomenon, intensionality, characterized by these three inferential patterns, has mostly been divorced from its origins in the theory of intentionality, and few authors have explicitly connected the project of giving a semantics for verbs that exhibit these features with the philosophical project of providing a theory of intentionality. So even though intensional verbs are relatively well-studied in linguistics, and linguists have developed a great deal of semantic machinery to

deal with their features, I suspect that reconnecting the study of intensionality to its history will come as a surprise to many working semanticists. However, in another sense, this confluence is exactly what we should have expected. In our theorizing about a phenomenon, we found that it had characteristic features, and we captured these features by noting the failures of certain kinds of inferences. Since these inferences fail for a certain class of reports, and these reports are made using certain verbs, it makes sense that these inferences can be used to delimit a distinctive class of verbs: those connected to representation.

A verb is considered transitive when it takes a noun phrase in its object position, occurring in sentences of the form **NP V NP'**.⁴ A transitive verb **V** is considered intensional when sentences of the above form exhibit some combination of the following three properties.

Nonexistence: **NP V NP'** has a reading which fails to entail **NP' exists**, where **NP'** is upward-entailing.⁵

Nonspecificity: **NP V NP'** has a reading that fails to entail **NP Vs a particular NP'**.

Opacity: **NP V NP'** has a reading that fails to entail **NP V NP***, where **NP'**, and **NP*** are extensionally equivalent.⁶

These three features are laid out as criteria for a *transitive* verb to qualify as intensional. However, it has been largely overlooked that they also serve to single out to intensional verbs that take clausal complements. Thus, arguably, these three features are characteristic of intensionality generally; showing that these features are present is intensional clausal verbs is merely a matter of changing the syntactic characterization.

⁴However, many verbs that are technically intransitive are treated as transitive when they occur in constructions of the form **NP V P NP'**, where **P** is a preposition. This is typically done when the combination of intransitive verb and preposition have a transitive verb as a near-synonym, such as in the case of “seeks” and “is looking for”. Many of these verb + preposition combinations behave identically to transitive verbs, and so unless otherwise noted, I’ll treat them as intensional transitives also.

⁵By “upward entailing” I mean to include positively quantified NPs like: “a dog”, “the men who robbed him”, “four gorgons”, “infinitely many numbers”, as well as proper names, and bare plural NPs. I mean to exclude negative NPs like “no dogs”, “no one”, *etc.* By “empty” I mean that nothing in the world answers to the NP.

⁶It’s important to note that there are many ITVs that do not have all of these features: typically, the presence of even one of these properties is sufficient to classify a verb as intensional. However, verbs like “seek” are paradigmatically intensional in that they exhibit all three of the properties. As will become clear, I am arguing that not only are our semantic verbs intensional, they are like “seek” in being paradigmatically intensional.

The thing to note is that these three inferential features—which play the role of data for which a semantic theory must account—are just the features of intentionality that we laid out above, in a linguistic guise. Thus, we find ourselves in a situation that should feel familiar from other areas of philosophy. On the one hand we find ourselves investigating a phenomenon of philosophical interest—intentionality—and isolating its core properties. But we also find that these properties are manifest in the various semantic features of verbs expressing the notions that we are interested in. Nonexistence, Nonspecificity, and Opacity are the three distinguishing metaphysical features of intentionality, and they are also the three core properties of intensionality, a purely linguistic phenomenon. Thus, the inferential criteria serve to couple the metaphysical and semantic levels of investigation. Insofar as we provide inferential criteria for the metaphysical features of intentionality, we tie our metaphysical investigation to a semantic investigation which attempts to provide truth-conditions for the reports in question that account for the inferential patterns.

This leads us to the view that intensionality, a linguistic phenomenon, is the mark of intentionality, the phenomenon that is plausibly characteristic of the mental. The view that intensionality is the hallmark of intentionality is an old view, and goes back to Chisholm, and was championed by Anscombe [1965]. Since the question of whether intensionality can serve as a criterion for intentionality is highly controversial, I don't wish to make any strong claims about the connection between them here.⁷ I won't, for instance, try to show that the intensionality of a report is a necessary and sufficient condition for a state to be intentional. However, I do think there is a strong connection between the two phenomena. I hold that the inferential characterization of the metaphysical features of intentionality is the best one we have, and it just so happens to overlap perfectly with the linguistic characterization of intensionality. In a case such as this, there seems to be very little difference between investigating the lexical semantics of intensional verbs and theorizing about intentionality itself; lexical semantics makes use of metaphysical insight, and metaphysics can likewise draw on the insights of lexical semantics. This might sound like a strong commitment to a linguistic approach to metaphysics, but it is not: I am denying that, in certain cases, there is an important distinction to be made between metaphysical investigation and lexical

⁷See Crane [1995] for a discussion of the relationship between intentionality and intensionality.

semantic investigation.

0.4 Semantics and Metaphysics: Anscombe's Project

If intensionality is the mark of intentionality, as Anscombe [1965] and Chisholm [1956] thought, then it points the way to a new project, one which Anscombe herself saw clearly and began to pursue. The project is as follows. If the three features characteristic of intentionality are manifested as the core features of intensional language, any theory of the nature of linguistic intensionality will likewise serve as an approach to the metaphysics of intentionality. Anscombe pursued this project in a very specific way: she wanted to use the semantics of perceptual locutions such as “see” and “perceive” to solve the problems of hallucination and illusion, which are themselves puzzles of intentionality. Her approach was to investigate what kind of thing could serve as the direct object of perceptual reports when no concrete existing object was a candidate for the job. Consider the following:

(10) John sees a dancing dragon.

On Anscombe's view, “sees” is an intensional transitive verb, and seeing is an intentional state. She held that if we can provide a semantics for sentences like (10), and more specifically, if we can find semantic values for NPs such as “a dancing dragon” that occur in the object positions of intensional transitive verbs, then we will have found the objects of hallucination, and taken a huge step toward solving the problem of hallucination. She extends this kind of reasoning to other perceptual verbs for which the problems of hallucination and illusion can be formulated. Thus, on Anscombe's view, providing the correct semantics for perceptual reports stands to illuminate the metaphysics of hallucinatory and illusory perceptual states, and perhaps perceptual states more generally.

Anscombe's insight forms the foundation for this project, but her project is limited in scope, and faces several difficulties that this thesis hopes to overcome. First, Anscombe's proposal is specific to perceptual verbs. But it is not merely perceptual verbs that exhibit the features of intentionality; rather, many, and perhaps all intentional verbs exhibit these features, and just as with perceptual verbs, investigation of their semantic features can in-

form our theories of their nature. However, I don't wish to hold that such a strategy can only be deployed in one direction, using semantics to uncover metaphysics. As I stressed above, in this case, among others, I do not think that there is a substantive difference between doing lexical semantic theorizing and giving an account of the nature of the thing expressed. Lexical semanticists deploy metaphysical intuitions just as readily as metaphysicians deploy semantic intuitions, and different intentional verbs may exhibit different combinations of the inferential features, and will exhibit different profiles of intensional semantic behavior. So while I hold that Anscombe's basic strategy is correct, my proposal is more general. I hold that the intensional semantic profiles of all intentional verbs will need to be reflected in the semantics we provide for reports of the phenomenon, and these profiles can be deployed in our understanding of the natures of the states that the verbs express.

Second, Anscombe's approach was originally worked out for perceptual verbs; her proposal concerned what she called a "grammatical feature" of verbs of sensation and perception: namely, their intensionality. However, it is hotly contested whether her basic semantic claims concerning perceptual verbs are true. Many theorists hold that perceptual verbs are fully extensional, and there has been a debate raging since well before Anscombe wrote over whether non-factive, intensional readings are available for perceptual reports. I myself think that perceptual verbs have intensional readings, and I have empirical work in progress that attempts to establish this claim. However, it is not essential to my project that any particular perceptual verb exhibit the features of intensionality. Even if some perceptual reports turn out to be fully extensional, some related verbs will be required to report the phenomena of hallucinatory and illusory perceptual experiences, and these verbs—for instance, "hallucinate"—verifiably exhibit the features of intensionality, and Anscombe's proposal can be generalized to them. Thus, even if we lose the ability to apply Anscombe's strategy to all perceptual reports, we can still deploy her insight to develop an account of hallucinatory perception. However, while the intensionality of perceptual verbs provides a rich area for further research, nothing in the remainder of this thesis will serve as an argument for the claim that perceptual verbs have intensional readings.

The strategy for investigating intentionality that emerges is a generalized Anscombian strategy. Since there is broad convergence between the properties of intentionality and

the features of intensionality, and the two categories are held together by the inferential criteria, we can expect there to be a fine-grained correspondence between our theories of the metaphysics of intentionality and our theories of the semantics of intensionality. Thus, not only do the semantic and metaphysical projects dovetail, but they stand to gain a lot from one another. Semantic accounts of intensionality can be seen as views on the nature of intentionality and intentional objects, and similarly, theories of the metaphysics of intentionality can be translated into semantic terms. This is in keeping with the tradition in philosophy on which giving a semantics for a particular term or collection of terms is integral to understanding the nature of the object of investigation, and plays an important role in argumentation for and against the view.

The style of investigation which intertwines semantic and metaphysics in this way is common in many fields: metaethics, epistemology, philosophy of mind, and metaphysics, just to name a few. Since this kind of approach is so common, it is surprising that modern analytic philosophers working on intentionality rarely seem to investigate the systematic connection between theories of intentionality and the semantics of the locutions used to report them. The Anscombian program has gained little traction. Perhaps this is due to the fact that the semantics for intensional transitive verbs in particular is a difficult semantic project. But similar projects have been undertaken on the nature of belief, knowledge, and the good, and the semantics for reports of these phenomena are equally complicated. In any case, an investigation of the semantic features of intentional reports stands to shed light on phenomena as diverse as perception, sensation, hallucination, illusion, imagination, belief (in), reference, satisfaction, and truth (of), among many others.

The next section will be devoted to categorizing views of intentionality in light of the generalized Anscombian strategy. This strategy will treat different approaches to the semantics of intensionality as potential views on the metaphysics of intentionality. What is most fruitful about this approach is that there is a wealth of semantic resources that have yet to be marshalled to address the problems of intentionality. Thus we can bring these semantic resources to bear on many verbs—as well as other parts of speech, such as nouns, adjectives, and prepositional phrases—that exhibit intentional notions, and these resources can provide us with the foundations of different views on the nature of representation,

intentionality, perception, sensation, reference, and myriad other phenomena expressed in intentional language.

0.5 Metaphysical Views and Semantic Views

Views of intentionality break down along two axes. First, views of intentionality can either be relational, non-relational, or mixed, where mixed views involve at least one relational and non-relational component. I count views of intentionality as relational if they take intentional states to be, or essentially involve, relations to intentional objects, whatever the nature of those objects turns out to be.⁸ Relational views of intentionality are paired with relational views of the semantics of intentional verbs, on which the intentional verbs that express the intentional state in question take arguments in both their subject and object positions. Roughly speaking, when an intentional state is ascribed with an intensional verb, the intentional object will be an, or the, internal, direct object argument of the verb, while the subject will be the external argument. Relational views of intentionality, as I will construe them, hold that the semantic value of the intensional complement of an intentional verb is the intentional object of the state that the verb expresses. Given a particular intentional verb that exhibits all three features of intensionality, the relational view holds that accounting for the peculiar features of intentionality is a matter of finding the correct type of object to assign as the semantic value of the complement.

Nonrelational views of intentionality are harder to characterize, and have often been overlooked. Roughly speaking, non-relational views hold that intentionality results from the non-relational, intrinsic features of a representation or a representational state. For instance, consider (11):

(11) Huey is searching for a superhero.

The nonrelational view of intentionality holds that Huey does not need to stand in a relation

⁸I am using “object” here in the broadest possible sense, to include entities, properties, relations, quantifiers, propositions: anything that we can dream up that can play the role of an argument in the direct-object position of a verb. However, I am restricting my usage of “object” to a metaphysical sense of “object”, and so I am not using “object” in the grammatical sense, although I think the two senses are related. By “direct object”, I think that most people mean something like the *theme* of the verb.

to anything at all in order for him to be searching for a superhero. He might, for instance, need to be in a certain state, and he also might, for instance, need to have a certain representation that guides his search, but importantly, there is no *thing* to which his search relates him. Searching is not a relational state, and although it may be directed, searches need not have direct objects. As we will see below, there are several ways of spelling out the nonrelational approach more precisely, and I will defend one such approach at length in chapter three. However, since much more ink has been devoted to spelling out relational views of intentionality, much of the rest of this introduction will focus on providing an account of the various nonrelational approaches.

There are also mixed views of intentionality that involve both a relational and a non-relational component. Such views hold that the intentionality of a representation involves a relational and an intrinsic feature of that representation. One notable kind of two-factor view is one which countenances both an intentional object and intentional content, where intentional content is meant as the “way in which the object is represented”. There are several views with this structure, the earliest of which is Brentano’s, as well as Anscombe’s. Tim Crane, who takes inspiration from both Brentano and Anscombe, also holds a two-factor view.

The second axis along which views of intentionality break down concerns the nature of the content assigned by the content-specifying phrases or clauses within the complements of intentional verbs. It has been common in the last 50 years of theorizing about language and the mind to treat all intentional states as propositional attitudes, or as states with propositional content. This is certainly true in the literature in the philosophy of mind, where the paradigm cases of mental states are beliefs and desires, both of which are construed as relations to propositions or representations with propositional contents. Non-propositional intentional states, such as “thought about” have been, for the most part, either overlooked or assimilated to propositional attitudes. This approach to intentionality takes a cue from the trend in philosophy of language to privilege propositional contents over other kinds of contents. For instance, both Quine and Lewis had worries about the assignment of contents to subsentential linguistic expressions. One reason for Quine’s skepticism was his view that the syntax of natural language is underdetermined, and that there is no fact of the matter

concerning the subsentential structure of English sentences. But as if that view weren't enough, Quine also held that reference was inscrutable, and that meaning was fundamentally indeterminate; he was a skeptic about the possibility of semantics in general. Lewis was not a skeptic about the possibility of semantics generally, but he held that assignment of reference to subsentential expressions was only holistically constrained by the convention of truthfulness and trust that links a community to a language, and that within these constraints, various assignments of subsentential contents were possible. Thus, Quine and Lewis both held that the only things that we have stable intuitive judgments about are the truth-conditions of sentences, and these judgments can be maintained in the face of various permutations of the references of the subsentential constituents. Thus, they view assignment of extension to subsentential constituents as merely instrumental to the computation of truth-conditions: semantic values are necessary for compositional construction of sentence meanings, but different assignments can fulfill this task equally well.

This approach has some immense benefits, the first of which is that it ties our theory of intentionality to theories that systematically assign propositional contents to sentences, i.e. formal semantic theories that assign truth-conditions to sentences. Insofar as having truth-conditional content is one way of being about the world, this seems close to what is wanted from a theory of intentionality. Second, at least as Quine and Lewis pushed for the view, it stands to solve the problem of Nonexistence by reducing it to the problem of falsehood. When we think of something that doesn't exist, we are having a propositional thought, and provided our account of propositions is not object-dependent, such thought poses no particular puzzle. It also points toward solutions to the other two problems of intentionality. Further, the propositional view of intentionality is paralleled by a view on the semantics of intensional transitive verbs also originally due to Quine: the view that all intensional transitive verbs are actually covertly propositional, or should be analyzed as propositional attitudes. The view was originally proposed by Quine [1956], and championed by the early Montague [1974b]. The view has been defended in a modern form by Larson et al. [1997].

However, viewing all intentionality as truth-conditional or accuracy-conditional is a massively simplifying assumption that seems to ignore very basic intuitions concerning how

our minds come to think about the world. For instance, when I think of a dog, or of the number three, is the content of my thought propositional? Am I thinking that such things exist, or have such and such a property? It seems not. It seems that I can think of objects simpliciter; not every thought *of* something is a thought *that* something is the case. Further, on the semantic side, there are good reasons to think that not all intensional transitive verbs can be reduced to instances of clausal complementation, despite the arguments given by Larson [2002], Larson et al. [1997].

A less restrictive, and more psychologically and phenomenologically realistic view of intentionality would be more ecumenical: it would allow that we can think about many different sorts of things, and we can do so without always having an attitude toward a proposition. There is no unique mechanism or type of content associated with intentionality, but rather, different kinds of intentional states may have different kinds of contents. On my view, the modes of intentionality are at least as multifarious as our common-sense intentional vocabulary, and we should take such common-sense intentional vocabulary at face value, rather than shoehorning it into a predetermined theoretical box. Similarly, I think we should approach the semantics of intensional constructions at face value, rather than by presuming that they can be paraphrased away into simpler and better-understood idioms. This approach promises a more realistic account of intentionality that is informed by the semantics of intentional locutions of various kinds. We should take the semantics of intentional reports seriously, and let them play a guiding role with respect to the nature of intentional objects and contents.

There seem to be three basic kinds of semantic values that theories of ITVs assign to their phrasal complements, and these three views yield three different metaphysical approaches to intentionality. Consider a paradigmatic intentional state: thinking about something. Let's suppose, for example, that Huey is thinking about a superhero, and let's presume that superheroes don't exist. We can either hold that the semantic value of this complement is an entity, a property, or a property of properties.⁹ The distinction between relational

⁹Often, in formal semantics, quantified NPs receive generalized quantifiers as their semantic values. Generalized quantifiers are properties of properties, but in this introduction I don't want to presume knowledge of any technical semantic or logical machinery, so I'll use the slightly less technical notion of a property of properties. The general idea is that quantified NPs will have distinctively quantificational semantic values, and such quantificational semantic values are one approach to forming a theory of intentionality.

and non-relational views will then determine whether the entity, property, or quantifier assigned to the complement is the intentional object of the state ascribed, or merely serves to characterize the state itself. In semantic terms, relational views will make the entity, property, or quantifier the direct-object argument of the verb, whereas on nonrelational views, they will bear a different relation to the verb: they will not be the verb's direct object.

Let's think about how the various relational views might account for this state. First, consider an entity-based view. On such a view, Huey might be thinking of a particular superhero, say, Superman. In this case, we might hold that Superman is a particular entity who just happens to not exist, and it is in virtue of being related to this nonexistent object that John is thinking of it. This is an example of a Meinongian view, that commits us to there being nonexistent objects. Superman is such an object. But John may not be thinking of an particular superhero; he may be thinking of a superhero, but not a particular one, or of a not-fully-specified superhero. On the entity view, we will then have to approach John's case as an instance of reasoning with arbitrary, or nonspecific objects. There are also some two-factor views that fall into the entity category. On most two-factor views, what John is thinking of is an object, but two-factor views allow that John can think of the object, or represent the object, in a particular way. Such views make a distinction between the intentional object of a state—the object that is thought of—and the intentional content: how it is represented. I take Brentano's conception of intentionality to be one of these views [Brentano, 1973], along with Anscombe's [Anscombe, 1965] and (possibly) [Crane, 2012].

The second approach to our example above is to hold that John is thinking of a property. When John is thinking of a superhero, he bears a relation to a property, which serves as the semantic value of the indefinite description in the object position of the intentional report, and on a relational view, the semantic value of the intensional complement of the ascription is the intentional object of the state the verb expresses. We can see this property-based view of intentionality as having one of its sources in Quine. While Quine was no fan of psychological idioms, his theory of linguistic intentionality can perfectly well be adapted into a general theory of intentionality. Quine's theory of linguistic intentionality is one given in terms of predicates and quantification. For instance, on Quine's view,

names are general terms—i.e. predicates—that denote sets of objects, although if we do a bit of violence to Quine’s own view, we can say that predicates denote properties. On this modified Quinean view, sentences like “Pegasus flies” are about Pegasus, roughly, because they are about something that satisfies a particular property: the property of being Pegasus. Accordingly, Quine paraphrases the sentence “Pegasus flies” using a predicate and a quantifier, as $\exists x(Pegasizes(x) \wedge Flies(x))$. Such views have been refined and defended by Fara [2015], who holds that both definite and indefinite descriptions are also predicates [Fara, 2001].

The view that we can develop from the Quinean approach to empty names is one on which the complements of intensional transitive verbs, such as “thinking of” in our example, have predicative type. This view has semantic precedent, although it has not been proposed as a view of *intentionality*. For instance, Zimmermann [1993, 2001, 2006] holds that all intensional transitive verbs have complements of predicative type, and that intensional transitives express relations to properties. This dovetails extremely well with Fara’s view of names and descriptions as predicates. On Fara’s view, names are predicates, and so contain free variables; in keeping with the ordinary treatment of predicates, the semantic value of a name is a property. When the name is used in a sentence, the free variable is implicitly existentially bound, and quantifier domain restriction is applied to yield or at least approximate uniqueness. Such a view could be applied uniformly to the complements of intensional transitive verbs, letting the fact that the predicate is interpreted inside of the scope of the verb suppress existential commitment.

The last relational view is one on which intentionality is a relation to a set of properties, or a quantifier. This was originally Montague’s proposal. Montague proposed that intensional transitive verbs like “seeks” were relations between a subject and an intensional quantifier: a function from worlds to sets of properties. This, Montague held, is a perfectly general account of the semantics of intensional transitive verbs, and moreover, one that accords with his approach to the semantics of noun phrases more generally. In assigning intensional quantifiers to phrases in the object-positions of ITVs, Montague assigns objects that seem to capture nonspecific searches, and searches for certain numbers of objects. Further, bearing a relation to an intensional quantifier doesn’t seem to

be existence-entailing, and intensional quantifiers are individuated finely enough to capture Opacity. Many subsequent theorists have followed Montague's lead, in particular Richard [2013], Moltmann [1997, 2008], and in certain ways, Forbes [2006].

Nonrelational views of intentionality have historically received little attention, largely due to the dominance of the view that our thoughts and representations have intentional objects. Part of the goal of this thesis is to remediate this oversight by an developing an alternative, non-objectual, nonrelational theory of intentionality. There aren't many well-known examples of non-relational theories of representation, but one prominent examples comes from the philosophy of perception; adverbialism about perception is the most well-known nonrelational theory of an intentional phenomenon. But adverbialism need not be seen as strictly a view on the nature of perception; rather, like we did above with the Anscombian strategy, we can consider what adverbialism looks like when considered as a view of intentionality generally. On the metaphysical level, Adverbialism about intentionality is the view that to think about, perceive, or sense something is to think, perceive, or sense in a particular *way*. If we take the label Adverbialism at face value, then specifying the particular way in which we think will be done by an adverb; in the case of Huey and his superhero, an Adverbialist account of Huey's intentional state would construe "a superhero" as an adverb that specifies the way that Huey is thinking.

However, the original motivations for Adverbialism, as a theory of perception, came from the desire to solve the problems of hallucination and illusion without recourse to sense-data. Treating the content-specifying phrases or clauses of perceptual reports as adverbs was one way to accomplish this goal; it was one version of a non-relational theory of such, but there are others. In particular, one need not assimilate the content-specifying phrases or clauses of perceptual and intentional reports to adverbs in order to give a nonrelational theory of intentionality. Rather, there are many linguistic resources that we can deploy to develop such a theory; treating the intensional complements of intentional verbs as adverbs was only a first attempt that was largely dismissed because it was lacking in linguistic and theoretical sophistication and plausibility.

There is, however, one theory of representation that is thoroughly nonrelational: Goodman's. Goodman holds that representing something is not a matter of bearing a relation

to a particular object, existent or otherwise, but is rather to be a representation of a particular *sort*. For example, for a picture to be a picture of Pickwick, from the Dickens novel, is not to represent some fictional object, but is rather for the very painting itself to have certain features, intrinsically. Goodman makes this idea explicit with a proposal for a paraphrase: “Picture of Pickwick” is best understood as a “Pickwick-picture”. The proposal can be generalized to a proposal for all intentional reports; chapter three develops just such a proposal in a linguistically rigorous way. But the main point to note here is that Goodman’s proposal presents one way of satisfying the desiderata that originally motivated the development of Adverbialism without invoking adverbs. Goodman does, however, use classifiers: the hyphenated paraphrase classifies the picture as one of a particular sort using an adjective, which is similar in spirit to how an adverb classifies, or helps to classify, an event. Thus Goodman’s proposal is similar to the adverbial view, but points to a broader approach to nonrelational intentionality: both adjectives and adverbs can serve as the basis for non-relational theories of intentionality.

Goodman’s view is not the only non-relational view on the market. There are several nonrelational proposals concerning the semantics of intensional transitive verbs that can themselves form the bases of nonrelational theories of intentionality, but are not explicitly adverbial. For instance, Graeme Forbes [2006] and Friederike Moltmann [2013] offer proposals concerning the semantics of intensional transitive verbs that do not treat such verbs as having a direct-object argument. Forbes, for instance, treats intensional transitive verbs within an event-semantic framework, and posits a special thematic role for the notional reading of an intensional transitive. This new thematic role takes as an argument a quantificational phrase, and the quantificational phrase, on his view, “characterizes” the event in question, for instance, a search, or a desire, or the relevant mental state, where “characterization” is spelled out in terms of satisfaction conditions.

Forbes and Moltmann’s semantics for intensional transitives are very different, but both treat the intensional phrasal complements of ITVs as supplying quantificational material to logical form. The quantificational material contributed by the NPs helps to specify *satisfaction conditions* for the intensional verb. However, neither view treats intensional NPs as direct-object arguments of the verb; the semantics that Moltmann provides treats

	Entity	Property	Quantifier
Relational	Meinong, Prior, Parsons, Crane, Brentano, Anscombe	Quine, Tarski, Fara, Zimmermann, Strawson	Montague, Richard, Johnston
Nonrelational	Goodman, Chisholm	Goodman, Chisholm, Dayal	Forbes, Moltmann
Mixed	Brentano, Crane, Anscombe, Johnston	Anscombe*	Johnston*

Table 1: Views on Intentionality, Categorized by Content-Type and Relationality

intensional positions syncategorematically, while Forbes treats them non-thematically. Syncategorematic treatments are useful in accounts of intensionality because they allow us to specify how an expression contributes to the truth-conditions of sentences containing it without supplying it with a semantic value. This strategy should be familiar from first-order logic: when we give the semantics for expressions of first order logic like the quantifiers, we do not assign them semantic values, but rather state how they contribute to truth-conditions in the metalanguage. The resulting categorization looks something like Table 1.

On the one hand, we can classify views of nonpropositional intentionality as relational or non-relational, and on the other hand, we can classify them in terms of the semantic value that their ascriptions assign to their object-positions.¹⁰ Or, in less semantic terms, we can classify them in terms of their relationality, and in terms of the nature of their contents. The problem cases for this classification are the two-factor views, represented by Brentano, Crane, and Anscombe. Since their views involve two factors, each of these factors may be of different sorts. For instance, on Anscombe’s view, every intentional state involves an intentional object, and intentional objects, for Anscombe, are something like objects under descriptions. So the intentional object, which serves as the semantic value of the intensional NP complements of intentional verbs, comprises an entity and a description, or an entity

¹⁰There is a problem with this way of categorizing the views. Moltmann’s semantics for ITVs is syncategorematic, which means that she does not assign a semantic value to the phrase in the object position of an ITV at all. But she does provide the entire construction with truth-conditions. However, in her syncategorematic specification, she uses a schematic letter whose instance is a quantifier, even though that quantifier is not, technically speaking, the semantic value of the NP. So I have classified her as assigning quantifier-type contents to the object-position NP, even though this only occurs as a part of the derivation of the final truth-conditions.

and some property of that entity. Thus, Anscombe’s two-factor view should actually take up two categories on each of our axes; it has a relational and a non-relational component, as well as involving both an object and a property. For simplicity, I have given mixed views their own category on the relationality axis, and then marked the two types of contents they involve with a second occurrence of the name under another content-heading.

0.6 Locating the Project

My project in this dissertation is to propose a new non-relational theory of linguistic, mental, and perceptual intentionality. The framing assumption is Anscombian: by addressing the semantics of intentional verbs, and providing them with a non-relational semantics, I develop a non-relational theory of the metaphysics of intentionality. The three chapters to follow each complete an important part of this project, and together they form the basis of a research project that explores the semantics of semantic, perceptual, and psychological verbs and uses these semantic insights to open new avenues of investigation of each of these areas. Below are summaries of the contributions that each of the three chapters make to the overall project, followed by a discussion of how they fit together into the foundation of an important research proposal.

0.6.1 Semantic Verbs are Intensional Transitives

In the first chapter, titled “Semantic Verbs are Intensional Transitives”, I argue that semantic verbs such as “refers (to)”, “applies (to)”, and “is true (of)” have all of the features of intensional transitive verbs, and discuss the consequences of this claim for semantic theory and the philosophy of language. One theoretically enriching consequence of this view is that it allows us to perspicuously express, and partially reconcile two opposing views on the nature and subject-matter of semantics: the Chomskian view, on which semantics is an internalistic enterprise concerning speakers’ psychologies, and the Lewisian view, on which semantics is a fully externalistic enterprise issuing in theorems about how the world must look for our natural language sentences to be true. Intensional Transitive Verbs have two readings: a *de dicto* reading and a *de re* reading; the *de dicto* reading of ITVs is plausibly

a nonrelational reading, and the intensional features peculiar to this reading make it suitable for expressing a Chomskian, internalist semantic program. On the other hand, the *de re* reading is fully relational, and make it suitable for expressing the kinds of word-world relations essential to the Lewisian conception of semantics. And since the *de dicto* and *de re* readings are plausibly related as two distinct scopal readings of the very same semantic postulates, we can see these two conceptions of semantics as related by two scopal readings of the very same semantic postulates.

The methods by which I argue for this claim are partly empirical and partly theoretical. I begin by discussing three empirical studies showing that “refers to”, as it is used in English to state speaker’s reference, has all three of the features of intensionality. I then argue that we have reason to treat “refers”, as it is used technically in semantic theorizing, as intensional as well. I argue for this on two grounds. First, on many views, semantic reference is ultimately determined by instances or patterns of speaker’s reference, and so will inherit the intensionality of speaker’s reference. Second, I argue that treating “refers”, as it is used technically in semantics, as an ITV has several important theoretical benefits, including that it points the way toward providing semantic values for empty names, and promises to provide semantic values for empty names that are as fine-grained as those of the NP complements of ITVs more generally. Since ITVs are often hyperintensional within their complements, this offers us the prospect of providing expressions with hyperintensional semantic values. And lastly, treating semantic verbs as ITVs allows us to satisfy a final Chomskian desiderata on a theory of reference: it allows for a notion of nonspecific reference.

0.6.2 Hallucination and the New Problem of Empty Names

In chapter two, titled “Hallucination and the New Problem of Empty Names”, I argue that the problem of hallucination and the problem of empty names are, at bottom, the same problem. I argue for this by reconstructing the problem of empty names in way that is novel, but implicit in much of the discussion on empty names. I then show how, once recast in this light, the two problems are structurally identical down to an extremely fine level of granularity, and also substantially overlap in terms of their content. If the problems are identical in the way I propose, then we should expect that their spaces of solutions

are also identical, and there is significant support for this conclusion. I characterize the space of possible responses to each of the problems, and then discuss the pairings between prospective solutions to each of the problems. However, there are some proposed solutions to the problem of hallucination that have been overlooked as potential solutions to the problem of empty names, and this realization opens new approaches to the problem of empty names, and to the nature of meaning more generally.

One notable option for the treatment of empty names that has been overlooked is what I call *semantic adverbialism*, the semantic counterpart of perceptual adverbialism. I propose a way of defending semantic adverbialism by invoking the idea defended in chapter one, on which semantic verbs are intensional transitives. One appealing view of the semantics for the intensional, *de dicto* reading of an ITV is to treat it as non-relational. On this view, noun phrases in the object-positions of ITVs serve as modifiers, helping to form complex predicates. Since adverbs are one kind of modifier, this allows us to formulate a theory of semantic meaning that is adverbial in spirit, if not in letter. The approach I propose is more general, and might be better termed “adjunctivism”. There are several specific semantic proposals that pursue this general idea: one due to Forbes [2006], another due to Moltmann [2008, 2013], and a third proposed in chapter three. Whichever nonrelational proposal we decide on, it will allow us to formulate a version of semantic adverbialism. I then discuss the possibility of generalizing this semantic approach to the problem, which would unify our approaches to the semantics of intensional verbs with our approaches to both the problem of empty names and the problem of hallucination.

0.6.3 Intensionality is Additional Phrasal Unity

In chapter three, titled “Intensionality is Additional Phrasal Unity”, I argue for a novel approach to the semantics of intensional contexts. At the heart of my proposal is the Quinean view that intensional contexts should, from the perspective of the semantics, be treated as units, with the material in them contributing to the formation of a single predicate. However, this proposal is subject to a number of objections, including the criticism that taken at face value, this would render intensional contexts, which seem to be fully productive, non-compositional. The paper begins by discussing the concept of the unity of the phrase,

and pointing to various ways that phrases can gain additional unity. It then proposes that the intensionality of intensional transitive verbs is best construed as a form of semantic incorporation; ITVs, on their intensional readings, meet all of the criteria for qualifying as incorporating the nominals in their object positions. Some of the criteria for qualifying as incorporated are even identical to the criteria for qualifying as intensional.

One form of incorporation that manifests itself in English is where an object-position nominal is moved to the front of the verb and compounded with it. Sometimes sometimes such compounding is marked with hyphenation, as in “apple-pick”, but other times the result is fully lexicalized, as with “babysit”. I propose to treat the semantics of intensional transitives on the model of semantics for these kinds of incorporated constructions. Following Dayal [2003, 2011], I treat the intensional NPs in the object-positions of ITVs as verbal modifiers, which combine with the verb to form a new, morphologically complex word. I show how such a proposal is compositional, and accounts for all three of the traditional features of intensionality: Nonexistence, Nonspecificity, and Opacity. I also gesture at how this proposal can help make sense of varying judgments concerning inference patterns within such contexts. I then extend the proposal to intensional verbs that take clausal complements, such as “believes” and “desires”, and also to intensional NPs, such as “picture of a house”. In the former case, I propose that the clausal complements function like phrasal compounds, in which a whole phrase serves as a unified modifier of the main verb. This proposal serves as one vindication of the Quinean approach to intensionality from within modern linguistics.

Chapter 1

Semantic Verbs are Intensional Transitives

1.1 Introduction

There is a strand of thought concerning the nature and subject matter of semantics on which semantics does not state relations between words and the world. On this view, semantics does not issue in truth-conditions, nor do its lexical postulates state relations between words and objects; instead, semantics is an internalistic enterprise that concerns the psychology of language users. This is the view proposed by Chomsky [1977, 1995, 2000], various versions of which are held by theorists working in the Chomskian tradition, including Pietroski [2003, 2005, 2006, 2008, forthcoming], Collins [2008, 2009, 2014], and Jackendoff [1983], among many others. On Chomsky's view, semantics attempts to explain how syntax interacts with our conceptual and intentional systems. It is only relative to an extremely detailed context, along with fine-grained aspects of speakers' intentions, interests, beliefs, and desires—which may turn out to be theoretically intractable—that we can ever say that a word picks out a particular object in the world. Further, Chomsky thinks that even so relativized, there are still often no objects that are suitable candidates to serve as the worldly referents of a word, but this does not in any way threaten to deprive words of semantic significance. I will not rehash Chomsky's arguments for these claims here, but his main point is clear: the

semantic features of words are not the result of their relation to any objects in the world, and the mind plays an important role in determining the semantic features of words in a way that renders assignment of objective reference either impossible or pointless.

But there is a competing line of thought according to which semantics does not concern speakers, their psychologies, or what makes them linguistically competent. Rather, semantics is a theory of the contents of natural language expressions, where such contents are ultimately found in the world, or constructed mathematically out of pieces of reality. On this view, semantics makes use of lexical postulates that express genuine relations between words and objects or collections of objects, and from these premises, semanticists derive theorems about what the world must look like for natural language sentences to be true. This is the Lewisian conception of semantics [Lewis, 1970, 1984, 1986], the main ideas of which are held by Soames [1987, 1989, 1992, 2002, 2005], Sider [2011], Williamson [2013], and Yablo [2014], among many others. Lewisian semantics is partly a metaphysical theory—it is a version of the theory of truthmaking.^{1,2}

In this paper, I argue for the adoption of a novel view of our foundational semantic notions that allows us to capture the core insights of each of these two views of semantics while also revealing how they conflict, and how they are systematically related. Philosophers of language and semanticists working both inside and outside of the Chomskian tradition have largely assumed that reference, application, and truth (of) are purely extensional, and state relations between words and particular objects or other pieces of reality. This is why some Chomskians have claimed that semantics should jettison the notions of reference and truth altogether, while many Lewisians have claimed that semantics should not concern itself with speakers, their psychologies, or what makes them linguistically competent.³ My

¹These two conceptions of the nature and subject-matter of semantics go by various names in the literature, but the distinction is ubiquitous. Sider [2011] aptly terms the two conceptions “linguistic semantics” and “metaphysical semantics”, although his conception of metaphysical semantics is a bit more specific than the one applicable here. The idea of truth-making comes largely from D. M. Armstrong [1997, 2004]. I recognize that there are important differences between the positions here, but each of them holds, roughly, that semantics states word-world relations, and that content is externalistic.

²Following Davidson [1966, 1967a], some theorists, most notably Larson and Segal [1995] have tried to maintain that semantics is both a theory of semantic competence and a theory of word-world relations, but these views remain problematic for various reasons. See Szabó [1997] and Gross [2006] for discussion.

³Importantly, not all Chomskians claim that we should jettison the notions of reference and truth altogether, although this is the lesson that some, including Pietroski [2003, 2005, 2006] at various points, draw from Chomsky’s arguments. I think that the best way of understanding Chomsky’s own comments in [Chomsky, 1995] and [Chomsky, 2000] is as endorsing a view on which semantics does make use of reference,

central claim is that there are strong empirical and theoretical reasons to treat the verbs we use in our semantic theorizing—including “refers (to)”, “applies (to)”, and “is true (of)” —as intensional transitive verbs (ITVs). ITVs have two readings: an intensional, *de dicto* reading as well as a relational, *de re* reading. Stating our semantic theory with the *de dicto* readings of our semantic verbs yields a theory that captures the core insights of the Chomskian approach to semantics, while the *de re* reading yields a theory that is fully relational, and issues in truth-conditions. These two approaches are related—and compatible—in that they are expressed by two different readings of the very same semantic vocabulary, and plausibly, the distinction between these two readings is one of scope.

A semantic theory stated with the *de dicto* readings of our semantic verbs can serve as a Chomskian semantic theory because it provides us with new, intensional versions of reference and application that satisfy several important Chomskian *desiderata*.⁴ The *de dicto* reading of a sentence involving an ITV can be true even when the noun phrase in its object position is empty, and also when that noun phrase does not pick out a specific thing. The object-positions of ITVs also resist substitution of even co-intensive noun phrases. This allows the theory to assign extremely fine-grained semantic values to expressions, even when those expressions are empty or do not pick out a particular object. The ultimate nature of this theory will depend on the semantics we provide for ITVs more generally, but on several plausible views, including the one I favor, the correct semantics for the *de dicto* reading is non-relational. However, ITVs also have a reading on which none of these intensional features are present: their *de re* reading. The *de re* reading of a sentence containing an ITV expresses a relation between the subject and a particular, existent object or collection of objects, and does so independently of how that object or those objects are characterized by the object position of the sentence.⁵ Stating a semantic theory with this reading of our

application, and truth, but construes them non-relationally, or intensionally.

⁴Many of the arguments that Chomsky gives for abandoning the relational conception of semantics are based on the fact that reference and application exhibit intensional features, and he seems to hold that expressions refer, but that reference is not a relation, at least to ordinary objects. Admitting a *de dicto* reading of our semantic postulates allows us to capture this view precisely. More on this explication of Chomsky’s view in §7.

⁵Chomsky also allows that we can introduce technical senses of reference, application, and truth that allow speakers to talk about the same stuff, for instance, in science (see Pietroski [forthcoming, p. 6] and references therein). Thus, this proposal should be particularly amenable to the Chomskian, particularly because, as we will see, I think that the relational readings of our semantic verbs are just this: technical readings that are stipulated and divorced from ordinary usage.

semantic vocabulary allows us to spell out the relational conception of semantics, on which semantics has metaphysical implications.

My argument begins by showing that, in English, semantic verbs like “refers (to)” and “applies (to)” exhibit all of the features of intensional transitive verbs. However, in English these verbs are used to report things that speakers do: they are used to report speaker’s reference and application. But when these verbs are used to state the semantic features of words, as they are in semantic theorizing, they are used technically. Accordingly, I provide several arguments that the technical usage should incorporate the intensional features of the natural language expressions. First, I argue that all theories of semantic reference appeal to speaker’s reference in their explanations—a fact that is rarely acknowledged—and so semantic reference should inherit the intensionality of speaker’s reference. I then argue that our technical terms need to ultimately be explained using non-technical vocabulary that we already understand, especially when such vocabulary is readily available. Further, in the case of our semantic verbs, there are practically no theoretical disadvantages of incorporating ITVs into our semantic theory because ITVs subsume the traditional, extensional semantic notions as special cases. Lastly, making use of intensional transitive verbs in our semantic theorizing is theoretically enriching in a number of important ways, one of which is that it allows us to capture and systematize the relationship between the two conceptions of semantics above. But the view also allows us to make headway on several recalcitrant problems in the philosophy of language and the foundations of semantics, including the problem of empty names and the Foster problem, along with its intensional variant.

1.2 Intensional Transitive Verbs

A verb is considered transitive when it takes a noun phrase in its object position, occurring in sentences of the form **NP V NP'**.⁶ A transitive verb **V** is considered intensional when sentences of the above form exhibit some combination of the following three properties.

⁶However, many verbs that are technically intransitive are treated as transitive when they occur in constructions of the form **NP V P NP'**, where **P** is a preposition. This is typically done when the combination of intransitive verb and preposition have a transitive verb as a near-synonym, such as in the case of “seeks” and “is looking for”. Many of these verb + preposition combinations behave identically to transitive verbs, and so unless otherwise noted, I’ll treat them as intensional transitives also.

Nonexistence: $\text{NP V NP}'$ has a reading which fails to entail NP' exists, where NP' is upward-entailing.⁷

Nonspecificity: $\text{NP V NP}'$ has a reading that fails to entail $\text{NP Vs a particular NP}'$.

Opacity: $\text{NP V NP}'$ has a reading that fails to entail NP V NP^* , where NP' , and NP^* are extensionally equivalent.⁸

To see these properties in action, let's consider a canonical example. We can see that the verb phrase "looking for" exhibits NONEXISTENCE by noting that (1) has a reading does not imply (2):

- (1) John is looking for the fountain of youth.
- (2) The fountain of youth exists.

This establishes that "looking for" exhibits NONEXISTENCE.

"Looking for" also exhibits the second property of ITVs, NONSPECIFICITY. Consider a case where "looking for" has an indefinite noun phrase in its object position, such as the following:

- (3) John is looking for a capable business partner.

Clearly, there is a reading of (3) that does not entail (4):

- (4) John is looking for a particular capable business partner.

John may merely be seeking to share his entrepreneurial tasks with someone he thinks will help his business, and he might be satisfied with a great number of different individuals.

We can bring this out with the following continuation:

⁷By "upward entailing" I mean to include positively quantified NPs like: "a dog", "the men who robbed him", "four gorgons", "infinitely many numbers", as well as proper names, and bare plural NPs. I mean to exclude negative NPs like "no dogs", "no one", *etc.* By "empty" I mean that nothing in the world answers to the NP.

⁸It's important to note that there are many ITVs that do not have all of these features: typically, the presence of even one of these properties is sufficient to classify a verb as intensional. However, verbs like "seek" are paradigmatically intensional in that they exhibit all three of the properties. As will become clear, I am arguing that not only are our semantic verbs intensional, they are like "seek" in being paradigmatically intensional.

(5) John is looking for a capable business partner—but no one in particular.

Lastly, “looking for” exhibits OPACITY: given two coextensive NPs, substitution of one for another within its complement does not preserve truth:

(6) John is looking for Ortcutt.

(7) John is looking for the shortest spy.

In this case, John might not know that Ortcutt is the shortest spy, and so the goal of his search may be to find Ortcutt and not the shortest spy. Thus (6) may be true while (7) is false, which means that “looking for” exhibits OPACITY.

These inferential tests indicate that there is a reading of sentences containing ITVs on which their object-position is not existence-entailing, can receive a nonspecific interpretation, and resists substitution of co-extensive expressions. However, there is also a reading that does not have these features. Consider John’s search for a capable business partner above. As we saw, John need not be looking for any particular person. However, he might be, and (3) can also be used to report just such a search. We can bring out this other kind of search with the following paraphrase:

(8) There’s a particular capable business partner for whom John is looking.

The truth-conditions of (8) differ from those of the reading which we brought out with (5) above. This indicates that (3) is ambiguous between two readings. I will call the reading brought out by (8) the *de re* reading of (3), and the reading brought out in (5) its *de dicto* reading. Distinguishing between these two readings is in keeping with a long tradition. Quine [1956] originally distinguished between what he called the notional and relational readings of sentences like:

(9) I want a sloop.

The relational, *de re* reading of (9) can be brought out with the following paraphrase:

(10) There is a sloop such that I want it.

The notional, *de dicto* reading can be captured by the idea that I seek “mere relief from slooplessness”, and brought out with the continuation in (11):

(11) I want a sloop—but no particular one.

In the current literature these two readings are often called “specific” and “nonspecific” or “extensional” and “intensional”.⁹ In what follows, I will continue to use the terms *de re* and *de dicto*, because they are somewhat more theoretically neutral than the other pairs of terms used to mark the distinction, and they don’t tie the two readings to any one of the properties of ITVs. Additionally, the distinction between *de re* and *de dicto* is commonly captured in terms of scope, which I think is the best way to capture the distinction between the two readings of ITVs.¹⁰ We will return to issues of scope below.

The non-equivalence of the (scopal) readings of a construction involving a transitive verb is sometimes seen as criterial for the intensionality of that verb, because the resulting ambiguity is not present in purely extensional verbs. A test for this non-equivalence often appears under the name “failure of quantifier exportation” [Moltmann, 1997].¹¹ If the

⁹The ambiguity is sometimes taken to arise only when the NP in object position is an indefinite description [Moltmann, 1997, Zimmermann, 1993, 2001, 2006]. But like Mark Richard [2013], I think this is mistake. While there may be a specific/nonspecific ambiguity that arises in connection with indefinite descriptions, this is simply a special case of the ambiguity that is characteristic of ITVs, which is much broader, and can occur with definite as well as indefinite NPs in object position: for instance, in “John imagined London” or “John needs the antidote”. Thus I differ from semanticists who take Nonspecificity as a necessary condition for a transitive verb to qualify as intensional.

¹⁰While a scopal analysis of the *de re/de dicto* distinction is plausible, the relationship between the *de re/de dicto* distinction is complicated. I adopt a two-way distinction here merely for ease of exposition, but recognize that, as Kripke [1977] showed, no two-way distinction can do justice to iterated intensional verbs, and the scopal readings they generate. Ultimately I believe that explaining the different readings of intensional sentences as different scopal readings is the correct explanation, and so may need to jettison the terminology I have chosen to use here. However, Janet Fodor [1970], in her dissertation, shows that intensional verbs have more than just two readings—she claims that in some cases they have four, and argues that they lack enough scopal readings to capture the four-way distinction. Fodor claims that the intensional status and the quantificational force of phrases in intensional positions can be evaluated independently. The four readings then correspond to each of the four possible combinations of (the presence or absence of) NONSPECIFICITY and OPACITY. If the basic scopal analysis holds, it would predict only two of the readings, since on the scopal analysis, the entire noun-phrase can scope only either over or under the verb, and thus, NONSPECIFICITY and OPACITY are predicted to co-occur. This indicates that there are not enough permutations of scope-bearing elements in intensional sentences to capture their readings, and so the different See [Keshet, 2008] for an overview, [Szabó, 2010] for a defense of Fodor’s specific opaque reading, and Keshet [2011] for a new scopal account of *de re* and *de dicto* that accommodates Fodor’s data.

¹¹Failure of quantifier exportation is an idea originally due to Quine [1956], but see [Kaplan, 1968] for a discussion. Richard [2013] calls the two scopal readings of intensional constructions the D-reading and the

quantifier in the verb’s complement fails to export, and can yield a falsehood when moved to a position where it takes scope over the verb, this shows the non-equivalence of the two readings, as in the following example:

- (12) John is looking for a unicorn. \nrightarrow
 A unicorn is such that John is looking for it.

In more generality, the inference that fails is :

- (13) **NP Vs Q N** \nrightarrow
Q N is/are such that **NP Vs** it/them. Moltmann [1997]

I will not take failure of quantifier exportation as criterial for intensionality, because I think the test is more coarse-grained than the tests mentioned above: quantified NP complements can fail to export either because they are non-specific, or because they are empty. However, I will sometimes treat the ability to elicit two distinct readings as weak evidence for the intensionality of a verb, due to the fact that judgments about intensionality can be subtle, and testing for the presence of a second reading provides us with another resource for its detection.

1.3 Representational Verbs

Analyses of ITVs often restrict themselves to considering just a few paradigmatically intensional verbs, for example: “seek”, “need”, and “want”. This can sometimes give the impression that the class of verbs which displays some combination of the above features is relatively small. However, Friederike Moltmann [2008] lists six categories of transitive verbs that have intensional readings:

1. (Simple) predicates of absence: *need, lack, omit, fit (into, onto)*
2. Psychological verbs of absence: *promise, desire, want*
3. Predicates of transaction and possession: *own, possess, owe, offer, buy, accept, have*

R-reading, and takes the presence of the ambiguity as criterial for intensionality.

4. Verbs of representation: *draw, paint, portray, imagine, represent, show, indicate, point (to), talk (about), signify*
5. Epistemic predicates: *see, recognize, find, discover, count*
6. Verbs of creation in the progressive: *is building, is creating, is putting together*

This shows that the category of intensional transitive verbs is surprisingly broad, especially considering that several of the “epistemic predicates” are usually taken to be paradigmatically *extensional*. However, for our purposes, the most important category of ITVs are the verbs of representation. Not only are these verbs intensional, in that they exhibit NONSPECIFICITY; they exhibit all three properties above, making them paradigmatically intensional.¹²

In what follows, I will argue that “refers to” and the other semantic verbs mentioned above are paradigmatically intensional. Their intensionality, together with the fact that semantic verbs clearly express intentional notions, makes it plausible that semantic verbs fall into the category of representational verbs. More specifically, my arguments will establish that, in English, “refers to” is roughly synonymous with “talks about” or “is about”. The intensionality of notions of aboutness and subject-matter is well-established [Martí, 1989, Perry, 1989], and it has been widely noted that “about” is an intensional preposition, which occasions intensional contexts [Montague, 1974b]. But while the intensionality of aboutness is well-known, the intensionality of “refers to” and “applies to” is surprising: what words

¹²To see this, consider the following examples:

- (14)
- a. The hammer and sickle represent a strong and industrious nation.
 - b. Dali drew a strange man.
 - c. The movie portrayed a pair of outlaws.
 - d. John imagined a distant city.
 - e. Newly developed economic metrics indicate a rise in stock prices.

It may well be the case that all strong and industrious nations are also unjust nations, and vice-versa, but the hammer and sickle need not represent an unjust nation. Similarly, all strange men may be sad men, and vice-versa, but Dali need not have drawn a sad man. Similar arguments can be made for the rest of the verbs. Thus representational verbs exhibit OPACITY. It is also quick to see that none-of these verbs are existence-entailing. What about NONSPECIFICITY? Clearly, the hammer and sickle need not represent a particular strong and industrious nation, Dali need not have drawn a particular man, and John need not have imagined a particular city, nor do the new economic metrics need to have indicated a particular rise in stock prices. Thus, these verbs display all three traditional features of ITVs, and are paradigmatically intensional.

refer to and apply to are typically taken to be thoroughly extensional notions, and “refers” is supposedly an extensional verb *par excellence*.¹³

One last verb deserves comment: “means” is also paradigmatically intensional, and this fact is highly suggestive. If semantics is supposed to be a theory of meaning (and what else could it be?), then it seems that any collection of semantic verbs that does not exhibit intensionality in the way that “means” does is certain to be inadequate for specifying a theory of meaning.¹⁴ Showing that semantic verbs have an intensional reading goes a long way toward showing that they can serve to state such a theory.

1.4 Speaker’s Reference and Application are Intensional

This section presents empirical data showing that “refers to”, as it is used in English to report speaker’s reference, is much closer to intensional than extensional with respect to all three of the core features of intensionality.¹⁵ To collect this data, I designed and ran three studies, each of which tested “refers to” for one of the three traditional features using the

¹³My proposal is connected to a point made by David Lewis, in his paper “Tensions” [Lewis, 1983], that has been drastically underappreciated. In the paper, Lewis shows that there is, in an important sense, no absolute difference between languages that are extensional and languages that are intensional. Instead, given a language in which every expression is assigned an intension, that language can be transformed into a language in that is fully extensional: just let each expression of the new language have, as its extension, the function that was the intension of the expression in the original language. Given a certain approach to the semantics of ITVs, this is what treating “refers to” and “applies to” does: it makes an expression’s *intension* its referent.

¹⁴This is closely related to points made by Davidson [1967a, 1976] in response to what has come to be known as the Foster Problem [Foster, 1976]. Foster famously showed that a theory of truth could issue in theorems that were not interpretive. He pointed out that the theorems of a truth-theory—biconditionals pairing sentences of the object-language with their truth-conditions—did not provide a tight enough connection to serve as meaning-theorems. For example, such a theory could have theorems that were true but obviously not meaning-giving, such as “‘Snow is white’ is true iff grass is green”. Davidson responded by claiming that the biconditionals needed to be laws of nature, and should be prefixed with an intensional operator, “Necessarily”, which rules out the simplest such cases. However, Soames [1989] recapitulates the Foster problem in the intensional setting, and tries to show that no theory that derives truth-conditions from reference and satisfaction clauses can suffice as a theory of meaning. Soames shows that given any reference and satisfaction clauses, ones that are *intensionally* equivalent can be constructed, and these clauses allow us to derive identical truth-conditions for sentences with obviously different meanings. In a sense, my strategy is the reverse. On my view, reference and satisfaction clauses have a reading that is hyperintensional, and so Soames’s intensionally equivalent clauses can be distinguished from one another. This means that derivations involving them will yield different meanings, although it is unclear whether these derivations will themselves provide specifications of truth-conditions. Truth-conditions can be derived from the other reading of our reference and satisfaction clauses: their extensional reading.

¹⁵The intensionality of speaker’s reference is not a new idea. However, this is, to my knowledge, the first time that the standard linguistic criteria for being an intensional transitive verb have been explicitly applied to “refers” to establish its intensionality. It is also, to my knowledge, the first empirical work has been done to support the conclusion.

inferential tests laid out above. Each study compared “refers to” to one paradigmatically intensional and one paradigmatically extensional transitive verb, and then took note of statistical differences with respect to one of the properties. As we will see below, “refers to” was closer to intensional in all three studies, and in the cases of NONEXISTENCE and OPACITY, did not differ statistically at all from “seeks”, a paradigmatically intensional verb.

1.4.1 Experiment 1: Nonexistence

The first study tested “refers to” for NONEXISTENCE.

Methods

In the study, 237 participants filled out a brief questionnaire.¹⁶ Each participant was randomly assigned to one of three conditions, either **Intensional**, **Refers**, or **Extensional**, and answered one question associated with that condition. In what follows, I will refer to the three conditions—intensional vs. refers vs extensional—as “verb categories”. The questions associated with the verb categories differed only in that they contained either a paradigmatically intensional verb (“search”), “refer to”, or a paradigmatically extensional verb (“touch”) as their main verb. To make sure the results were not peculiar to one particular noun phrase, participants were then assigned to one of four vignettes (**Unicorns**, **Elves**, **Magical Fountains**, or **Dodos**).

Each participant was asked to suppose that they knew that a certain kind of entity did not exist, but that their friend, John, didn’t. For example, the first vignette consisted only of the following sentence:

Unicorns Suppose that you know that unicorns do not exist, but your friend John doesn’t.

The participant was then asked one of the three questions below, depending on the condition to which they had been assigned:

Intensional Is it possible for John to search for a unicorn?

Refers Is it possible for John to refer to a unicorn?

¹⁶Participants were recruited using Amazon’s Mechanical Turk. Sample was 51.25% male, mean age 35.5.

Extensional Is it possible for John to touch a unicorn?

The other vignettes, and the questions associated with them, differed only in that they had a different indefinite noun phrase in place of “a unicorn”; instead, the other three used “an elf”, “a magical fountain”, and “a dodo”, respectively. Participants responded to the questions on a 7-point Likert scale, with an answer of 1 indicating a response of “definitely not” and 7 indicating “definitely yes”. Thus, if a participant responded with a high score on a question, it indicated that the participant took the verb to exhibit NONEXISTENCE, whereas a low score indicates the opposite.

Results

The average rating for each of the verb categories across the four vignettes can be found in Figure 1.1. The key thing to notice is that the ratings for “refers” were much closer




Figure 1.1: Mean ratings by condition in Experiment 1. Error bars show standard error.

to those given for the paradigmatically intensional verb than the extensional verb. When these averages were compared, the rating for “refers” did not differ significantly from the intensional case, but was significantly higher than that for the extensional case ($p < .001$),

and this difference was consistent across the four vignettes. Further, comparing the means for the **Refers** and **Extensional** conditions revealed a large effect size.¹⁷

Discussion

The results are striking, and seem to establish unequivocally that “refers” exhibits NONEXISTENCE: it appears to pattern completely with “search for”, and bear very little similarity to “touch”. Consider the following sentences:

(15) John is referring to a unicorn.

(16) John is referring to an elf.

On the supposition that an affirmative response to the **Refers** question above indicates that (15) has a reading that does not entail the existence of unicorns, then the results indicate that sentences such as (15) and (16) exhibit NONEXISTENCE. This, I think, should be somewhat surprising; it is often the case that philosophers of language take genuine reference to require existence. These results pose a dilemma for such theorists: either they are flatly wrong about the nature of reference, or the version of reference with which they are concerned is not the one that ordinary speakers make use of and have intuitions about. I will discuss this question at length in §5 and §6.

1.4.2 Experiment 2: Nonspecificity

The second experiment tested “refers” for NONSPECIFICITY.¹⁸

¹⁷The results were analyzed using a 3 (verb category: intensional vs. “refers” vs. extensional) x 4 (vignette) ANOVA. As expected there was a significant main effect of verb category, $F(2,225) = 66.6$, $p < .001$, but there was no significant main effect of vignette, $F(3,225) = 2.2$, $p = .084$, and no significant interaction, $F(6,225) = 1.2$, $p = .3$. To explore the differences between the intensional case, the extensional case, and the case of “refers”, I used Tukey’s post-hoc tests. Unsurprisingly, participants gave higher ratings in the intensional condition ($M = 5.35$, $SD = 2.1$) than in the extensional condition ($M = 1.96$, $SD = 1.78$), $p < .001$. Ratings for “refers” ($M = 4.99$, $SD = 2.16$) were significantly higher than those for the extensional case, $p < .001$, $d = 1.53$, but not significantly different from those for intensional case, $p = .495$.

¹⁸Two anonymous referees point to the fact that NONSPECIFICITY is itself a property that is slightly unclear. For instance, if John is looking for a dog, he might be looking for a specific *property*, even if he is not looking for a specific *dog*. I take NONSPECIFICITY to be present in cases where an agent is not related to any particular *entity*. The idea that an ITV might relate the subject to a specific *property* is, I believe, a piece of theory that attempts to explain the basic intensional datum, which is that when indefinites appear in the object position of an ITV, they need not pick out particular entities: they need not refer, or provide an entity that serves as the argument to the verb. Further, such indefinites are not merely instances of what

Methods

In the study, 236 participants answered three questions each, one question for each of the three conditions, **Intensional**, **Refers**, and **Extensional**.¹⁹ The **Intensional** question contained a paradigmatically intensional verb (“look for”), the **Refers** question involved “refers to”, and the **Extensional** question contained a paradigmatically extensional verb (“touch”). The questions were presented in a random order. Each question asked the participant to suppose that the subject was involved in a particular activity or in a particular state, and then queried whether it was possible for the activity or state to be directed toward something nonspecific. The activities and states were all characterized using an indefinite NP in the object position of the main verb, so the questions assessed whether a nonspecific interpretation was available for the indefinite. To make sure that answers did not depend on the specific NPs used in the questions, each participant was assigned to one of five vignettes (**Dog**, **Person**, **Book**, **CC-Cookie**, or **Cigarette**) at random, each of which involved a different indefinite noun phrase. For instance, the first set of questions was as follows:

Intensional Suppose that John is looking for a dog. Is it possible for John to be looking for a dog, but not a particular one?

Refers Suppose that in a conversation, John is referring to a dog. Is it possible for John to be referring to dog, but not to a particular one?

Extensional Suppose that John is touching a dog. Is it possible for John to be touching a dog, but not a particular one?

The other sets of questions differed only in that they contained a different indefinite NP within the complement of the transitive verb. Since each respondent answered a question containing a verb from each category, they were able to compare the three questions asked, and adjust their answers accordingly. Participants responded to the questions on a 7-point Likert scale, with an answer of 1 indicating a response of “definitely not” and 7 indicating

Zimmermann [2001] calls “unspecificity”: namely, cases where what particular entity the indefinite picks out is left *unspecified*. I take the inferential test I introduced above and the questions in Experiment 2 to be genuine tests for nonspecificity, rather than unspecificity. When an object is left unspecified, the continuation “but no particular one” is not appropriate.

¹⁹Participants were recruited using Amazon’s Mechanical Turk. Sample was 63.1% male, mean age 25.5.

“definitely yes”. Thus, if a participant responded with a high score on a question, it indicated that the participant took the verb to exhibit NONSPECIFICITY, whereas a low score indicates the opposite.

Results

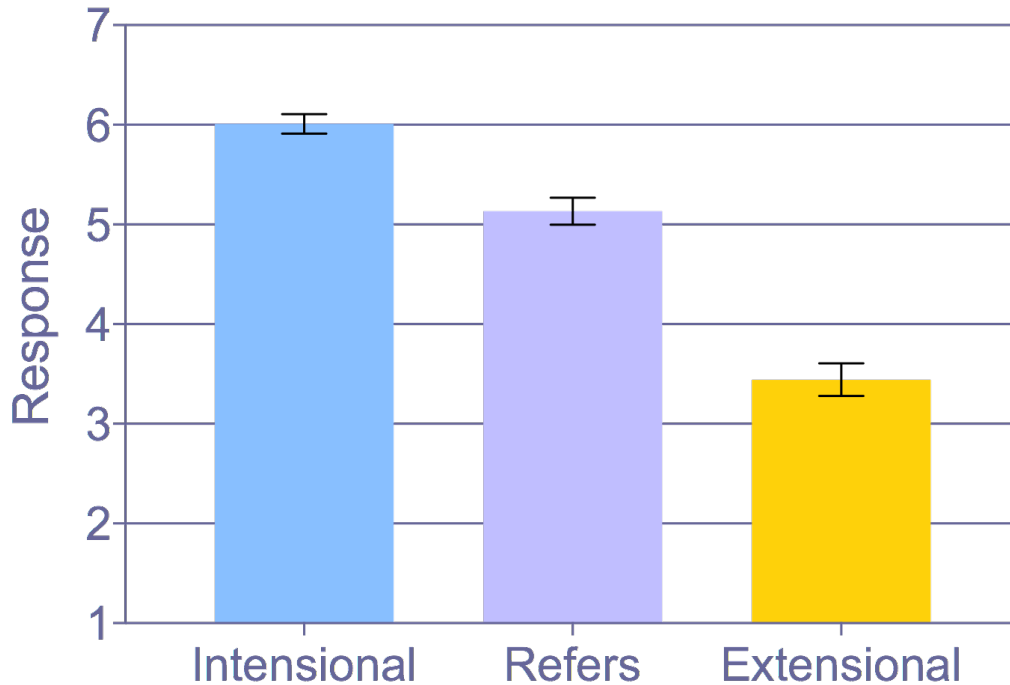


Figure 1.2: Mean ratings by condition in Experiment 2. Error bars show standard error.

As Figure 1.2 shows, the average rating for “refers” was intermediate between the intensional and extensional cases. Ratings for “refers” were significantly lower than the intensional verb, and significantly higher than the extensional one.²⁰ However, the average was still closer to intensional than extensional, and the effect size when comparing the mean for **Refers** to the for **Extensional** was larger than the effect size when comparing **Refers** to **Intensional**.²¹

²⁰The data were analyzed using a mixed-model repeated measures ANOVA, with verb category (**Intensional** vs. **Refers** vs. **Extensional**) as a within-subject variable and vignette (**Dog** vs **Person** vs. **Book** vs. **CC-Cookie** vs. **Cigarette**) as a between-subject variable. As we would expect, there was a significant main effect of verb category, $F(2, 231) = 100.4, p < .001$. There was no significant main effect of vignette, $F(4,231) = 1.5, p = .192$. There was a significant interaction, $F(8,462) = 3.2, p = .002$.

²¹To further explore the effect of verb category, and establish whether “refers” is intensional or extensional, I ran separate ANOVAs comparing each pair of verb categories. As we would expect, ratings for the intensional verbs ($M = 5.99, SD = 1.53$) were higher than those for the extensional verbs ($M = 3.45, SD =$

Discussion

The results show that “refers” differs significantly from both paradigmatically intensional and paradigmatically extensional verbs with respect to NONSPECIFICITY; the average for “refers” was intermediate between the intensional and extensional cases. This indicates that with respect to NONSPECIFICITY, certain verbs can have an intermediate status. This intermediate status poses a question for standard ways of categorizing verbs as intensional vs. extensional. Most semanticists take the distinction to be binary: either an intensional reading is available or it is not. The results show that a more nuanced approach is required. It may be that intensional readings are heard by some speakers and not others, or it may be that many speakers hear a genuinely intermediate rating. Preliminarily, the variance in responses for “refers” indicates that some speakers get the nonspecific reading while others do not. But then this poses a problem for the lexical semantics of intensional verbs: when speakers are divided about a verb’s intensionality, should a lexical-semantic theory encode it?

However, this problem is not too worrying in the case of “refers”. Across the vignettes, the mean for “refers” was much closer to intensional than to extensional. If we keep to the relatively standard assumption that the distinction between intensional and extensional is binary, then it seems we have good grounds for saying that “refers” exhibits NONSPECIFICITY. The averages clearly fall on opposite sides of the midpoint of the scale, which is 4.

Overall, this is an even more surprising result than that of the first experiment.²² Even when we restrict ourselves to consideration of speaker’s reference, reference is ordinarily

2.45), $F(1,231) = 155.7, p < .001$. Ratings for “refers” ($M = 5.14, SD = 2.07$) were significantly higher than those for extensional verbs, $F(1,231) = 86.25, p < .001, d = .74$. Ratings for “refers” were also significantly lower than those for intensional verbs, $F(1,231) = 24.05, p < .001, d = .48$. Looking at the differences between vignettes, we found that, in contrast to the first experiment, there was an interaction between which indefinite NP was involved in the vignette and whether a nonspecific reading was available for “refers”. In particular, in one of the vignettes—**CC-Cookie**—the ratings for “refers” were closer to extensional, although still intermediate, while in the other four vignettes the ratings were closer to intensional. There are two possibilities for explaining this interaction. One is that the **CC-Cookie** vignette was an anomaly. The other is that there is a genuine interaction between the NP in the object position of a verb, and whether a nonspecific reading is available.

²²It is also worth noting that the sample size for this experiment was quite large: in contrast to the other experiments, each of the participants in this experiment answered 3 questions: one for the intensional condition, one for “refers”, and one for the extensional condition, and so there were 237 data points available for each question, as opposed to around 80 in the other experiments.

presumed to be fully specific. The fact that a nonspecific reading is often available gives us strong reason to think that reference is not a relation between a speaker (or a word) and an object. Nonspecificity cannot be explained by positing nonexistent objects, as is often done to explain NONEXISTENCE, or by positing senses or conceptual covers, as is often done to account for OPACITY.

1.4.3 Experiment 3: Opacity

The third experiment tested “refers to” for OPACITY.

Methods

In the study, 231 participants filled out a brief questionnaire.²³ Each participant was randomly assigned to one of three conditions, either **Intensional**, **Refers**, or **Extensional**, and answered one question associated with that condition. As with the previous experiments, the questions associated with the verb categories differed only in that they contained either a paradigmatically intensional verb (“search”), “refer to”, or a paradigmatically extensional verb (“touch”) as their main verb. For the sake of generality, each participant was randomly assigned to one of four vignettes (**Art Collector**, **Chief Justice**, **Murderer**, or **Spy**). As an illustration, the first vignette was the following:

Art Dealer Suppose that the person with the largest art collection in the country just so happens to be the national record holder in the high jump.

Participants were then asked to answer one of the following three associated questions:

Intensional Now suppose that Mary is looking for the person with the largest art collection in the country. Does it have to be true that Mary is looking for the national record holder in the high jump?

Refers Now suppose that in a conversation, Mary is referring to the person with the largest art collection in the country. Does it have to be true that Mary is referring to the national record-holder in the high jump?

²³Participants were recruited using Amazon’s Mechanical Turk. Sample was 46.9% male, mean age 33.5.

Extensional Now suppose that Mary is touching the person with the largest art collection in the country by shaking his hand. Does it have to be true that Mary is touching the national record-holder in the high jump?

The other vignettes, and the questions associated with them, differed only in that they made use of a different pair of definite descriptions. Participants answered their question on a 7-point Likert scale, with an answer of 1 indicating “definitely not” and an answer of 7 indicating “definitely yes”. In this case, lower scores indicated a higher degree of OPACITY.

Results

The average rating for the three questions corresponding to the different verb categories can be found in Figure 1.3. As in the first study on NONEXISTENCE, but in contrast to the




Figure 1.3: Mean ratings by condition in Experiment 3. Error bars show standard error.

second study on NONSPECIFICITY, the average rating for “refers” did not differ significantly from the average rating for the intensional case. But as with both previous studies, “refers” differed significantly from the average rating for the extensional case, with a moderate effect

size.²⁴ In contrast to the previous two studies, however, a significant main effect of vignette was observed, and like the study addressing NONSPECIFICITY, we observed a significant interaction between vignette and the status of “refers”.

Discussion

As we can see from the figure above, “refers” again patterns with the paradigmatically intensional verb as opposed to the paradigmatically extensional verb. However, this data is less clear than the previous two experiments. We observed an overall effect of vignette on participants responses, and also an interaction between the vignette and participants ratings for “refers” relative to the controls. I think these facts are due to a slight anomaly in two of the four intensional questions, which received higher ratings than the others. The effect of the intensional questions getting these lower scores was to bring the overall intensional average up, and closer to that of “refers”. Thus, while the average for “refers” does not differ significantly from that of the intensional condition, my conclusion is that with respect to OPACITY, “refers” is best construed as intermediate between intensional and extensional, rather than patterning perfectly with verbs of search.

Even though it is likely that the results only support an intermediate status for “refers” with respect to OPACITY, this intermediate status is still surprising. This shows that our ordinary notion of reference differs from the technical notion of extension with respect to granularity: reference is, to some degree, dependent on description. The idea that reference is description-dependent is even more surprising considering that opaque contexts are often *defined as* contexts in which coreferential terms are not substitutable. If we keep this definition, but treat “refers” as opaque in its object position, it may turn out that no contexts are opaque. They are all fully extensional; it is just that “refers” itself is opaque. But alternatively, we can define an opaque context as one in which co-extensive expressions

²⁴The results were analyzed using a 3 (verb category: **Intensional** vs. **Refers** vs. **Extensional**) x 4 (**Art Collector** vs **Chief Justice** vs. **Murderer** vs. **Spy**) ANOVA. There was a significant main effect of verb category, $F(2,219) = 16.67, p < .001$, and a significant main effect of vignette, $F(3,219) = 6.47, p < .001$. To explore the differences between the intensional case, the extensional case, and the case of “refers”, I used Tukey’s post-hoc tests. Unsurprisingly, participants gave lower ratings in the intensional condition ($M = 3.61, SD = 2.52$) than in the extensional condition ($M = 5.65, SD = 2.00$), $p < .001$. Ratings for “refers” ($M = 4.19, SD = 2.45$) did not differ significantly from the intensional case, $p = .215$, but were significantly lower than those for the extensional condition, $p < .001, d = .64$. We also observed an interaction effect of verb category and vignette, $F(6,219) = 3.98, p = .001$.

are not substitutable, although this will force us to divorce reference from extension.

1.4.4 Conclusions

Together, the results of the three above studies indicate that speakers recognize intensional readings of sentences such as

(17) John is referring to a unicorn.

(18) John is referring to a dog.

(19) Mary is referring to the person with the largest art collection in the country.

On its intensional reading, (17) can be true, but does not entail the existence of unicorns. Similarly, speakers seem to recognize a reading of (18) on which John is referring to a dog, but not to a particular one. And finally, speakers recognize a reading of (19) on which Mary need not be referring to the national record-holder in the high jump, even if that person happens to also be the person with the largest art collection in the country. But “refers”, like “seeks”, also has an extensional reading: there is a reading of (18) on which John is referring to a particular dog, just as he may be seeking a particular dog. As mentioned above, I will call the intensional reading of a sentence involving and ITV its *de dicto* reading, and I will call its extensional reading its *de re* reading, and I will call the forms of reference reported by these two readings “reference *de dicto*” and “reference *de re*”.

However, one might worry that, even given the data above, speakers’ judgments are not being driven by the presence of a genuine *reading* of the sentence whose presence needs a semantic explanation, but instead are being driven by pragmatics. After all, it is common practice to attempt to explain, or explain away, substitution failure within the contexts of attitude verbs by appealing to pragmatics. However, there are several reasons why such an approach is unlikely to succeed in this case. First, “refers” patterns quite closely with a paradigmatically intensional verb, “seeks”. It is possible that the intensional features of “refers” recorded in the studies are due to pragmatics, but if this is true, why should we not say the same for “seeks”? It doesn’t seem that there is a reason why we should treat the intensional features of the two verbs differently. Thus, one can hold that the

intensionality of “refers” is due to pragmatic features only if one believes the same thing about all intensional verbs, which amounts to the denial that there *any* genuinely semantic intensional phenomena in natural language. But most working semanticists believe that some verbs are intensional, and that this intensionality is a datum for which a semantic theory must account, and I am content for my view to depend on this view, which seems to be the consensus view among semanticists.

But there are also more specific reasons why the intensionality of “refers” cannot be pragmatic. Ordinarily, the intensional phenomenon that philosophers try to explain away using pragmatics is OPACITY. And I think that, restricting attention to this particular aspect of intensionality, these attempts at pragmatic explanation are plausible, for it seems that the phenomenon of OPACITY does not interact with the rest of our semantic machinery. It seems to be a relatively isolated phenomenon. However, the above studies show that “refers” also exhibits NONSPECIFICITY, which does interact with the rest of our semantic machinery, and these interactions give us reason to think that the phenomenon is semantic. First, the presence of a nonspecific reading for an object-position indefinite licenses a peculiar form of quantification that has come to be called *special quantification*, on which a quantifier replaces the entire quantified NP in object-position:

- (20) a. John is referring to a ruby.
b. John is referring to something.
c. John is referring to something valuable.

Special quantifiers are ordinarily existential, and are formed from combinations of a determiner and the morpheme “-thing”. Accompanying the possibility of special quantification are restrictions on the kinds of anaphora licensed by nonspecific indefinites. Notably, (20-a) does not license anaphoric reference with ordinary pronouns, nor does it entail readings on which the indefinite takes scope over the verb:

- (21) a. \nrightarrow John is referring to it.
b. \nrightarrow There is a ruby to which John is referring.
c. \nrightarrow A ruby is such that John is referring to it.

Rather, intensional indefinites only license special anaphora, which makes use of special pronouns and descriptions, such as “the same thing”, “one”, “what”, and possibly “that”:

- (22)
- a. John is referring to what Bill is referring to.
 - b. John is referring to one, too.
 - c. John is referring to a ruby. Bill is referring to the same thing.
 - d. Johh is referring to that (?)

The inferential behavior above seems to indicate that “refers” has two readings, only one of which licenses a nonspecific interpretation for the indefinite, and neither of which entail the other. Further, the two readings seem most naturally treated as two scopal readings, particularly in light of the sentences in (21), and scope is a distinctively semantic phenomenon. These inferential patterns are not easily explained pragmatically, and thus we have reason to take the intensional features of “refers” at face value.

One final point is in order: it is not just “refers” that is intensional in English. It is plausible that verbs expressing speaker’s predication, such as “ascribes” and “attributes”, also exhibit intensional features in both their direct and indirect object positions. Consider the following sentence:

- (23) Jack ascribes supernatural powers to a relic.

In (23), “a relic” can be read either specifically or non-specifically. Further, Jack may not ascribe healing powers to a small piece of wood, even if that’s just what the relic is (perhaps it’s a shard of the cross). Additionally, there may be no such property as possessing supernatural healing powers, and even if there is, Jack need not ascribe it to a relic by any other name. Thus, both the direct and indirect object positions of (23) are intensional. Further, if we replace “ascribes” with “attributes” in (23), the exact same arguments suffice to show that it is intensional as well.

1.5 Semantic Reference and Application are Intensional

All of the sentences in the studies above have speakers as subjects, which means that the studies have an important limitation: they only establish that “refers” is intensional when we use it to report what speakers are referring to. Borrowing a distinction from Kripke [1977], the studies show that speaker’s reference is intensional, but semantic reference may well still be intensional. This section will provide arguments that semantic reference does in fact inherit the intensionality of speaker’s reference. However, statements of semantic reference are less common in ordinary language than statements of speaker’s reference, and so the semantic verbs that figure in these statements are best considered technical terms. As a result, surveying native speakers about their features will not help us understand how they function; native speakers can’t be expected to have intuitions about technical terms, and further, since these terms are technical, semanticists are at liberty to stipulate their features. In light of this, the next three sections will provide some arguments that semantic reference *does* inherit the features of speaker’s reference, along with arguments that we *should* make use of intensional semantic vocabulary in our theorizing.

Let’s start with some metasemantic arguments. First, on many views, what a word refers to is ultimately determined by how speakers use that word: linguistic intentionality is explained in terms of the intentionality of thought. On such views, semantic reference will be determined by instances or patterns of speaker’s reference. But the last section showed that speaker’s reference is intensional: it is much closer to a paradigmatically intensional notion than it is to an extensional one. Accordingly, on views that privilege the intentionality of thought, it is natural to expect that the intensional features of speaker’s reference will carry over into our account of semantic reference. If semantic reference does inherit the intensionality of speaker’s reference, semantic reference will come in two forms: reference *de re* and reference *de dicto*. These two forms of reference are a generalization of the traditional notion of reference, which subsumes the traditional notion as a special case. Reference *de dicto* is a novel form of reference that exhibits the intensional features characteristic of the intensional readings of ITVs, while reference *de re* is the traditional, extensional notion of

reference that is transparent, specific, and existence-entailing.²⁵

Following Tim Crane [2012, p. 113], we can distinguish four main theories of linguistic intentionality: descriptive, pictorial, causal, and functional. I think it is reasonably clear that on the descriptive, pictorial, and functional accounts of intentionality, instances or patterns of speaker's reference ultimately determine what a word refers to, so I will only discuss them briefly. However, causal theories of reference also ultimately rely on speaker's reference, but I think this fact is less obvious, and rarely appreciated, so I will discuss the causal theory at greater length. In the case where an image fixes what a linguistic expression is about, it is plausible that the image is one associated with the expression by a speaker, and further, that no image could be about an object without an agent who intends or takes it to be so.²⁶ Similar things can be said about the descriptive case: presumably, if a linguistic expression refers to something in virtue of being associated with a particular description, it is because some particular speaker or group of speaker's associates this descriptive content with the expression. Thus the reference of the word will be dependent on instances of speaker's reference involving a description, which as we showed above, exhibit the features of intensionality.

The idea behind functional theories of reference is that a word's reference is determined by fulfilling a certain function: say, allowing an agent or group to successfully navigate their environment. On a view such as that in Millikan [2004], a word refers to an object just in case, roughly, taking it to refer to that object confers an advantage on an agent or group. This is explicitly a case in which patterns of speaker's reference serve to determine semantic reference. Thus, while I do not take these observations to remove all possibility that on one of these views, semantic reference could be determined independently of speaker's reference, I do take them to make the involvement of speaker's reference plausible, and thus make it

²⁵I am open to the possibility that the three features of intensionality can come apart; in fact, it is my belief that they do come apart, and can be treated separately. So like Fodor [1970], I do not think a two-way distinction is adequate to explaining them. However, I am using the terminology of *de re* and *de dicto* more to streamline the discussion than because I think the terminology captures a deep distinction. Rather, I think the important thing is just that semantic verbs are ITVs, and that ITVs have different readings that can be accounted for in terms of scope. The *de re/de dicto* distinction is only adequate for describing scopal distinctions in simple cases.

²⁶See Putnam's example of a likeness of Winston Churchill that happens to have come to be accidentally in the sand.

plausible that semantic reference inherits its intensionality.²⁷

Accounts of the causal theory of reference ordinarily begin with a discussion of a baptism: a case where a speaker initially uses a word and attempts to attach it to a piece of non-linguistic reality. Although various theorists differ on the details of how baptism works, and baptisms themselves come in various forms, their canonical form is when a speaker uses a word to pick out a piece of non-linguistic reality with which they are in causal contact, and this causal or perceptual contact then serves to fix the reference of the term on its subsequent uses. On other views, the baptist is seen as fixing a condition; when an object uniquely satisfies the condition, that object is fixed as the referent of the expression. However, causal theories struggle to account for cases of reference to abstract objects, such as numbers, where there is no causal connection. Given a causal theory of reference, it is not plausible to think that we refer to the number 6 in the same way that we refer to Barack Obama.

However, it has been largely overlooked that every case of baptism involves a speaker intending to refer to something, and thus involves an act of speaker's reference. This provides at least the beginnings of a solution to the puzzle of how we can refer to things with which we are not in causal contact. As we saw in the last section, speakers can successfully refer *de dicto* whether or not they are genuine causal contact with an object, and whether or not the condition they specify is uniquely satisfied, or satisfied at all. As a consequence, speakers can initiate causal chains with acts of speaker's reference *de dicto*, even if they fail to refer *de re*. Causal or perceptual connection to an object might be required for the term to have a *de re* reference, and when such connections are present, the baptized term will come to have both a *de dicto* reference and a *de re* reference: roughly speaking, an intension and an extension. However, when such connections are absent, the baptized term will not have an extension, but the *de dicto* speaker's reference will the name or term in question with distinctive intensional semantic features. That is to say, the act of speaker's

²⁷One anonymous referee makes the point, however, that it is plausible that speaker's thoughts are intensional and idiosyncratic—at least in terms of their granularity—in a way that meanings in a public language are not. Even though different speakers may all refer using, for instance, a mode of presentation, the extension is all that such uses have in common, and so we ought to assign the extension as the reference of that expression. However, this can be accommodated on my view, for saying that semantic reference is intensional is merely to say that it has an additional, intensional reading. The extension of a linguistic expression can be assigned with the *de re* reading of a reference clause.

reference built into every baptism can guarantee that the term is not semantically trivial. In a case like that of Leverriere and Vulcan, this means that we can capture Leverriere's act of reference with (24):

(24) Leverriere used "Vulcan" to refer to the planet responsible for the irregularities in Mercury's orbit.

And his successful act of speaker's reference initiated a particular usage for "Vulcan", which we can capture with the *de dicto* reading of (25):

(25) "Vulcan" refers to the planet responsible for the irregularities in Mercury's orbit.²⁸

It is this true reading that allows us to explain why (25) sounds true, while (26) is totally wrong:

(26) "Vulcan" refers to Phlogiston,

or why (27) seems to be a fine way of specifying one aspect of Vulcan's semantic profile:

(27) "Vulcan" refers to Vulcan, not to Nibiru!²⁹

Neither Vulcan, Phlogiston, nor Nibiru exist, and so if (25) is false on account of Vulcan's non-existence, we are left without a way to distinguish why (25) sounds so much better than (26), and why (27) seems like a good way of capturing one of "Vulcan"'s semantic features.

Similarly, in ordinary speech, if someone asks me what "unicorn" refers to, I would obviously respond by saying that it refers to unicorns. Whether or not unicorns exist seems to be totally beside the point.³⁰ The following seems to capture my willingness to respond that way:

²⁸Mark Sainsbury [2005] states reference clauses using universally quantified biconditionals. His reference clauses can be paraphrased in the following way: for all x , "Vulcan" refers to x if and only if x is identical to Vulcan. This allows for uniformity in our reference postulates. However, Sainsbury operates with a negative free logic, which makes all statements containing empty names false. This, in my opinion, is an intolerable result, for it makes sentences like "Sherlock Holmes is famous" false.

²⁹Nibiru is a planet that was supposed to collide with Earth at the end of the Mayan calendar in 2012, resulting in our planet's destruction. Thankfully, Nibiru does not exist.

³⁰Compare this point to Parsons [1979, 1980].

(28) “Unicorn” refers to unicorns.

In contrastive cases, the intuition is even stronger:

(29) “Unicorn” refers to unicorns, not to flying horses generally!

(30) “Sherlock Holmes” refers to a famous literary detective, not to a Tolkien character!

Thus it is plausible to think that statements of semantic reference made true by a causal chain inherit the NONEXISTENCE of the act of speaker’s reference that originated the causal chain.

Similar arguments can be made to show that semantic reference is opaque. Suppose that a Babylonian sees a star in the evening and baptizes it “Hesperus”. This is an act of speaker’s reference, and speaker’s reference, as we argued above, exhibits OPACITY. This means that the Babylonian can refer to Hesperus while not referring to Phosphorus or Venus, from which it seems to follow that (31) has a false reading:

(31) The Babylonians used “Hesperus” to refer to Phosphorus.

But if we trace the semantic features of “Hesperus” back to its original uses in acts of Babylonian speaker’s reference, there is nothing to prevent those features from being preserved in our statements of semantic reference, and neglecting them completely seems to be an oversight. Speaker’s reference is ambiguous between a *de re* reading and a *de dicto* reading, and we lose nothing if we treat semantic reference as inheriting both of these readings. Rather, they make available a more general notion of semantic reference. If we see semantic reference as anchored in an act of speaker’s reference, then (32) will have both a true reading and a false reading:

(32) “Hesperus” refers to Phosphorus.

The true reading is inherited from the transparent, *de re* form of the Babylonian’s acts of speaker’s reference, while the false reading is inherited from their opaque, *de dicto* reference.

Inheritance of these two readings allows us to satisfy several important constraints on

a theory of semantic reference. First, it is often considered important that statements of semantic reference are obvious, or even *a priori*. This is important because this is the only way they can appropriately figure into a theory of meaning, into psychological explanations, and into explanations of communication. To see this, consider (33):

(33) “Hesperus” refers to Hesperus.

If our theory of meaning is fully extensional, then (33) will entail (32). But (32) obviously does not state a fact about the meaning of “Hesperus”. Thus our semantic theory will have consequences that are themselves not statements of meaning, and that no reflection on our knowledge of meaning could ever reveal. Further, suppose that we try to explain an act of successful communication with the sentence “Hesperus is bright”. Suppose that John utters the sentence to Bill. Bill looks up in the sky and sees Hesperus, comes to agree with John, then forms the belief that Hesperus is bright. One part of our explanation for how Bill came to have that belief is that John uttered the word “Hesperus”, and “Hesperus” refers to Hesperus. But were we to state our explanation by saying that “Hesperus” refers to Phosphorus, our explanation would be a bad one. It would not explain how Bill came to have his belief on the basis of John’s linguistic act. Similar things can be said concerning successful communicative interactions with empty terms. This shows that explanations of communication are intensional, and so if reference is to play a role in a theory of communication, reference must be intensional also. Allowing statements of semantic reference to have a *de dicto* reading accomplishes both of these tasks: it allows us to specify the meanings of expressions in a way that is independent of such metaphysical facts, and at a degree of granularity that is appropriate for a theory of meaning and a theory of communication. This seems like the right result, since “means” itself is an intensional transitive verb, and the approach promises to unify our theory of meaning, theory of communication, and theory of reference through a simple mechanism.³¹

³¹It is instructive to point out the connections between this argument and the arguments given by Chomsky [1995, 2000]. His basic claim is that what we are referring to depends in intricate ways on our intentions, goals, interests, and other aspects of our psychology, not on a pairing with an external object. This is just another way of saying that reference should cohere with the rest of our psychological and communicative lives.

1.6 Technical Terms and Ordinary English

The previous sections argued that in English, semantic verbs are intensional transitives, and gave metasemantic arguments showing that the technical notion of semantic reference does, or at least should, inherit their intensionality. But this does not establish that we must make use of the English terms in theorizing about the meanings of English expressions. You might think that even if semantic verbs are intensional in English, we should still do our linguistic theorizing with totally extensional vocabulary. It is surely the prerogative of the semanticist, you might argue, to define technical terms as she sees fit, and to stipulate that “refers to” and “is true of” hold only between linguistic expressions and specific, existing objects or collections of objects.

In general, I agree that theorists are at liberty to define their terms how they see fit. However, if semantics is going to define “refers to” and “true of” so that they do not resemble their natural language counterparts, semanticists need to have a reason for this divergence. Perhaps these technical definitions are more fruitful than employing ordinary intensional language, or the intensional language is not clear enough to be suitable for theorizing. But there can be no such reasons, because traditional, extensional semantic postulates are just one reading of the ambiguous, intensional semantic postulates. Traditional, word-world connections are stated by the *de re* readings of constructions that are systematically ambiguous between *de dicto* and *de re* construals. Thus, semantic postulates stated in English *subsume* the technical reading of those postulates *as a special case*. The English words are simply more flexible, and more general. Accordingly, while we lose some univocality by stating our theory with intensional semantic vocabulary, the flexibility gained allows us to recapture the traditional notions of reference and truth-of, while also allowing for a pair of new notions corresponding to the *de dicto* readings of our semantic postulates.³²

³²Further, if we treat our semantic vocabulary as technical, and divorced from ordinary usage, then we are forced to posit a lexical ambiguity between the colloquial and technical uses of our semantic verbs. But if, instead, we use state our semantic theory with the terms as they are used in English—as ambiguous between *de dicto* and *de re* readings—it keeps our semantic verbs lexically univocal. What would have been two separate senses of our semantic verb phrases emerge as merely two different readings of our lexical semantic postulates, and are no more ambiguous than “seeks”. Further, if we think, as is plausible, that the *de re/de dicto* ambiguity is to be captured in terms of scope, then we can hold that the difference between the colloquial and technical senses is a structural ambiguity. Accepting such a structural ambiguity seems much more palatable than holding that our concepts of reference and truth are ambiguous between colloquial and technical senses. Consider a comparison. Imagine that we are proposing to give a theory of action that

1.7 Consequences

On my view, our semantic vocabulary is structurally ambiguous between two readings: a *de re* reading and *de dicto* reading. These two readings yield two ways of doing semantics. On the one hand, we can do semantics by specifying semantic significance using the *de dicto* readings of our semantic locutions: we can do semantics *de dicto*. Or, on the other hand, we can do semantics by reading our semantic locutions *de re*. Stating our semantic theory with the *de dicto* readings of our semantic postulates allows our semantic theory to serve as a theory of meaning or semantic competence in exactly the way envisioned by those working in the Chomskian tradition. This is made possible because semantics *de dicto* provides us with a novel form of semantic evaluation that overcomes several problems faced by views that treat semantic evaluation as purely extensional. First, the *de dicto* readings of our semantic verbs are not existence-entailing, so they can provide distinctive semantic values for empty NPs, including empty names, and they also allow for a form of non-specific reference.³³ Second, since ITVs are hyperintensional within their complements, they can assign semantic values to expressions that are much more fine-grained than ordinary extensions; semantics *de dicto* is able to assign hyperintensional semantic values.

The semantic values that semantics *de dicto* assigns to empty NPs will be of the same type as the semantic values of the complements of intensional transitive verbs generally, when they are read *de dicto*. The exact nature of these semantic values will depend on what the best semantics for ITVs turns out to be. If we were to adopt Montague's view that the

makes use of the notions of belief and desire. Clearly, "believe" and "desire" are intensional verbs; in fact, they are paradigmatically intensional. Noun phrases in their clausal complements can be interpreted either inside or outside the scope of the verb, yielding *de dicto* and *de re* construals of the beliefs and desires, and they resist substitution within their complements. But now suppose that we insisted on stating our theory with the extensionalized, technical, *de re* readings of "believes" and "desires". This would save us the trouble of having to come up with a theory that captures the intensionality of these verbs; and correspondingly, we might think that the theory is clearer, because we know quite well what it is to believe something about a particular object, as opposed to bearing a relation to some kind of finer-grained intermediary. But on the other hand, using only this vocabulary would drastically distort, and severely cripple, our proposed theory of action. It seems like many of our actions can only be adequately explained by the finer-grained, intensional readings of belief ascriptions, and extensionalizing merely limits the theory's expressive resources, likely making the theory empirically inadequate. If extensionalizing has these consequences for a theory of action, why should we extensionalize the vocabulary with which we state a theory of meaning?

³³Of course, there are already some views on the correct semantic values for empty names. For instance, Kripke [1973] and van Inwagen [1977], hold that many empty names, particularly fictional names refer to fictional characters instead of ordinary objects, and that fictional characters exist. But this view of empty names commits these theorists to drastically unintuitive claims like: "Sherlock Holmes exists" and "Vulcan exists". For other creationist views of fiction, see [Salmon, 1998], [Searle, 1979], and [Thomasson, 1999].

semantic value of an ITV's complement, when read *de dicto*, is an intensional quantifier, then the *de dicto* readings of “refers to” and “applies to” would assign intensional quantifiers as the semantic values of names and predicates. Alternatively, with Zimmermann [1993, 2001, 2006], we might hold that the *de dicto* readings of ITVs specify relations to properties, in which case all names in our language would have semantic values of predicative type. This would pair well with the view, advocated by Fara [2015], that names are predicates. Or we might even hold that the *de dicto* reading of an ITV is non-relational, treating such complements as adverbial modifiers, as in Forbes [2006, ch. 5]. On this latter view, and any view which holds that the *de dicto* reading of an ITV should be understood nonrelationally, the semantic values of names will not serve as ordinary arguments of the verb “refers”. Rather, they will serve to modify either an underlying event or state, or to form a complex predicate. Whatever semantics for ITVs turns out to be correct, it will have to account for the intensional behavior of the object position on the *de dicto* reading, and so will assign a semantic value to that position that accounts for the three features mentioned above.

In providing semantic values for empty names, semantics *de dicto* helps us make progress on the problem of empty names; it shows us how empty expressions can be meaningful, and make non-trivial contributions to the meanings of sentences in which they figure. And moreover, the fact that the *de re* readings of our lexical postulates for empty expressions are false allows us to retain a sense in which these expressions are genuinely empty. Many views that provide semantic values for empty names fail to fulfill this *desideratum*. Further, since the *de dicto* reading of an ITV is hyperintensional, the semantic value assigned to the expressions in its object position will be extremely fine-grained. This allows semantics *de dicto* to overcome problems of insufficient granularity that have historically kept truth-conditional semantics from issuing in interpretive theorems.³⁴

The ability to accomplish these related tasks is part of what make semantics *de dicto* such a good candidate for playing the role of a Chomskian theory of semantics. But there

³⁴The problem of truth-conditional theories not being sufficiently fine-grained to issue in interpretive theorems has come to be called the Foster problem. It was presented as a problem for Davidson's truth-theoretic approach to semantics [Davidson, 1967a] by John Foster [1976], and was recapitulated in the intensional setting by Scott Soames [1989]. Both Foster and Soames's arguments rest on an extensional construal of the premises from which T-sentences are derived. Soames's argument in particular depends explicitly on premises which involve predicates that are necessarily satisfied by the same objects. If these premises are hyperintensional, it blocks Soames's derivations.

are several further reasons. First, many of the reasons that Chomsky gives for rejecting relational reference and application are that these semantic notions behave intensionally as opposed to extensionally. For instance, Chomsky [1995, p. 21] offers the example of “al-Quds” and “Jerusalem”, which are both supposed to be names for the same city: Jerusalem. The force of the example comes from the intelligibility of the proposal to move al-Quds to a site north of Jerusalem. In the example, these names co-refer on a *de re* construal of reference, but on the *de dicto* construal of reference, they refer to different things, which allows us to talk about moving one without moving the other. Insofar as Chomsky’s point is that we can often refer to an object under one name but not under another, or move between relational and non-relational senses of reference, my proposal captures Chomsky’s point perfectly. Similar claims can be made about many of Chomsky’s other examples. Chomsky often points to the instability and abstractness of referents as evidence for the non-relationality of reference. For instance, London might be reduced to dust and be built in another place, but we can refer to it all along. If we were referring to the concrete object in the first place, it would pose a serious puzzle for how the referent of “London” could move from a concrete object, to an abstract object, and then back to a concrete one. A non-relational view of reference solves this puzzle. Lastly, Chomsky [2000, p. 178] often challenges advocates of the Lewisian view of semantics to give an account of the reference of expressions like “Joe Sixpack”, which seem to refer non-specifically. Semantics *de dicto* handles this case straightforwardly, because it countenances a form of non-specific reference that results from the non-specific reading of an ITV.

A final reason to think that intensionality is the right way of capturing Chomsky’s view of semantics is that the truth of many intensional constructions, particularly reports of searches, desires for, and beliefs in, depend crucially on facts about the intentions, beliefs, and interests of the subject, and this is exactly what Chomsky claims is the case for reference and application. Consider the case where London is destroyed and rebuilt in another place. According to Chomsky, the conditions under which the rebuilt city is considered London are determined by both psychological and social factors—they are not to be accounted for metaphysically. Insofar as Chomsky thinks that fine-grained psychological factors play a role in what it is to which words refer and apply, he seems to be saying that reference and

application are intensional, in that they depend on facts about the subject's psychology, or are mind-dependent in some way. Just as "John seeks a dog" says something about John's psychology—i.e. John's intentions, goals, beliefs, and desires—claims about reference may partly concern psychological and social facts, and do not report relations to particular objects.

Chapter 2

Hallucination and the New Problem of Empty Names

2.1 Introduction

The problem of hallucination is relatively simple: we often have perceptual experiences of things that don't exist.¹ A man in a desert, due to a host of physiological and environmental factors, may have a perceptual experience of an oasis, when in fact all that is in front of him is an expanse of hot sand. The man's experience still represents the world as being a certain way—as containing an oasis—and yet there is no oasis that the man perceives. The content of his perceptual experience cannot be dependent on an oasis, for by hypothesis there is no such thing, and so we need to find an alternative account of how his perceptual experience comes to represent what it does. Once we have such an account, we feel the real force of the puzzle: why shouldn't we generalize the account to all perceptual experiences, so that no perceptual experiences depend on the objects perceived? This problem, together with its cousin the problem of illusion, has delimited the space of possible views in the philosophy of perception: the major views are individuated in terms of how they respond to the puzzle.

¹Some philosophers may object to this phrasing, claiming that we can't have perceptual experiences of things that don't exist. This objection is likely based on their commitment to a regimentation of English where a "perceptual experience of" and other locutions for reporting intentional states are relational, and existence-entailing. I am making use of ordinary English, and in ordinary English this is a perfectly acceptable and common thing to say. The locution "to perceptually experience" is one of English's many forms perceptual ascription that has an intensional reading. Much more will be said about this below.

The central claim of this paper is that the problem of hallucination is identical to a problem in the philosophy of language: the problem of empty singular reference. We can state the problem of empty singular reference roughly as follows: what is the semantic value of an empty singular referring expression? This is, on its face, quite similar to the question of: what is the content of a hallucinatory perceptual experience? But despite this similarity, the problems have been addressed within separate philosophical subdisciplines, and while both have developed sizable bodies of literature, these literatures have remained largely disconnected.² The problems are rarely discussed together, and no one, at least for the better part of a century, has come close to explicitly identifying them.³

My goal in this paper is to show that this is a serious mistake; much theoretical progress can be made through the recognition that the two puzzles are identical. The problem of hallucination and the problem of empty singular reference are derived from the same two general principles concerning representation, and have the very same conclusion. The problems follow from the principles of SIGNIFICANCE and UNIFORMITY. The principle of significance is the principle that empty representations are still contentful, and their content is not trivial. The principle of UNIFORMITY is the principle that our account of the content of empty and non-empty representations should be uniform, from which it follows that empty and non-empty representations have the same sort of content. These two principles push us toward the same conclusion in both the case of perceptual experience and the case of reference: representational content is not dependent on (existent) objects. Insofar as externalism is a thesis about the object-dependent nature of content, the conclusion of this

²As an illustration of how the problems are treated separately, a recent volume on hallucination: *Hallucination: Philosophy and Psychology* [Macpherson and Platchias, 2013] does not mention empty names even once. Similarly, but in the opposite direction, an important collection on empty names, *Empty Names, Fiction, and the Puzzles of Non-Existence* [Everett and Hofweber, 2000], does not mention hallucination one in any way that connects to empty names.

³The problems are often treated as faintly analogous, and many philosophers may already have the feeling that the problems are related. If you do have such a feeling, the virtue of this paper will lie in articulating, and making explicit, that feeling, and then drawing out the consequences. The closest thing I have found to an identification of the problems of hallucination and empty names is made by Russell [1951, §XII]. However, much still has to be inferred from what he says the basis of the fact that he treats the contents of perception and the contents of (not logically proper) names as descriptive. See also [Russell, 1921, §X]. The connection between language and perception in Locke and Berkeley is discussed by Ian Hacking [1975, Ch.s 3-5], although the parallel between hallucinations and empty names is not drawn explicitly. In addition, Gareth Evans [1982, Ch. 1] draws parallels between certain problematic accounts of perception and problematic accounts of Fregean senses, from which the analogy between hallucinations and empty names can be extrapolated.

argument can be seen as a form of internalism.

However, it is important to note that I am *not* endorsing this argument, nor am I endorsing any form of internalism. Rather, I am demonstrating that one argument underlies both problems, and that this identity has deep philosophical consequences. If the problems are identical in the way I propose, we should expect there to be a clear mapping between their spaces of possible responses, and this is exactly what we find: each response to the problem of hallucination has a corresponding response to the puzzle posed by empty referring expressions, and so each view of perceptual content has a semantic counterpart. However, while many of these semantic views have already been developed in the philosophy of language, others have not. One notable option that has been overlooked is what I call *semantic adverbialism*, the semantic counterpart of perceptual adverbialism. I discuss semantic adverbialism, and propose a way of elaborating and defending semantic adverbialism that draws on previous work. The route I propose is to treat semantic verbs such as “refers to”, “applies to”, and “is true of” as intensional transitive verbs. Two of the most fully developed approaches to the semantics of intensional transitive verbs provide them with a non-relational semantics, and bringing these views to bear on our semantic verbs allows us to develop a non-relational theory of the content of empty names, and of semantic content more generally.

I will end the paper by showing that even if one is not inclined to accept semantic adverbialism, the Anscombian approach that I adopted above in developing it is theoretically enriching. In response to the problem of hallucination, Anscombe [1965] showed that verbs of perception and sensation are intensional transitive verbs, and claimed that the correct response to the problem of hallucination required determining the correct semantics for the object positions of such verbs. Given the identity between the problem of hallucination and the problem of empty singular reference, we can, and should, treat semantic verbs in the same way: “refers”, just like “sees”, “senses”, and “hallucinates”, is an intensional transitive verb, and the semantic value of empty referring expressions, just like the intentional object of a hallucination, will be given by determining the semantic contribution of the NP in the object position of intensional transitive verbs.⁴ As a consequence, each pair of responses

⁴I have argued for the claim that “refers to” and “applies to” are intensional transitive verbs at length

to the problem of hallucination and the problem of empty names can be seen as resulting from a particular view on the semantics of intensional transitive verbs. Thus, our account of the semantics of intensional transitive verbs can inform our theories of both linguistic and perceptual intentionality.

2.2 Intentionality

2.2.1 The Problem of Non-Existence

The problem of hallucination and the problem of empty singular reference are both versions of what is perhaps the central puzzle of intentionality: the Problem of Non-Existence. The problem can be stated as follows: how do we manage to represent, talk about, think about, *etc.* things that do not exist? There are two general approaches to accounting for intentionality—the relational approach and the non-relational approach—and each of them yields a different family of answers to this central problem.⁵ The relational approach accounts for intentionality in terms of intentional objects: mental states and other representations have their intentional features in virtue of a relation to an object, and this object is what they are about or represent. On the non-relational approach, representations can exhibit intentionality without being related to an object which they represent; they have their intentional features intrinsically.

The relational approach has been around as long as the puzzle of nonexistence itself—at least since the middle ages. Many medieval philosophers accounted for aboutness in terms of intentional objects: our thoughts have intentional objects, which are what our thoughts are about, or represent [see Pasnau, 1997, Perler, 2001a,b, Priest and Read, 2004]. On a

elsewhere, on primarily empirical grounds.

⁵In the philosophy of mind it is common to view intentionality as primarily, or solely, propositional, and to make “intentional content” more or less synonymous with “propositional content”. With Tim Crane [2014b], I think that this is a serious mistake. I see no reason why intentionality—or representation or aboutness or directedness on an object—should be ascribed primarily to states expressed by propositional attitudes as opposed to states expressed by intensional transitive verbs, such as “seeks”, “wants”, “fears”, “hopes”, and “needs”, when these verbs take NP complements. These states seem like paradigmatic instances of intentional mental states, insofar as they all have objects toward which they are directed. In what follows I will primarily be considering non-propositional forms of intentionality, that are expressed with intensional transitive verbs such as “represents” or “is about”. This is due to the fact that the problem of empty names and the problem of hallucination, as I will reconstruct them, concern the intentional content of names, or perhaps noun phrases, rather than the contents of whole sentences. But much of what I say about relational vs. non-relational views of intentionality carries over straightforwardly to the propositional case.

relational approach to intentionality, the main project is figuring out the nature of these intentional objects—do they exist? are they concrete or abstract? what type of objects are they [Crane, 2006]?⁶ According to some traditional views, since many of our thoughts and representations are about things that don't exist, and intentional objects are what these representations are about, some intentional objects will fail to exist. These days this view is typically attributed to Meinong, who sums up what he takes to be the objectual approach nicely:

That knowing is impossible without something being known, and more generally, that judgments and ideas or presentations (*Vorstellungen*) are impossible without being judgments about and presentations of something, is revealed to be self-evident by a quite elementary examination of these experiences. [Meinong, 1904]

Meinong holds that every representation represents something, and holds that “something” ranges over objects. According to Meinong, the fact that there are such objects is what validates inferences like the following:

- (1) I imagined a unicorn.
Therefore I imagined something.

- (2) John searched for a golden mountain.
Therefore John searched for something.

Given an objectual treatment of the object-position quantifier, it follows from these inferences that some things do not exist.⁷ There are also relational approaches to the problem of nonexistence that are non-Meinongian. One such approach is to claim the intentional objects of our thoughts do exist, but they are just different than we thought they were: we might, for instance, treat Sherlock Holmes as an existent abstract object—a cultural artifact [see Kripke, 1973, van Inwagen, 1977, Salmon, 1998, Soames, 2005]. This route provides a representation with an intentional object at the cost of denying that there is a

⁶The modern rehabilitation of the theory of intentional objects comes from Brentano [1973, Ch. 1].

⁷I will use the label “Meinongianism” for the position that there are objects that don't exist. See [Parsons, 1979, 1980], [Priest, 2000], Priest and Read [2004], [Routley, 1983], [Zalta, 1988].

genuine problem of nonexistence. A third option is to revert to what some claim is the original interpretation of Brentano's claims, holding that intentional objects have intentional *inexistence*, and so are somehow immanent to the mind [see Brentano, 1973, Crane, 2006, Perler, 2001b].⁸ Within a framework where we explain representation in terms of intentional objects, the problems of hallucination and empty singular reference can be stated together simply: what are the intentional objects of hallucinations and empty singular referring expressions?⁹

The other approach to problems of intentionality treats intentionality as non-relational. On this view, a representation can represent something, or be about something, without there being some object (or any other entity—of any type) which it is a representation *of*. That is to say, on this view, representation, or aboutness, is not a relation [Goodman, 1976, Crane, 1998, 2006, 2009, 2012]. Recall Meinong's quote from above, and the inferences that followed. According to the non-relationalist, the inferences in (1) and (2) are good ones, and the claims in the quote from Meinong are true, but the non-relational view holds that we have no reason to treat "something" as an objectual quantifier, or as generalizing over an entity that serves as an argument to the verb. In fact, the non-relationalist points to the fact that there is much evidence suggesting that "something" is not an ordinary objectual quantifier, but a *special quantifier*, which quantifies over the semantic contribution of the NPs in the object positions of intensional transitive verbs, whatever that contribution turns out to be [Moltmann, 1997, 2003a, 2004, 2008, 2013, Zimmermann, 1993, 2001, 2006].

⁸It's important to note that we could also take intentional objects to be properties or collections of properties. When I say "intentional object", I mean to include any entity (in the broadest sense) that can serve as a relatum of the representation relation.

⁹Even though it seems obvious that names and hallucinations are representations, it is rare to hear the problem of empty singular reference posed in these terms. We don't often inquire about the intentional objects of names—we are more often inquire about their referents. I imagine that this is because stating the problems this way—in terms of intentional objects—poses a dilemma: if hallucinations are genuinely hallucinations, and empty names are genuinely empty, then it seems that we cannot provide them with intentional objects without countenancing objects that do not exist. The dilemma is this: either we make hallucinations into illusions and empty names non-empty, for instance, by positing that they actually relate us to existent but abstract objects, or we countenance non-existent objects. That is to say, we can't address the problems of empty names and empty singular reference within an intentional object framework without being Meinongians.

However, stating the problem in terms of intentional objects is still profitable, because it acknowledges the fact that empty names have distinctive representational features, and that one way of accounting for such features is to give a theory of intentional objects, some of which may fail to exist. Much the same is true for hallucinations: asking about the intentional objects of hallucinations acknowledges that hallucinations represent in non-trivial ways. Different hallucinations will represent different things, and so asking for the intentional objects of hallucinations is one way of accounting for their distinctive representational features.

More will be said about this below.

On the non-relational view of representation, what a representation is of is an intrinsic property of that representation. Consider an example: a picture can be of a unicorn, or represent a unicorn, without there being a unicorn that it represents, and it represents the unicorn in virtue of its own properties, not in virtue of some relation that it bears to that unicorn. The non-relational view takes the fact that unicorns don't exist very seriously: since unicorns don't exist, they can't be the relata of any relations, and so *a fortiori* they can't serve as intentional objects, the putative relata of the representation relation. Rather, on the non-relational view, the painting might represent *pictorially*, or *iconically*, but not in virtue of a relation to an object represented.

A prominent version of the non-relational approach to representation is descriptivism. In what has turned out to be a microcosm of the relational/non-relational dispute, Russell [1905] responded to Meinong's theory of objects [Meinong, 1904] by treating empty names as disguised definite descriptions. Russell analyzed these descriptions as contributing both quantificational and predicative material to the content of sentences containing them. Because descriptions serve to introduce a mix of quantificational and predicational material, they do not represent objects because of any relation that they bear to the object. Rather, they help to specify a condition that object may or may not meet, and they represent what they do independently of whether any object happens to satisfy that condition.¹⁰

Within the non-relational framework, the problem of hallucination and the problem of empty names manifest themselves differently. What is required is that we provide empty names and hallucinations with *content* that is not object-dependent: we need to supply names with semantic values other than their referents, and hallucinations with content other than the objects that they represent. Pairing different names with different descriptions is

¹⁰It seems extremely plausible that Russell's treatment of names as definite descriptions, and his analysis of definite descriptions as contributing quantificational and predicational material to the truth-conditions of a sentence, undergirds much of the tendency to treat intentionality as a primarily propositional phenomenon. According to Russell, definite descriptions are "incomplete symbols", and do not represent any object, or exhibit any form of intentionality, on their own, outside of the context of a sentence. Rather, it is the entire sentence that is representational: it represents the world as being one way or another. This move allowed Russell, and a host of philosophers following him, to treat the problem of non-existence merely an instance of falsehood: the problem of non-existence was just a case of incorrect propositional representation. There are, however, descriptivist views of names on which names are not incomplete symbols, and on these views, names are self-standing, but non-relational representations.

one way of providing empty distinctive representational contents, even in the absence of objects satisfying these descriptions. For the descriptivist about names, “Pegasus” and “Zeus” both represent, and do so distinctively, but not in virtue of bearing a relation to objects which they represent. They both have representational, in this case descriptive, content. In what follows, I will use the term “content” broadly, and as neutral between relational and non-relational approaches to intentionality, even though it is sometimes used more narrowly for a “way of representing an object”.¹¹ More specifically, below I will sometimes switch between saying that a representation or experience has content and that it “represents something”, “exhibits intentionality”, or “is about something”.

2.2.2 The Anscombian Approach

Anscombe took the biggest steps toward connecting the two puzzles, and realizing that the problem of hallucination and the problem of empty names were at bottom the same problem of content. Anscombe realized that the problems of hallucination and illusion are really problems concerning certain kinds of verbs: verbs of perception and sensation such as “see”, “hear”, “smell”, *etc.* Anscombe pointed out that these verbs have readings that are true even when the NPs in their object positions are empty, and that such cases are just reports of hallucinations.¹² Once this connection is made, traditional intentionality

¹¹“Intentional content” is used this way by one of the main non-relational views in the philosophy of perception: intentionalism. Inspired by Anscombe, intentionalists, most notably Tim Crane [1998, 2006, 2009, 2012, 2014a], hold that perceptual experience is fundamentally a form of representation. According to Crane, every representation has an intentional object: this object is what the representation is of, or what it is directed towards. But following Brentano [1973], intentionalists also distinguish between a representation’s intentional content and its intentional object. Intentional content, according to the intentionalist, is the way a representation represents its object, whereas its intentional object is what it represents. On its face, intentionalism, like the Anscombian view that inspired it, seems to be relational: it holds that every representation has an intentional object. This would mean that hallucinations pose a problem for the intentionalist, as they do for relational theories generally. But here the intentionalist falls back on two related points. First, again invoking Anscombe, the intentionalist uses “intentional object” in the old sense of “object”, as in “object of thought” or “object of desire”, to merely designate what it is toward which a state is directed. They then couple this notion of “intentional object” with the possibility of giving a reductive account of intentional objects: it may well turn out that what is really going on in cases of representation is to be explained intrinsically [Crane, 2012, pp. 133-135]. On such a reductive view, talk of intentional objects, and the relational conception of representation to which they seem tied, do not provide a fully perspicuous account of the nature of representation, in terms of surface form. Rather, the intentionalist allows that we might give an explanatory reduction of truths about the non-existent to truths about the existent, and this explanation will turn out to be non-relational.

¹²In modern terminology, Anscombe showed that verbs of sensation are intensional transitive verbs (ITVs). There is some debate, however, over whether “see” in fact has an intensional reading. There is a longstanding debate over whether, when Macbeth hallucinates that there is a dagger in front of him, whether he in fact sees a dagger. Thus the peculiarity of Macbeth’s question: “is this a dagger which I

puzzles like the problems of hallucination and illusion can be understood as problems about the semantic contribution of the noun phrases in the object-positions of these verbs. An example will be helpful here. On Anscombe's view:

(3) John hallucinated a dancing dragon

and:

(4) John saw a dancing dragon

are on a par with

(5) John searched for a dancing dragon.

Searches are a traditional instance of the puzzle of nonexistence, and so unsurprisingly, accounts of search are often given in terms of intentional objects. On such a view, what John is searching for is the intentional object of his search: a dancing dragon. Anscombe keeps this terminology when she moves to discussing both verbs of search and verbs of perception, claiming that "a dancing dragon" gives the *intentional object* of all three of the states reported above. According to Anscombe, the key to solving the problem of nonexistence is to give an account of the semantic contribution of the "a dancing dragon", as it occurs in the object positions of these verbs; an account of its contribution will serve as an account of the intentional objects of both searches and perceptual experiences.^{13,14}

see before me?". For a recent argument that perceptual ascriptions have intensional readings, see Bourget [2016].

¹³Although I don't wish to take a strong interpretive stance here, Anscombe seems to think that the contribution of the NP in the object position of an ITV to the truth-conditions of sentences like the ones above is an object under a description. And insofar as Anscombe thinks that this semantic approach is the right way to approach the problem of nonexistence, she seems to think that the intentional object of all intentional states will also be an object under a description. If this interpretation is correct, then Anscombe is still operating within the relational framework, and accordingly, her view faces the same problem that we mentioned above: if one's account of intentionality importantly involves an intentional object, then in cases where that intentional object fails to exist, Meinongianism seems to be the only option. Thus, I think Anscombe's account, unmodified, suffers from the same defect that all accounts that make use of intentional objects suffer from: they threaten to collapse into Meinongianism. However, Anscombe's view of the semantics of ITVs, as involving an object under a description, is not essential to her more general point.

¹⁴We might also be able to avoid a relational construal of Anscombe's remarks by giving the right account of an object under a description, or a so-called "*qua*-object". Some views of *qua*-objects, for instance, that in Szabó [2003a], do not treat *qua*-objects as objects at all. Other views, such as those in Landman [1989a,b] and Fine [2003] do.

The most important conclusion of Anscombe’s paper is the following: an account of hallucination, and perhaps of intentionality more generally, can be given by providing a semantics for intensional transitive verbs. And as Anscombe acknowledges, just as there are relational and non-relational approaches to intentionality, there are also relational and non-relational semantics for ITVs. It follows that each different view on the semantics of intensional transitive verbs will yield a different theory of intentionality.¹⁵ The details of how different views on the semantics of intensional transitives yield different accounts of intentionality will be discussed in §6. But once we realize that the problem of hallucination and the problem of empty names are identical, Anscombe’s insight concerning verbs of sensation can be straightforwardly extended to verbs that state the semantic features of words. Just as “sees” and “senses” can report hallucinations, “refers” can report the semantic features of empty names. As a consequence, the correct account of the semantics of empty names will likewise be given by providing a semantics for intensional transitive verbs. The next section will reconstruct the problems of hallucination and empty names in a way that exhibits their identity.^{16,17}

¹⁵Provided that we view our semantics in a realist way.

¹⁶Anscombe’s claims about verbs of sensation and perception were not completely without precedent. Richard Cartwright [1987], in his classic article on Macbeth’s Dagger, cites G.E. Moore as distinguishing between two senses of a perceptual verbs like “sees”. One sense is that on which it is true to say that Macbeth is seeing a dagger when he hallucinates; this is what Moore calls *directly seeing*. Similarly, Cartwright also cites Ayer as distinguishing between two senses of “sees something”. On one reading of “sees something”, Ayer claims, there need not be something that one sees. Both the idea that “sees” has two senses, and the idea that the quantifier in the object position need not export, can be seen as further confirmation of Anscombe’s idea that “see” has a reading that is true even when the object position NP does not pick out an existent object. Both of these properties are features of intensional transitive verbs.

¹⁷Gilbert Harman [1990], in the course of arguing for a somewhat different point, reiterates some of the claims made by Anscombe years earlier. Harman emphasizes that perceptual experiences such hallucinations are instances of the broader category of intentional states, which includes beliefs, desires, imaginings, and paintings, and many others. An account of the representational features of hallucinatory experience, he claims, will fall within the purview of a theory of intentional objects more generally. Remarking on Macbeth’s dagger, Harman claims:

If a logical theory can account for searches for things that do not, as it happens, exist, it can presumably also allow for a sense of “see” in which Macbeth can see something that does not really exist [Harman, 1990, p. 38]

Thus Harman sees theory of hallucinatory perceptual experience as located within a theory that tries to specify the intentional contents of representational states more generally. And similarly to Anscombe, Harman thinks that this theory will be intimately connected to the project of providing a semantics for both perceptual verbs like “sees”, and other intentional verbs like “seeks”, both of which are intensional transitive verbs.

2.3 Reconstructing the Arguments

This section will provide reconstructions of the problem of hallucination and the problem of empty names, within the framework outlined above for discussing problems of intentionality. Once we have reconstructed the arguments in this way, we will see that the two main premises in the arguments are identical: they are two general principles concerning representation.

2.3.1 The Argument from Hallucination

The traditional argument from hallucination is most often presented within the intentional-object framework. Consider the following version of the argument from hallucination, due to Jeff Speaks, which I take to be representative:¹⁸

1. In every experience there is an object of your awareness.
 2. In the case of a hallucination, the object of your awareness cannot be a material thing.
 3. The objects of awareness are the same in the case of hallucinatory and veridical experience.
- C1. Material things are never the objects of experience.
- C2. Every perceptual experience has something other than a material thing as its object.

The phrasing of the first premise hides two important theoretical choices. First, it commits us to the view that experiences are of *objects*, and second, that experiences consist in a relation to these objects: the relation of awareness.¹⁹ The most general form of the argument, as I will show, jettisons both of these assumptions.

¹⁸The original argument appears in Berkeley [1710]. Another version appears in Moore [1953], and is used to argue for a sense-data theory of perception. Other versions appear in Crane and French [2016] and BonJour [2016].

¹⁹However, if we construe “object” in “object of your awareness” in a suitably deflationary way—perhaps in the old way, as Anscombe does—then [1.] does not commit us to awareness being a relation. But in general, when we talk about “objects of thought” or “objects of awareness” or “intentional objects”, it is not enough to merely say that such objects do not make the state in question relational. Rather, like the intentionalist, we must tell a story about why their surface form is misleading, and why we are not to construe “object” in the modern sense.

First, we need not be committed to the idea that awareness is awareness of objects, or even that awareness is a relation. As we saw above, accounting for the representational features of experience in terms of objects is just one possible approach. The incontrovertible premise of the argument from hallucination is not that there are objects of which we are aware in experience; rather, it is that in every experience, we are aware of *something*, or, phrased another way, every experience is about something, just Meinong claimed in the passage quoted above. Further, this something is distinctive: we can ask, of every experience, “what are you aware of?” And “something” generalizes over whatever it is that our answer provides. But it is a further step to say that “something” is an objectual quantifier, which generalizes over an object that serves as an argument for the verb. In fact, there is evidence that such quantifiers are not objectual: the quantifier “something” is a *special quantifier*, and there is a significant debate about what these quantifiers range over, or whether they are substitutional. Few linguists think that they behave in the same way as ordinary objectual quantifiers like “some”. Consider an example: in a hallucination, I might be aware of a vulture circling the room. It follows that I am aware of something. But it does not follow that there is a vulture, or that “a vulture” contributes an object to the report of my hallucination, and so it does not follow that “something” is an objectual quantifier. This leaves us no reason to think that the semantics for special quantifiers should be objectual, and accordingly, we have no reason to state the first premise in terms of objects, rather than in a way that remains neutral on the semantics for special quantifiers.

Second, the argument is cast in terms of awareness, but the idea that in every experience we are aware of something is meant to capture the idea that every experience of ours is *of* something, has a subject-matter, or represents something.²⁰ All three of these locutions are simply attempts to get at the idea that the experience exhibits intentionality; that is to say, “object of awareness”, as it is used in the argument, is basically synonymous with “intentional object”. But as we saw in the last paragraph, intentionality need not be explained within an intentional-object framework: it need not be construed as a relation,

²⁰In saying that every experience represents something, I intend to use “represents” in a pretheoretical sense. I don’t intend to commit myself to representationalism, or a representational theory of mind, or any particular view of the mind at all, other than that mental states have intentional features. If “represents” has come to be laden with too many theoretical commitments, feel free to substitute “exhibits intentionality” or “is of something”.

and the same goes for awareness. In what follows, if an experience or a representation exhibits intentionality, I will say that it has *content*, but I use the term “content” in a way that is neutral between relational and non-relational approaches to intentionality discussed above.

With these modifications in hand, we can go back and revise the argument to make it more general, freeing it from these theoretical assumptions. I have kept the terminology of “awareness”, but the argument could just as well have made use of the notions of representation or aboutness.

- 1.' In every experience, we are aware of something.
- 2.' In the case of a hallucination, what we are aware of is not a material thing.
- 3.' What we are aware of in veridical experiences is the same (in kind) as what we are aware of in a hallucinatory experience.
- C1. We are never aware of material objects.
- C2. We are always aware of something other than a material object.

Premise [1.'] modifies [1.] by using “something”, a special quantifier, and remaining neutral on its semantics. Relational theorists are free to specify the premise by treating “something” as an objectual quantifier, but it need not be treated that way. Further, in premise [1.'], “something” serves as a generalization over whatever semantic contribution is made by the noun phrase in the object position of “aware of”. If, for instance, I am aware of a dog, then “something” generalizes over the semantic value of “a dog”, as it occurs in the object position of the verb.²¹ This shows that “something” here generalizes over things of which I am aware, but not in the sense of objects. Rather, a non-relationalist might interpret such things as “ways in which I am aware”.

Premises [2.'] and [3.'] make use of the pronoun “what” as a pronoun referring to the thing of which you are aware. In these premises, “what” is being used as a *special pronoun*.

²¹On many accounts, the semantic value of “a dog” is seen as a second-order property: a generalized quantifier. But not every view of the semantics of special quantification sees “something” as generalizing over a generalized quantifier, although Zimmermann [2001] does. Moltmann [2003b, 2008, 2013] treats such quantifiers as quantifiers over the entities denoted by a certain kinds of nominalizations, which she thinks are tropes.

Similarly to special quantifiers like “something”, its values may not be objects: instead, its values are whatever contributions are made by intensional NPs that ordinarily appear in the object-positions of intensional verbs [Moltmann, 1997, 2003a, 2004, 2008, 2013]. In general, special quantifiers serve as the antecedents for special pronouns. The key point to take from this is that “something” and “what” can range over distinctive non-objectual specifications of what we are aware of in experience. They allow for a non-relational way of formulating the argument, while still capturing the way in being *of* a particular object allows it to have distinctive representational features.

With the argument recast in this form, we can now go on to consider the arguments in favor of its premises. The main premises are (1) and (3), while premise (2) is a definitional premise. In what follows I’ll adopt the following terminology: I’ll call the first premise the principle of INTENTIONAL SIGNIFICANCE, and I’ll call the third premise the principle of EXPERIENTIAL UNIFORMITY. As we will see below, EXPERIENTIAL UNIFORMITY actually has two versions. On the one hand, it can be the claim that a single experience can have the same content as a counterpart that differs only in that it is empty, or it can be the claim that perceptual experiences generally have the same kind of content. The latter version is, I believe, stronger, and implies the former, but it is correspondingly harder to establish. These two different forms of uniformity will become important below when we discuss arguments in favor of EXPERIENTIAL UNIFORMITY.

Arguments for Intentional Significance

Defending INTENTIONAL SIGNIFICANCE requires us to demonstrate that all of our experiences, including the hallucinatory ones, represent something, or have content. Below I will offer two arguments that hallucinatory perceptual experiences have content: the psychological argument and the substitution argument.

The psychological argument begins by considering the role that hallucinations play in our psychology: their cognitive role. The cognitive role of a representation is its causal role; mental representations, including perceptual experiences, interact causally with each other and other aspects of our psychology, and capturing these causal interactions is an important way in which a psychological theory helps to explain behavior [Segal, 2000]. A

representation's cognitive role is the way it figures into this network of causal connections. Hallucinations play an important cognitive role, and thus will figure into the explanations of behavior given by a psychological theory. Most importantly, they can only fulfill this explanatory roll if we attribute content to them.

Consider our example from the introduction: a man in a desert, due to his dehydration, exhaustion, and the heat, hallucinates an oasis. This perceptual experience is important to understanding his behavior: it causes him to run, screaming and overjoyed, toward a particular spot in the desert where he perceives there to be an oasis. The man's hallucination plays an important role in explaining why he acted as he did: the direction and exuberance of his run are only explicable because of *what it was* that he hallucinated, and *where* he thought it was located. Moreover, different hallucinations will cause different kinds of behaviors. If our traveler lost in the desert had hallucinated a pack of coyotes instead of an oasis, he would likely have run away from where he perceived them to be, rather than toward. This scenario shows that hallucinations are contentful; the man is overjoyed because his hallucination has content: it is *of* an oasis.²²

The second argument that hallucinations have content comes from substitutions of NPs within reports of hallucinations. Suppose that John is having a hallucination of a giant albatross. We can report his hallucination with (6):

(6) John is having a hallucination of a giant albatross.

But now consider a different hallucination that John might have had, one of a dancing dragon, which we could report as in (7):

(7) John is having a hallucination of a dancing dragon.

On the face of it, John's mental states—his hallucinations—represent the world as being two very different ways. The first represents the world as containing a very large bird while the second represents the world as containing a scaled, possibly fire-breathing creature

²²Perceptions are sometimes not seen as exhibiting intentionality until further downstream in the cognitive process: they are sometimes seen as "raw feels". On this view of content, insofar as the representation's content is connected to its explanatory role in a psychological theory, even a "raw feel" will have content. Thus I'm using the notion of content broadly.

that has momentarily taken a break from whatever dragons normally do to dance. But if hallucinations lack content, then the object-position of these sentences will not be able to distinctively characterize John’s mental state, for his mental state is not contentful. His hallucinations would not be *of* anything. But this is problematic for two reasons. First, the sentences above seem true in the relevant circumstances, and they explicitly state what his hallucination is *of*. This indicates that hallucinations exhibit intentionality. Second, if hallucinations were not contentful, we should be able to substitute “a giant albatross” for “a dancing dragon” without a change of truth-value.²³ But in the circumstance where John is ducking and screaming “albatross”, it would be absurd to say that he is hallucinating a dancing dragon. Any plausible account of the contents of perceptual experience is going to have to account for the fact that these two hallucinations are of different things. This is the substitution argument in favor of the INTENTIONAL SIGNIFICANCE of hallucinations.

Another way of making this argument is by appealing to a specific account of representational content. Many philosophers of perception take the contents of perception to be *accuracy conditions* [Siegel, 2010a,b]. If we treat representational features as accuracy conditions, it seems impossible to deny that the two hallucinations represent in non-trivial ways: what is required of the world for my perceptual experience of an albatross to be accurate is very different than what is required of the world for my perceptual experience of a dancing dragon to be accurate. Thus, taking on board the most widely accepted account of perceptual content leads us directly to the truth of INTENTIONAL SIGNIFICANCE.^{24,25}

²³On the assumption, that is, that expression content is compositional. See Szabó [2008] for a discussion.

²⁴There are, however, several ways of denying that experiences have accuracy conditions. The first is to be an Adverbialist, *à la* Chisholm [1956]. Another is to hold, with Charles Travis [2004], that perception requires “taking” in order to be assessable for accuracy: accuracy is a notion that only comes in further down the cognitive stream. Or, we might, with Ned Block [1990, 1996], hold that perceptual experience yields only raw feels.

²⁵Accuracy conditions are typically captured propositionally, and so they lead to views on which perceptual content is always propositional. I think such views are mistaken, for the following reason: if we think, with Anscombe, that perceptual experiences are often reported using ITVs, and that their intentional features can be captured by providing understanding the semantic contribution of the NP in the direct object position of these reports, then a propositional view of perceptual content would require that these reports in fact be covertly propositional. Such a view would require the reduction of intensional transitive verbs to propositional attitude verbs. While many people defend this view [Quine, 1956, Larson et al., 1997, Montague, 1974c], I do not think it is plausible as a general view of the semantics of intensional transitive verbs. And more generally, I reject the idea underlies this sort of reduction: that all intentionality must be propositional: I can think *of* something without thinking *that* anything is the case. See Crane [2014b] and Szabó [2003b] for arguments to this effect.

Arguments for EC-Uniformity

The third premise of the argument from hallucination is what I have called the principle of EXPERIENTIAL UNIFORMITY. The idea underlying this principle is that there is a kind of content that is uniform across empty and non-empty representations. But as we saw earlier, uniformity comes in two varieties. On the one hand, we can consider the content of a single, non-empty representation and then compare it to an empty counterfactual counterpart. Or we can consider the entire category of mental representations, and claim that within this category, veridical and hallucinatory experiences should be supplied with contents of the same kind. Below I will go through several arguments in favor of both forms of uniformity.

When we consider a single experience, arguments for uniformity make use of a Twin Earth thought experiment. The arguments I give here are adaptations of the arguments given in Segal [2000] for the claim that empty natural kind terms have content that is not object-dependent.²⁶ Consider our situation from the introduction where our traveler, John, is hallucinating an oasis. Now consider Twin Earth, and his counterpart in the desert on Twin Earth, Twin John, who is molecule for molecule identical to John. Suppose further that the only difference between Earth and Twin Earth is that on Twin Earth, there is actually an oasis in the desert that John perceives. I take the situation on Twin Earth to be one that could very well arise. In defending INTENTIONAL SIGNIFICANCE above, I argued that hallucinations have content. What follows will be three arguments that the content of John's perceptual experience is the same as the content of Twin John's perceptual experience.

[EU1] The first is based on the notion of a cognitive role discussed above. Consider the causal role played by John's hallucinatory perceptual experience on Earth, P_1 , and compare this causal role to that of its veridical counterpart on Twin Earth, P_2 . P_1 and P_2 seem to interact causally with other aspects of John and Twin John's psychologies in identical ways.

²⁶Segal's arguments are part of a sustained argument for internalism about mental content for kind terms. As I mentioned above, I am not endorsing internalism here, since I am merely outlining the problem of hallucination, and cataloguing what I take to be the best arguments for its premises. However, the argument that underlies both the problem of hallucination and the problem of empty names can be seen as an argument for internalism. In §5, I construct a master argument for the conclusion that the content of a representation is not object-dependent. Segal's Twin Earth arguments can be seen as an application of this argument to the case of empty kind terms.

For instance, in both of the cases, the perceptual experience causes its subject to run in a particular direction, and to be excited at the prospect of quenching its thirst. That is to say, the hallucinations figure into identical psychological explanations, and in order to adequately play a role in a psychological explanation, we must attribute content to the hallucinations. If the attribution of content is guided by a mental state's causal role, then the two representations, P_1 and P_2 , will share whatever content is relevant to psychological explanations. Of course it will be true that one hallucination is veridical but one is not, and so *in one sense*, one is of something that the other is not, but in another sense they are of the same thing: an oasis. Thus they will still share an important form of content: namely, the kind of content that is relevant to psychological explanations.

[EU2] The second argument concerns the supervenience base of the content of a hallucination. Consider again John and Twin John. They are molecule for molecule identical, but John is hallucinating while Twin John is having a veridical perception. Consider the properties on which the content of John's hallucination supervenes. These properties form the supervenience base of the content of John's representation. This might include any number of his physical or psychological properties, together with relations to his environment (the heat that has caused his exhaustion, etc). But importantly, in John's case, there simply is no oasis. Given that the only difference between the Earth and Twin Earth is that in the latter, there is an oasis, it follows that the supervenience base for the content of John's hallucination is duplicated on Twin Earth. Supervenience bases are by definition sufficient to necessitate things that supervene on them. This means that the content of John's hallucination is also present in Twin John's case. Thus the two perceptual experiences share intentional content, and this content is independent of the existence of the oasis.

[EU3] The third argument concerns the causal antecedents of John and Twin John's perceptual experiences. Consider how John's hallucination on Earth came about: through a mix of environmental and physiological factors, John's brain was in a state such that it appeared to him that there was an oasis in the desert. Thus there is a causal story we can tell about how it is that John came to have the hallucination of an oasis, and given INTENTIONAL SIGNIFICANCE, how John came to have a contentful perceptual experience. But now consider Twin Earth, where John and his environment are identical except for

the presence of the oasis. All of the conditions that are jointly sufficient for John to have a contentful perceptual experience of an oasis on Earth are also present on Twin Earth. Thus, the sufficient causal conditions for having a contentful perceptual experience are shared across the two cases, and so plausibly, the content for which these conditions are sufficient is likewise shared across the two cases. These are the arguments for the first form of EXPERIENTIAL UNIFORMITY.

But there is a second form of EXPERIENTIAL UNIFORMITY, one which concerns the entire class of perceptual representations. This type of uniformity, which premise 3' attempts to capture, requires us to assign the same kind of contents to all perceptual representations, so that perceptual representations are all of, for instance, objects or properties or collections of properties, or have uniformly descriptive or uniformly adverbial contents. What kinds of arguments can we give for this form of uniformity? What reasons do we have to think that we should treat the intentional features of all perceptual experiences similarly?

[EU4] Suppose that we are trying to give a theory of the contents of perceptual experiences: we are trying to state the representational features of perceptual experiences, rather than merely talking about them at a high level of abstraction. This seems like an important part of a psychological theory, since attributions of contents play an important role in psychological explanations. But once we do try to state the representational features of particular perceptual experiences in any way that is remotely specific, something important happens. In giving this kind of theory, the theorist must distinguish between the representations that are empty, and to be given one treatment, and the representations that are non-empty, and are to be supplied with representational contents that are dependent on existent objects. But the ability to make this kind of distinction is dependent on the theorist's beliefs about which of the perceptual experiences are veridical and which ones are not. But now suppose he is wrong, and some of the representations he thought were veridical, and object-dependent, turn out not to be. The theorist will be forced to revise his assignment of content to the mental state that he thought was veridical: he will be forced to assign it content of an altogether different *kind*, because of how an empirical detail turned out. In taking this approach, we make our theory of mental content beholden to our current ontological beliefs. Psychologists must wait for the deliverances of physics and metaphysics

in order to assign contents to mental states. This seems like a problematic conclusion.

But even if one is willing to bite that bullet, there is another argument in favor of the second form of uniformity. Many traditional versions of the argument from hallucination rely the notion of *subjective indistinguishability*. Veridical and non-veridical experiences can be subjectively indistinguishable—that is to say, they can be phenomenally identical. An agent may not be able to tell a veridical experience apart from a phenomenally identical but hallucinatory experience, nor may she be able to separate her veridical perceptual experiences from her hallucinatory ones. This form of phenomenal indistinguishability is often taken as support for veridical and non-veridical experiences being of a psychological kind. Tim Crane encapsulates this thought nicely:

When two conscious experiences are indistinguishable for a subject, then the experiences are of the same specific psychological kind. So for example, if my genuine perception of the snow-covered churchyard and my hallucination of the snow-covered churchyard are indistinguishable for me, then these experiences are of the same specific psychological kind [Crane, 2005, p. 6]

That is to say, the inability of a subject to distinguish between veridical and non-veridical experiences is taken to support what Martin [2004, 2006] calls the Common Kind Assumption (CKA):

CKA Whatever fundamental kind of mental event occurs when one veridically perceives, the very same kind of event could occur were one hallucinating.

Subjective indistinguishability is often invoked as a premise in other versions of the argument from hallucination precisely because it is seen to support the CKA: if the CKA holds, we should expect veridical and non-veridical mental states to have contents of the same type. Thus, subjective indistinguishability yields a fifth argument for EXPERIENTIAL UNIFORMITY:

[EU5] When two experiences are subjectively indistinguishable, they are experiences of the same specific psychological kind. When two experiences are of the same specific psychological kind, our account of their content should be the same. More specifically, if one

of the experiences has content that is not object-dependent, then the same should be true of the other experience. Veridical experiences are sometimes subjectively indistinguishable from non-veridical experiences, both counterfactually and actually. One veridical experience can be phenomenally identical to a counterfactual, hallucinatory experience. And further, veridical experiences are indistinguishable—in terms of their veridicality—from hallucinatory ones. That is to say, one cannot determine, on the basis of their phenomenology, which actual experiences are veridical and which ones are not. Thus we should give the same account of the contents of all actual experiences.

I myself find this argument the least convincing of the five, because I find the connection between indistinguishability and the CKA to be questionable. It is perfectly conceivable that subjectively indistinguishable experiences may turn out to be of different psychological kinds. In fact, this is exactly what the knowledge-first epistemologist denies: even though we cannot tell knowledge from belief, knowledge is its own distinctive kind of mental state Williamson [2000]. Of course, this is a slightly different claim, in that it concerns the attitude in question (knowledge vs. belief) as opposed to the kind of content to which the attitude relates us (object-dependent vs. non-object-dependent). But it is still perfectly possible for the disjunctivist about perception to deny that hallucination and veridical perception are of the same psychological kind even though we can't tell them apart from the inside. As we will see, analogs of each of these arguments are available in the case of empty names, and their relative strengths and weaknesses carry over to the semantic case as well.

2.3.2 Empty Singular Reference

The problem of empty names is often presented as a problem for classical first-order logic. Given an empty name in English like “Vulcan”, and a translation of that name into a first-order logic as a , the axioms governing identity yield:

$$(8) \quad a = a$$

The standard rules governing existential inferences then yield the following generalization:

$$(9) \quad \exists x(x = a).$$

This statement is to be interpreted as saying that Vulcan exists. But (9) is surely false, for Vulcan does not exist: Einstein was right and LeVerriere was wrong.

One standard response to this puzzle is to move to a free logic.²⁷ Free logics are logics that admit empty names, but modify the inferential and semantic rules of classical logic to avoid manifestly false conclusions such as (9). While there are many different ways to set out a free logic, most versions modify the natural deduction rules of first-order logic by restricting the instantiation of universally quantified statements to non-empty names. Accordingly, when a is empty, we can never use $a = a$ as a premise, because that would require us to instantiate the universally quantified axiom of identity— $\forall x(x = x)$ —with an empty name. But beyond modifying the inferential rules, the free logician can offer three different types of semantics: positive, neutral, and negative, which yield what have come to be called positive, neutral, and negative free logics. Positive free logics allow some atomic sentences involving empty names to be true, neutral free logic stipulates that all such sentences lack truth-values, and negative free logic forces all such sentences to be false.²⁸

But while the logical status of empty names is an important issue, it is not the core problem that empty names pose. The core problem they pose is this: empty names, like non-empty ones, are semantically significant—they often affect the truth-conditions of sentences in which they figure—and we need an account of how they do so. Call this the principle of SEMANTIC SIGNIFICANCE. The fact that empty names are semantically significant indicates that names have content or exhibit intentionality. By the “content” of a linguistic expression, I mean the aspect of an expression’s linguistic meaning, relative to context, that contributes to the truth-conditions of sentences in which it figures. A compositional semantic theory, if it is to be empirically adequate, must capture this contribution to truth conditions by assigning content to empty names.

²⁷But it is by no means the only response. Many respond instead by claiming that names are not genuinely singular terms, but are rather general terms. While this may clash with the intuitions that names are singular referring expressions, it allows us to keep the logic for our language classical. But there are several different versions of this view. The view that names are actually abbreviated definite descriptions was Russell’s, but the idea that names should be translated as descriptions in which the name itself occurs in the predicate component of that description can be found in Tarski [1983], and is elaborated in Quine [1940, 1950, 1953c] (although Quine himself cites Russell as his inspiration, rather than Tarski). The view is defended in its modern form by Delia Graff Fara [2015], who holds, roughly, that names have the same semantics as common nouns, and also by Zoltán Szabó [2015].

²⁸For more on free logic, see [Bacon, 2013, Crane, 2012, Lambert, 2003, Sainsbury, 2005], among many others.

The way a semantic theory captures the semantic significance of an expression it is by assigning that expression a *semantic value*; semantic values serve to *model* the contribution an expression makes to the truth-conditions of sentences containing it. That is to say, semantic values are modeling devices that a semantic theory assigns to subsentential expressions that allow the semantic theorist to illustrate how the truth-conditions of sentences depend on the meanings of their constituents. Within an intensional semantic framework, semantic values will rarely be ordinary objects or sets of ordinary objects; more often, they will be functions that vary their values with various pieces of information, such as world, time, and context [Lewis, 1970].

The relationship between the semantic value of an expression—a function—and its content is contentious; as with many modeling tools, there is much debate over how realistically such functions should be construed. Many semanticists and philosophers view semantic values as mere devices whose primary goal is to allow for the derivation of the truth-conditional content of declarative sentences.²⁹ These philosophers view the semantic values of subsentential expressions instrumentally. In their minds, a semantic theory has done its job if it assigns semantic values to subsentential expressions that yield the correct sentential truth-conditions, whatever those semantic values are. In what follows, for ease of exposition, I will assume that such instrumentalism is not correct; I will assume that the semantic value of an expression in an empirically adequate semantic theory must reflect that expression's content. On this view, the principle of SEMANTIC SIGNIFICANCE requires a semantic theory to pair empty names with semantic values that capture their distinctive contributions to the propositions expressed by sentences containing them. If you find this assumption unpalatable, I invite you to revise the arguments below so that they concern only the contents of expressions, rather than their semantic values. This revision will not change the important features of the arguments.

Once we assign a semantic value to an expression in order to model its content, we can consider how this semantic value would differ, if at all, in the counterfactual situation where the name was non-empty. The view that the name should be assigned the same semantic value as its empty counterpart is one form of what I will call the principle of SEMANTIC

²⁹This view is clearest in Quine [1960] and Davidson [1979].

UNIFORMITY. But in addition to this kind of uniformity, we might also think that the semantic value of names ought to be uniform across the lexicon, so that the meaning of an empty name does not differ *in kind* from that of a non-empty name. This is a second form of the principle of SEMANTIC UNIFORMITY. With this terminology in place, we can state the problem of empty names as follows:

1. Names, including empty ones, make non-trivial contributions to the truth-conditions of sentences containing them. A semantic theory must account for this significance by providing them with non-trivial semantic values [SEMANTIC SIGNIFICANCE].
2. By hypothesis, the semantic value of an empty name is not an ordinary object; plausibly, it's semantic value will be of a different, perhaps higher, type.
3. The semantic value of an empty name is the same as that of its counterpart in a counterfactual situation that differs only in that the name is non-empty. [SEMANTIC UNIFORMITY 1]
- 3'. The grammatical category of singular referring expressions is semantically uniform: expressions of the same syntactic category should be assigned semantic values of the same type [SEMANTIC UNIFORMITY 2].
4. No name, empty or non-empty, can have an ordinary object as its semantic value.
5. All names have something other than an ordinary object as their semantic value.

This is what I like to call the New Problem of Empty Names. Now let's turn to the arguments for its premises.

Arguments for Semantic Significance

What reasons do we have to think that empty names make a significant contribution to the truth-conditions of sentences, or contribute distinctive contents to the propositions expressed by sentences? There are two arguments, the first is the argument from communication, and the second is the argument from substitution.

First, consider two people who both think that think that the devil—Lucifer—exists (and if you believe that the devil does exist, then momentarily suppose that he doesn't). If the devil doesn't exist, then these two people are making use of an empty name. In spite of this, they manage to communicate thoughts with sentences involving “the devil” and “Lucifer”. Like many other people, they come to form beliefs on the basis of utterances of these sentences. Like many Christians, they likely have lots of beliefs about the devil: that he exists, that he tempts people, wreaks havoc on Earth, torments the damned, *etc.* These beliefs may figure prominently in the explanations of many of their behaviors. If we think that an adequate psychological theory needs to explain behavior by ascribing contents to mental states, then these beliefs are contentful mental states *par excellence*, even if they are about something that doesn't exist. Given that these mental states are contentful, it seems overwhelmingly plausible that the contents of these beliefs are what are expressed by sentences like “The devil tempts the righteous”. If such sentences did not express contents, and more specifically, if “the devil” were not itself contentful, then it would seem impossible to explain how such communication and belief formation occurs.

Thoughts about the nonexistent make up a significant portion of our cognitive lives, and the sentences that express them form a correspondingly large portion of our discourse. This is particularly true as we look further back into history. For instance, if you're an atheist, nearly all theological thought falls into the category of thought about the non-existent. But on pain of having a radically incomplete explanation of the behaviors of many, if not most human beings, we need to allow that we're having contentful thoughts about the non-existent. But now suppose that we hold that empty names are not significant. It follows that the majority of religious assertions do not express propositions. It also follows that there is nothing semantic that distinguishes “Lucifer” from “Jesus”. But we used those very sentences to express and form beliefs thousands of times over, even without the sentences expressing any content. This borders on absurdity. And the problem is magnified if we consider how many false theories scientists have come up with over the centuries. If empty terms are semantically vacuous, then there is a gaping hole in our ability to explain the most basic communicative practices.

The second argument that empty names are semantically significant is what I call the

argument from substitution. Consider the following examples:

- (10) Sherlock Holmes is famous.
- (11) Siegfried is an unappealing hero.
- (12) Sherlock Holmes is more famous than Dr. Manhattan.³⁰

Absent anterior theoretical commitments, it seems that these sentences contain genuinely empty names and are nonetheless true. And moreover, if we substitute one empty name for another, we are often left with the strong feeling that the resulting sentence is false:

- (13) Dr. Manhattan is famous.
- (14) Superman is an unappealing hero.

This indicates that there is something relevant to the truth of the sentence that changes when we substitute one empty name for another. The only difference between the sentences is the name, which seems to indicate that the difference in truth-value must be on account of the semantic contribution of the name.³¹ Now suppose that the semantic value value of each of the names above is its extension. By virtue of the fact that they are *empty* names, their extension is either a dummy object like the null set, or nothing at all. So, on the assumption that our semantics is compositional, when we substitute one extensionally equivalent name for another, the truth-value of the whole sentence must remain the same; this follows because

³⁰The fact that you might ask yourself who Dr. Manhattan is is itself a testament to the truth of (12). He is a character in Alan Moore's graphic novel *Watchmen*.

³¹Even though these contexts look fully extensional, one might be tempted to respond by claiming that the predicates in (10) and (11) actually create intensional contexts, and that this somehow absolves us from needing to provide the names with distinctive semantic values. But suppose that this is true. On what grounds do we ordinarily call a context intensional? Ordinarily, we call a position intensional when expressions that occupy that position can be empty while the sentence remains true, and that the substitution of co-referential expressions in that position can affect the truth value of the sentence. These tests indicate that the expressions in these positions do not contribute their extensions to the truth-conditions of the sentence. Perhaps they contribute their intensions. But this is exactly the phenomenon that we are pointing to: "Superman" and "Dr. Manhattan" have null extensions, and given our intuitions about the above examples, that cannot be what they contribute to the truth-conditions of (13) and (14). So calling the contexts above intensional does not defend against the points made; far from it: it makes the point. I take this as an argument for SEMANTIC SIGNIFICANCE: substitution tests show that names, including empty ones, contribute non-trivially to the truth-conditions. This shouldn't surprise us; empty names still represent, and it seems plausible a name's representational features will manifest themselves by affecting the truth-conditions of sentences containing them.

compositionality forces the truth-value of the whole to be a function of the semantic values of the ultimate constituents. But this clearly contradicts our intuitions about the sentences above. As a consequence, if we keep the assumption that our semantics is compositional, the semantic values of empty names cannot be their extensions; different names must have different semantic values of a higher type. This is the substitution argument for the semantic significance of empty names, which pushes us to provide empty names with non-trivial semantic values.

Arguments for Semantic Uniformity

Once we have decided on semantic values for empty names, which is a semantic theory's way of giving an account of content, there are reasons to think that this account should be uniform across empty and non-empty names. Above I formulated two versions of SEMANTIC UNIFORMITY, one that comes from considering a single empty name and its hypothetical non-empty counterpart, and another claiming that the category of names should be lexically uniform. This section will provide arguments for both in succession. Once again, the arguments concerning a single name and its hypothetical non-empty counterpart will be adaptations of the arguments given in Segal [2000], and make use of a Twin Earth thought experiment.

Consider an empty name: take "Lucifer". Presume, if you don't already believe it, that there is no such fallen angel. SEMANTIC SIGNIFICANCE implies that "Lucifer" contributes importantly to the truth-conditions of sentences containing it. That is to say: it exhibits intentionality, or has content. Now consider a world in which everything is exactly identical to ours, except for that fact that the creature named in the Bible and the apocrypha exists.³² I will offer five arguments that the semantic value of "Lucifer" is the same across the two possible cases.

[SU1] Suppose that an expression's content plays an important explanatory role in a theory of communication: how speakers manage to communicate with that expression is

³²I think this is a possible situation, but Kripkean wisdom has it that if a name is empty, then it is necessarily empty. Given that Lucifer does not exist, it follows that he could not have existed. I take this to be a drastically counterintuitive claim, but if you are an orthodox Kripkean, just frame the argument in terms of epistemic, rather than metaphysical, possibility.

(at least partly) to be explained by the intentional features of that expression. From it follows that the semantic value of “Lucifer” is uniform across the two worlds mentioned above. Here is the argument. In both of the worlds described above, the name “Lucifer” behaves identically in discourse. It expresses and conveys beliefs that play identical roles in psychological explanations. That is to say: the term plays the exact same causal role in the communicative process. Whatever kind of content allows the empty term to play its role in the actual world is also present in the second world in which the term is non-empty. This means that the fact that Lucifer exists is not relevant to its role in communication. Given our supplementary premises, it follows that the two names have the same semantic value.

[SU2] SEMANTIC SIGNIFICANCE shows that “Lucifer” has content. But now consider the supervenience base of the content in the actual world, in which “Lucifer” is empty. This supervenience base might comprise a set of properties like: the beliefs, expectations, and intentions of many competent users of the name, and perhaps properties of groups of such speakers that form the basis of a convention to use the name in the same way as other members of the community. But this supervenience base will be duplicated in a world that differs from the actual one only in that Lucifer exists. But this means that there is shared content across the empty and the non-empty names, and they differ at most in their extensions.

[SU3] Lastly, we can consider the causal conditions necessary for a name in a language to have content. Still supposing that “Lucifer” is empty on Earth, we can consider the diachronic conditions that are sufficient for the name to have content. These might include many initial uses of the name, perhaps accompanied by false beliefs, causal chains of uses extending from these initial uses, and perhaps certain kinds of conventions of using the name to refer to the same thing that others in your community use it to refer to. Then we can consider Twin Earth, which differs only in that Lucifer exists. All of these sufficient conditions are still met on Twin Earth, and so the diachronic sufficient conditions for the name to have content are also satisfied on Twin Earth. Given that the conditions sufficient for content are identical across the two cases, and these conditions are shared, the names should share content across the two cases as well.

But there is another version of SEMANTIC UNIFORMITY that is stronger than the first

version. Instead of focusing on a single empty name and a possible non-empty counterpart, it asserts that the grammatical category of names should be given a single kind of semantic value: names should be semantically uniform. Within a type theory, this principle amounts to the requirement that all names should be assigned semantic values of the same type. This requirement embodies the often-endorsed constraint that there be a homomorphism from the set of syntactic types to the set of semantic types. Uniformity of this sort should be familiar from Montague’s PTQ [Montague, 1974b], where he assigns types of the highest sort necessary to each syntactic category, and then extensionalizes them in special cases using meaning postulates. In theories that are not type-theoretical, it might mean that our semantic clauses state the semantic values of expressions in the same ways, and that our reference clauses ought state truths whether or not the names are empty. For example, in a Davidsonian theory, we might specify semantic significance using universally quantified biconditionals.³³

There are two main arguments in favor of the lexical semantic uniformity of names. [SU4] The first argument comes from the epistemic constraints faced by semantic theorists. In distinguishing empty from non-empty names, a semantic theorist does so to the best of her knowledge. However, the theorist is almost certain to miscategorize at least one name, because she, and theorists more generally, are almost certain to be wrong in at least one of her beliefs about what exists. Now suppose that we decide to treat empty and non-empty names differently in our semantics. This means that how the theorist applies the different treatments depends on her current state of our scientific or metaphysical knowledge: we make our semantics beholden to our current ontological beliefs. But this is surely unacceptable. If we were to handcuff our semantics in this way, every time we made a new ontological discovery, we would be forced to revise both the meaning of a word as well as the *type* of meaning it has. With each new ontological discovery, we would likewise discover that our semantic theory in fact specified a word’s contribution to compositionally determined meanings in the wrong way. When we discovered that Vulcan didn’t exist, we didn’t go back and revise the lexical meaning of “Vulcan”, or claim that in fact, “Vulcan” actually

³³On some views, such biconditionals will still require the move to a free logic, as in Sainsbury [2005]. However, Larson and Segal [1995] state their semantic theory in a Davidsonian way without a free logic, but hold that the correct semantic theory may be false on account of the fact that some names are empty.

turned out to be a description. Rather, it seems clear that the meaning stayed the same and the world simply turned out differently than we thought. Perhaps this point can be encapsulated by saying that our semantic postulates need to be obvious, since knowing them does not require speakers to know complicated or hard-to-come-by scientific or ontological facts.

[SU5] The second argument actually serves as an argument for both versions of the uniformity principle. It is plausible that one goal of semantic theorizing is to capture what speakers know about their own language. However, speakers cannot distinguish any particular empty name from a possible non-empty version of the same name. Nothing in my semantic competence nor anything in our pretheoretical conception of meaning, can separate the empty from the non-empty names. Thus, with respect to speaker's semantic knowledge, and all of the relevant linguistic information, the content of a name like "Lucifer" is identical with that of its non-empty counterpart. This consideration counts in favor of supplying names with semantic values other than objects. But more broadly, insofar as semantics is supposed to reflect a speaker's knowledge of meaning, it cannot give a non-uniform semantics for the grammatical category of names. It seems implausible that speakers are drastically and consistently mistaken about the meanings of their own words, even though they may be drastically mistaken about what exists. The only way to accommodate this fact within a semantic theory is to give a semantics for the category of referring expressions that is uniform and not object-dependent.

Thus we have five total arguments in favor of SEMANTIC UNIFORMITY, three that apply to a single name and its counterfactual counterpart, and two that draw on general considerations about the grammatical category of names.

2.4 Identity of the Arguments

Above I provided reconstructions of the problem of empty names and the problem of hallucination as puzzles of intentionality, and presented several arguments for their premises. While I'm sure you've already noted many of the similarities between the puzzles, I'll now attempt to make them explicit, and in so doing, I'll construct one central argument from

which both particular arguments can be derived. This central argument can be seen as a master argument against the object-dependence of content. Insofar as internalism is the thesis that mental and linguistic content do not depend on objects in the world, the argument can be seen as a master argument for content internalism. Keep in mind, however, that I am not endorsing this argument; it is simply instructive to have such a master argument because responses to it help us categorize the various views on intentionality.

Both the problem of hallucination and the problem of empty names concern representations—in one case perceptual experiences, and in the other case names—which fail to bear any relation to ordinary material objects. The first premise in both arguments is what we can call the principle of SIGNIFICANCE: despite not bearing any such relation to ordinary material objects, empty representations, like all representations, exhibit intentionality. That is to say: every representation represents something.

The second premise is a definitional premise. By definition, empty representations do not represent ordinary objects. In the perceptual case, hallucinations are just perceptual representations where there is no ordinary object represented—the world fails to comply, and likewise in the case of empty names—empty names are representations that fail to have an ordinary object corresponding to them.

The third premise is the principle of UNIFORMITY: our account of the intentional features of empty and non-empty representations should be uniform. In both arguments, there are two versions of the UNIFORMITY premise. The first version considers a single empty representation, and making use of its significance, argues that this significance is also present in a counterfactual situation in which the representation is nonempty. The second version of the UNIFORMITY premise concerns the entire category of representations in question: in the first case, perceptual experiences, and in the second, names.

Together, these premises entail the conclusion that no representations for which the problem of nonexistence can arise represent ordinary material objects. On most views of intentionality, the ability to represent the non-existent is a pervasive feature of representation, and on some traditional views of intentionality, the possibility of representing the nonexistent is even criterial for intentionality.³⁴

³⁴See Caston [2001] for a historical discussion of this and other features of intentionality.

But this confluence of the premises of the arguments is not the only evidence we have for their identity. The arguments for each of the premises are also nearly identical. Consider the arguments for the perceptual and linguistic versions of SIGNIFICANCE. The first argument for each version notes the role that empty representations play in their respective domains. In the perceptual case, empty representations play a significant role in the explanation of behavior, and in the linguistic case, empty names play an integral role in a theory of communication. The second argument for significance involves how empty NPs contribute to the truth-conditions of reports of hallucinations, and sentences involving empty names. Substitution of one empty NP for another in both cases affects the truth-conditions of the sentences, which indicates that these NPs contribute important semantic material to the truth-conditions of sentences containing them. This serves as a second argument that both hallucinations and empty names are significant.

The arguments for uniformity in the two cases are also nearly identical: each of the arguments EU1-EU5 correspond closely to the arguments labelled SU1-SU5. The first three arguments in each case are based on a Twin Earth scenario. We can call them, respectively, the argument from explanatory role, the argument from supervenience, and the argument from causal antecedents. Arguments EU4-EU5 and SU4-SU5 both concern actual experiences and words, without resorting to counterfactuals, and they try to establish that the *kind* of content that a theory should assign to representations should be uniform across the entire relevant category.

The fact that the arguments can be reconstructed in such close forms, and the fact that even the arguments in favor of the premises are nearly identical, points to an underlying identity. The problems appear to come from the same general principles concerning representation, simply applied to two different domains. Of course, there will be differences that arise from the application of the principles to two different kinds of representations, but the core ideas remain the same: empty representations are contentful, empty representations share (their kind of) content with non-empty representations, and so non-empty representations cannot have object-dependent content. This argument is a brief version of a master argument against object-dependent content. As we will see in the next section, many views on perceptual and linguistic content can be categorized according to where or

whether they reject the argument. In this sense, isolating this core argument serves an important unificatory role: it unifies what seemed to be several disparate problems in two different subdisciplines in philosophy into a single, succinct argument.

Let me ward off two potential objections. First, you might object that I have not shown that the problem of hallucination and the problem of empty names are identical, and that instead I have merely shown that there is a structural analogy between them. Structural analogies, you might claim, are not identities, and further, structural analogies are cheap, and can be found merely looking at the problems in a sufficiently abstract way. These are legitimate worries. However, the core of the worry comes from the fact that it is unclear how to argue for the identity of two *puzzles* in any way other than by exhibiting structural analogies. It is also unclear what the identity conditions for puzzles or arguments really are. What has been established is that the problems are structurally identical down to a shockingly fine level of detail, and it doesn't matter one way or another whether one thinks this is sufficient to have established an identity. My goal is to show that the two problems have spaces of solutions that are identical, and so should be approached and prospectively solved together. All that is necessary for this is that the premises and conclusion are sufficiently similar to show that each way of responding to the argument—by either denying a premise or accepting the conclusion—is available in both the perceptual and linguistic case.

Second, one might object that I have based my arguments for the identity of the problems on my own reconstructions of the problems, and these reconstructions might not be faithful ones. Further, one might even claim that I have even curated the arguments in favor of the premises so that they would come out looking maximally similar. With respect to the first claim, I don't deny that there was a process of reconstruction involved, but I do deny that the reconstructions I have given, particularly of the problem of empty names, in any way distort their most important aspects. In fact, the opposite is true. My account of the problem of empty names makes it out to be a problem about our semantic account of names generally, which makes it broader and more pressing than most other versions. If I can be accused of distorting the problem, it is by presenting the strongest version of the problem. And even if the problem I presented is not what people ordinarily have in mind

when they think of the problem of empty names, what I have presented is still *a* problem of empty names. Lastly, it is definitely true that I have selected the arguments in favor of each of the problems so that they match. But this is not a criticism—it is a testament to the identity of the problems. The fact that the very same arguments work to establish the conclusions in both the case of hallucinations and the case of empty names shows just how similar the problems actually are. The ease of transfer from the perceptual to the semantic case supports the underlying identity, rather than calling it into question.

2.5 The Space of Possible Responses

The major views in the philosophy of perception can be individuated in terms of how they respond to the problem of hallucination. Following Tim Crane, we can isolate four such views: sense-data theory, intentionalism, adverbialism, and disjunctivism. We can further divide these views into relational and non-relational views: sense-data theory and disjunctivism are committed to relational views of perception, while intentionalism (arguably) and adverbialism are non-relational theories. Each of these views responds to the problem of hallucination in a distinctive way, and then develops the response into a general theory of perceptual content. But given the identity of the problem of hallucination and the problem of empty names, each response to the problem of hallucination has a corresponding response to the problem of empty names, and as a consequence, each general view of perceptual content has a corresponding view on the semantic value of a name. That is to say: the space of possible solutions to each of the problems is identical. The goal of this section is to map each perceptual view onto its semantic counterpart, and show how this yields some unrecognized views on the semantics of empty singular referring expressions.

Sense-data theory accepts both the premises and the conclusion of the argument from hallucination. The sense-data theorist accommodates the conclusion by claiming that the objects of perception are uniformly non-physical, mental entities called sense-data. These entities are what we are aware of in both cases of hallucination and cases of veridical perception, although in the veridical case, we may also be indirectly aware of ordinary objects. There is a perfectly analogous view in the philosophy of language that results from

accepting both of the premises of the argument from empty names. It is the view, famously attributed to Locke, that words refer to ideas. These ideas are mental entities, and they play the role of intermediary between words and the world. In the case where a name does not have a worldly bearer, it still refers, as all names do, to ideas in our minds. Thus, Locke's view of names introduces a "veil" of ideas between words and the world, just as sense-data theory introduces a veil of immaterial, mental entities between mind and world.³⁵

The second relational response to the problem of hallucination is disjunctivism. Disjunctivism, characterized minimally, is the view that whatever the contents of hallucinatory experiences are like, they are different in kind from the contents of veridical perceptions. Thus, disjunctivism can be seen as the outright denial of EXPERIENTIAL UNIFORMITY.³⁶ The goal of denying EXPERIENTIAL UNIFORMITY is to maintain that veridical perceptual experiences are genuinely relations to material objects. Thus the disjunctivist keeps relationality at the expense of uniformity. There is a clear analogue to this view in the case of names. Many semanticists and philosophers of language wish to treat ordinary names as bearing a causal relation to ordinary objects, treating them as mere tags. This is the Millian view of names. But no such possibility is available for empty names, unless one tells a very complicated story about how we interact with abstract or non-existent objects. The other option, which is an options sometimes adopted, is to hold that empty names have a totally different semantics than non-empty ones: perhaps they are not even genuine referring expressions. This is the disjunctivist view of empty names, minimally characterized, and perfectly analogous to the view in philosophy of perception. One version of this view is to hold the view, discussed by Szabó and Thomason [Fortchoming], that empty names have descriptive contents while non-empty names have objects as their contents, and to make this a principled epistemological distinction.

The first non-relational view in the philosophy of perception is intentionalism. The intentionalist claims that, at bottom, experience is not a relation, but is rather a form of

³⁵Locke uses the terms "primary signification" and "secondary signification". On Locke's view, a word's primary signification is an idea, and its secondary signification can be an object in the external world. For a modern incarnation of the Lockean view, see Davis [2003], who holds that the meaning of a word is an idea.

³⁶Depending on the disjunctivist's views on hallucinatory experience, the disjunctivist may also deny INTENTIONAL SIGNIFICANCE. If one thinks that hallucinatory experiences do not have genuine accuracy conditions because they lack objects, then they must deny INTENTIONAL SIGNIFICANCE.

representation, and the world may or may not be how we represent it to be. According to the intentionalist, what a mental state is about, or what it represents, is, by definition, its intentional object. As Tim Crane [2009], the main proponent of intentionalism, puts it: “since intentional objects are what we think about, desire, or hope for, and we can think about or desire or hope for things that do not exist, some intentional objects do not exist. These can be called ‘mere intentional objects.’” Treating mental states in this way—as always about, or directed toward, or representing an object—seems to make mental states relational. But Crane uses intentional “intentional object” to capture the idea that every mental state is *of* something, without claiming that this surface relationality is the ultimate metaphysical story about perception. Rather, the intentionalist holds that thought about the non-existent can ultimately be explained reductively, in a way that is non-relational. Crane [2012] argues that this reductive explanation will be a psychological one, but will be particular to different cases.

There is a nearly perfectly analogous view on the semantics of names. The intentionalist maintains that mental states, including hallucinations, have intentional objects, but that hallucinations have non-existent, “merely intentional” objects. The corresponding linguistic view is that of Parsons [1980], on which all names denote objects, with empty names denoting non-existent objects. Parsons likewise gives a reductive view of what it is for linguistic representations to be about the non-existent. According to Parsons, each non-existent object is in fact just a collection of properties that is not instantiated. However, the intentionalist might not want to hold that hallucinations have intentional objects, in order to avoid commitment to merely intentional objects. This is a view suggested in Crane [2005]. This sort of view is one on which hallucinations have content, but they do not have intentional objects. This view also has an analogous linguistic view: it is analogous to the orthodox, non-Evansian Fregean view on which empty names have senses but do not have referents.

The last response to the problem of hallucination is adverbialism. Adverbialism is a thoroughly non-relational approach to experience. The adverbialist’s view is that to sense, or experience something white is not to experience an object, but rather to experience in a particular *way*: whitely. In general, it treats verbs (and verb + PP combinations) that seem

to express binary relations—such as “senses”, “experiences”, “is aware of” and “sees”—as in fact complex, one-place predications. It does this by treating the object-positions of these verbs as modifiers, modifying the verb, rather than as noun-phrases which contribute an object that acts as an argument to the verb. The NPs in the object-positions of the verbs rather serve to modify how it is that we sense, experience, or are aware.

While this view is a relatively well-known view in the philosophy of perception, few people have recognized that there is a corresponding view available in the philosophy of language which treats names as referring in particular ways, as opposed to referring to things. This, which I call “semantic adverbialism”, is the view that reference is not, or is not always, a relation. On a non-relational view of reference, when I say I am referring to Superman, I mean that I am referring in a particular way—the “Superman” way—but there need not be any object to which I am referring. Further, when I say that a particular word—like “Superman”—refers to Superman, I am not saying that there is an object to which this word bears some relation, the reference relation. Semantic adverbialism can be seen as a way of working out the idea, famous from Quine, that linguistic expressions are meaningful, without having *meanings*:

I feel no reluctance toward refusing to admit meanings, for I do not thereby deny that words and statements are meaningful. McX and I may agree to the letter in our classification of linguistic forms into the meaningful and the meaningless, even though McX construes meaningfulness as the having (in some sense of “having”) of some abstract entity which he calls a meaning, whereas I do not. I remain free to maintain that the fact that a given linguistic utterance is meaningful (or significant, as I prefer to say so as not to invite hypostasis of meanings as entities) is an ultimate and irreducible matter of fact; or, I may undertake to analyze it in terms directly of what people do in the presence of the linguistic utterance in question and other utterances similar to it [Quine, 1953c].

Quine, of course, thought that the right way to do away with unwanted reference was to translate names into a regimented language as predicates—“Pegasus” as “Pegasizes”—

and then treat them as Russell did, as contributing both quantificational and predicative material. But this version of the non-relational approach to reference is not forced on us. There are other ways to represent non-relationally. Think again of our example from section 2, of how a painting represents. It represents by having certain intrinsic features: the locations of the brushstrokes, the colors, *etc.* It represents by being a painting of a certain *kind*, or of a certain *sort* [Goodman, 1976, Ch. 1]. Much the same is true for semantic adverbialism: “Superman” refers to Superman by being a representation of a particular sort, or representing in a particular way, and these features—the way it represents or the kind of representation it is—are intrinsic.³⁷ Thus, the identity of the space of possible responses to the problem of empty names and the problem of hallucination has paid dividends: we now have a new, relatively unexplored view of the semantic of empty names: semantic adverbialism.³⁸

2.6 Semantic Adverbialism and Intensional Transitive Verbs

Recall from above Anscombe’s claim that the problems of hallucination and illusion are really problems of how to provide a semantics for intensional transitive verbs. My claim is that given the identity, or close structural analogy, of the problem of hallucination with the problem of empty names, this strategy is should be extended to verbs that we use to specify the semantic significance of names. Anscombe claims that perceptual verbs like “perceives”, “sees”, and “senses” are intensional transitive, and giving an account of the semantic contribution of the NPs in their object positions, on their intensional readings, yields an account of hallucination.³⁹ I have argued in other work, on both empirical and theoretical grounds, that the same is true for the **V + PP** combinations “refers to”, “applies to”, and “is true of”; each of these has the semantics of an intensional transitive verb. This means we can apply the Anscombian proposal to solve the problem of empty names; whatever the best semantics for the intensional reading of an ITV turns out to be will yield

³⁷Ways and sorts are both forms of modifiers: deciding between them rests primarily on our decision of whether to make use of an event-based semantics.

³⁸For an early adverbialist view about perception, see [Chisholm, 1956]. See also: [Tye, 1975, 1984], and [Quine, 1953b, 1960, 1969].

³⁹For a new argument that perceptual verbs such as “sees” and “perceives” have intensional readings, see Bourget [2016].

an account of the semantic value of an empty name.

Just as there are relational and non-relational approaches to intentionality, so there are relational and non-relational approaches to the semantics (for the intensional reading) of intensional transitive verbs, and each different view of the semantics of intensional transitive verbs yields the foundation for a different theory of intentionality. Thus, each pair of corresponding solutions to the problems of hallucination and illusion can be paired with a single approach to the semantics of ITVs. For the purposes of our discussion, we can restrict our attention to “perceives” and “refers”. On relational views of intentionality, the primary question to answer is: what is the nature of intentional objects? The very same is true for relational approaches to intensional transitive verbs: there are many approaches, and they are differentiated by what objects they take to be the relata of the relation expressed by an ITV. Consider one very early and influential relational treatment: according to Montague [1974b], on their intensional readings, ITVs like “seeks” express relations from worlds to sets of properties: they express relations between the subject and an intensional quantifier. This approach can be applied to yield views of both hallucinations and empty names. We can, with Mark Johnston [2004], treat the object of hallucination as a collection of properties: the collection of properties that a scene or object is experienced as having in the hallucination. Similarly, one approach to the semantics of empty names treats them as denoting intensional quantifiers. Thus, Montague’s approach to ITVs pairs with, or maybe even yields, a view on both the metaphysics of hallucination and an account of semantic value of an empty name.

On the other hand, two recent well-developed approaches to the semantics of intensional transitives, due to Forbes [2006] and Moltmann [2008, 2013], provide ITVs with a semantics that is nonrelational. Since these views can be pursued to yield the foundation for an adverbial view of content, I want to dwell on them at a bit more length. On Forbes’ view, intensional NPs, serve as arguments to a special thematic role in logical form, but importantly, the argument they provide is not the theme of the verb. Instead, it serves as a modifier. On his view, Intensional verbs are *theme suppressed*. Moltmann proposes a semantics for ITVs on which the intensional complements of ITVs are given a *syncategorematic* semantics, and so do not provide an argument to logical form at all. Rather, they contribute to the

specification of an overall truth-condition for the sentence. This provides a second non-relational option. In what follows, I will focus primarily on Forbes' view, but I think that Moltmann's view would also be employed to develop an adverbial theory of content.

Consider a paradigmatic intensional ascription:

(15) John is searching for a dancing dragon.

On the approach outlined in Forbes [2006], the logical form of (15) is (16):

(16) $\exists e(\text{Looking-for}(e) \ \& \ \text{Agent}(e,\text{John}) \ \& \ \text{CHAR}(e,\text{a dancing dragon}))$.

We can paraphrase this as follows: John is engaged in a search that is characterized by being for a dancing dragon. The key part of this semantics is the special thematic role CHAR, which is a relation between the event and the semantic value of “a dancing dragon”, which in this case is a generalized quantifier. CHAR serves to specify satisfaction conditions for the search, effectively playing the role of a modifier; I will discuss how CHAR serves to specify satisfaction conditions for searches below. However, the most important thing is that CHAR does not provide the *theme* of the search; “a dancing dragon” is not the direct object of the ascription, and its semantic value is not what John is looking for. In other words, the generalized quantifier denoted by “a dancing dragon” is not the theme, or object, of the search. In nonspecific searches, there is no theme; instead, CHAR serves as an adjunct or modifier. This approach is similar to the proposal for *that*-clauses made by Pietroski [2000], and echoed by Hawthorne and Yli-Vakkuri [forthcoming 2017] in their discussion of narrow content. The idea is that content-ascribing clauses are associated with a special thematic role, which they call *CONTENT*, that specifies the content of a particular clausal psychological verb, such as “believes” or “thinks”. We can see CHAR as a content-specifying modifier.

“Refers to”, like “look for”, is an intensional transitive verb, so we can apply this semantics to “refers to” to develop a non-relational theory of reference. First, consider a report of speaker's reference:

(17) John is referring to a dancing dragon.

We can give a semantics for this as follows:

(18) $\exists e(\text{referring-to}(e) \ \& \ \text{Agent}(e,\text{John}) \ \& \ \text{CHAR}(e,\text{a dancing dragon}))$.

This semantics yields that John is the agent of an event of referring-to, and that event is characterized by being one of referring to a dancing dragon.

But what about the references of words? How might this proposal help with their features? In order to apply this view to words, we need to invoke states, along with events. This is a proposal defended by Parsons [1990]. Consider (19):

(19) “Unicorn” refers to unicorns.

We could provide a semantics for this as follows:

(20) $\exists s(\text{referring}(s) \ \& \ \text{Subject}(s,\text{“Unicorn”}) \ \& \ \text{CHAR}(s,\text{unicorns}))$.

Roughly speaking, this semantics interprets (19) as saying that there is a state s that is a state of referring, the subject of the state is the expression “Unicorn”, and the referring is characterized as being to unicorns.

However, this proposal makes little progress unless we can give an informative explanation of how the novel thematic role CHAR works. Both Forbes and Moltmann’s proposals face the problem that they do not provide truth-conditions for intensional transitives completely at the level of logical form. In order to specify truth-conditions, they must resort to stating a separate condition on what it takes for the entire construction to hold, invoking what has come to be called the *decompositional* approach. In Forbes’ case, for instance, he states a meaning postulate that connects $\text{CHAR}(e,Q)$ to a complex condition that specifies what it takes for a quantifier to characterize an event.

(21) $(\text{char}(Q))(e)$ iff $\Box(\text{for any } \bar{e} \text{ such that } R\bar{e}e, \text{ for } Qx, \text{ there is some } e' \text{ that is part of } \bar{e} \text{ such that } Fe' \text{ and } x \text{ is a theme of } e')$

This schematic biconditional links CHAR, in the logical form above, to the idea of success conditions for the search, which, on Forbes' view, are necessary conditions, as represented by the "Necessarily" operator above. In the schema above, F stands in for a verb that would specify what success would amount to; in the case of "searching" or "looking for", the instance of F would be "finding". The idea is that, in any event that qualifies as a successful F (or search), one must find Qxs . Thus, if I am searching for three dogs, in order for my search to be successful, there must be a course of events \vec{e} that has three parts, in each one of which I find a particular dog.

This view provides us with an interesting advance in the study of adverbialism. In the philosophy of perception, Adverbialism was traditionally proposed as an alternative to the so-called "act-object" account of perception. Instead of perceptual states being relational, they were monadic states of particular sorts. In being monadic, it has traditionally been held that adverbialism is inconsistent with accounts of content based on truth- or accuracy-conditions. This is due to the fact that truth-conditions seem to require, for their specification, relations to objects in the world. The conditions can only be specified if we pick out those objects, and say how things stand with them. Thus, adverbialism is typically seen as at odds with the dominant view of perceptual content, which is given in terms of accuracy conditions.

This worry is even more pressing in the semantic case, where plausibly, the goal of the theory is to state the truth-conditions of English sentences. What good would a theory of the contents of empty names be if it did not help us to specify the contribution that empty names make to truth-conditions? But Forbes' approach to the semantics of CHAR provides a mechanism for assigning accuracy-conditions to perceptions, and so is also capable of serving to assign reference- or satisfaction-conditions to empty and non-empty linguistic expressions. The general idea is that, on the *de dicto* reading of either a perceptual ascription or an ascription of reference, the NP occurring in the object-position of the verb is not the verb's theme, but rather characterizes the way in which the subject of the ascription refers or perceives. The details of how to assign accuracy- and reference-conditions to such ascriptions remain to be filled out, but this provides a clear framework in which such an adverbialist project can be pursued. And lastly, this proposal can also accommodate someone who is

inclined toward a traditional adverbialist approach, on which perceptual states do not have accuracy conditions. In such a case, one can simply drop the meaning postulate above, and leave CHAR as a primitive.

I take the above outline to merely be suggestive of how we might pursue an adverbial theory of semantic content. However, even if one is not inclined toward adverbialism, treating the various approaches to the problems of hallucination and empty names as tied to the semantics of intensional transitives is still theoretically enriching. Other approaches to the two problems emerge from competing views of the semantics of ITVs. For instance, one view of the semantics for ITVs treats them as denoting ordinary objects that may fail to exist. This is the approach in Priest and Read [2004]. This pairs with a view of intentionality that treats representations as about objects that may fail to exist: the Meinongian view. On this view, both hallucinations and empty names will have intentional objects, and these intentional objects will be non-existent. We might also treat ITVs as expressing relations between the subject of a sentence and a property, as in Zimmermann [1993, 2001, 2006]. There are corresponding views of intentionality: consider the view that names have predicative type, and actually denote properties [Fara, 2015, Szabó, 2015], or the view that the content of a perceptual experience is a property Siegel [2010a,b].

One last example will suffice to demonstrate the connection between the semantics of ITVs and theories of intentionality. One major strand of thought concerning intentionality is that intentionality is something fundamentally propositional. One way of putting this thought is that all content is propositional content. This view is not often articulated, but it is implicit in much talk of content.⁴⁰ This type of view is entailed by views that take perceptual content to be accuracy conditions. Accuracy conditions, in being assessable for correctness, or truth, differ from non-propositional contents, and it is this assessability that makes content distinctively propositional. Such propositional views of intentionality have a corresponding theory of the semantics for intensional transitives, the view that the correct semantics for ITVs is one that reduces them to propositional attitudes. This view is often called “Propositionalism” about ITVs, and has been championed at one point by Montague

⁴⁰Byrne [2001], McDowell [1994], Sellars [1956], Siegel [2010a,b], Tye [1995], Hawthorne and Yli-Vakkuri [forthcoming 2017].

[1974c] and by a host of others: Larson et al. [1997], Larson [2002], Quine [1956] all argue for propositionalism about ITVs.

While there are many other views on the semantics of ITVs, and many other views on the nature of intentionality, the connections I have drawn above between different accounts of the semantics of ITVs and different theories of intentional content should suffice to show that the semantics of representational verbs like “perceives” and “refers” stands to play an important role in a theory of intentionality. Since both the problem of hallucination and the problem of empty names are instances of intentionality’s central puzzle, the problem of non-existence, the semantics of ITVs has a central role to play within both semantics and the philosophy of perception.

2.7 Conclusion

I take myself to have shown the following: the problem of hallucination and the problem of empty names are both puzzles of intentionality. But they have many more specific similarities: the most general forms of the arguments come from the very same premises, and the arguments for these premises coincide. This points to an identity between the two problems, or at least a very fine-grained structural analogy. If the two problems are in fact identical, their spaces of possible solutions should likewise be identical; and this is largely borne out: all of the major responses to the problem of hallucination in the philosophy of perception have corresponding views that serve to respond to the problem of empty names. All other things being equal, an argument for a particular response to the problem of hallucination serves as an argument for the corresponding response to the problem of empty names. Lastly: my preferred response to both the problem of hallucination and the problem of empty names is the Anscombian one: the solution to both problems is to provide a semantics for intensional transitive verbs. Just as Anscombe argues that “perceives” and “sees” are intensional verbs, I have argued that “refers” is an intensional verb; a solution to the problem of empty names hinges on our finding its correct semantics.

Chapter 3

Intensionality and Phrasal Unity

3.1 Introduction

Quine is famous for doing violence to grammar in order to avoid metaphysical extravagance. He had a particular aversion to intensions, and engaged in much logical and grammatical wrangling in order to avoid commitment to intensional entities. Of course, grammatical violence did not worry Quine in the slightest; in general, he was not concerned with the semantics of natural language. Rather, he was interested in regimenting natural language in ways that would make it suitable for scientific theorizing, a purpose for which he considered natural language irremediably unclear. And so as one would expect, many of the grammatical hoops through which Quine and his nominalist contemporaries jumped in their regimentations have been dismissed by analytic philosophers as bad linguistics.

However, there is one strategy that Quine endorsed out of his distaste for intensionality that was dismissed too soon; I will argue that this strategy is the correct way to understand intensionality generally, and it is vindicated by modern linguistics. In an attempt to ban attitudinal objects from his ontology, Quine [1960, Ch. 6] endorsed a view on which propositional attitudes were not relational, as the surface form of their ascriptions seems to suggest. Rather, on this view, the logical form of propositional attitude ascriptions is that of a complex but unary predication: the material inside the attitude verb forms a unit with the verb. Quine [1956] also gestures toward an analogous view for intensional transitive verbs like “hunt” and “seek”, although he stops short of endorsing it. Such views did not begin

with Quine; rather, they are versions an idea developed by Nelson Goodman [1976] and Israel Scheffler [1963]. Their idea was as follows. Sentences such as “John is hunting a lion” and “John believes that the earth moves” have intensional readings, and on these readings, the expressions in the verb phrase are bound together more tightly than on readings that are fully extensional. Thus Goodman and Scheffler hold that where there is intensionality, there is additional phrasal unity, and they often try to bring out this unity with hyphenated paraphrases, as in “John is lion-hunting” and “John believes-that-the-earth-moves”.

Heightened phrasal unity is a phenomenon that manifests itself in several places in natural language. For instance, we observe it in various forms of compounding, in idioms, and in incorporation. The idea undergirding heightened phrasal unity is that sometimes, for the purposes of semantic interpretation, phrases behave like words; syntactically complex phrases receive interpretations that we would ordinarily expect to provide for morphologically complex lexical items such as compounds. Quine and Goodman’s paraphrases bring out exactly this point: phrases such as “hunting a lion” or “hunting lions” have readings that are best interpreted in the same way as “lion-hunting”, which is a complex lexical item denoting a unified kind of activity.

The way I will argue for this view is by showing that intensional verb phrases such as “seeks a lion” and “needs a meal” are instances of semantic incorporation. Semantic incorporation is the semantic counterpart of the morphological (or syntactic, depending on who you ask) phenomenon in which a nominal in the object position of a transitive verb forms a unit with the verb, and comes to denote a unified activity. Examples of this phenomenon in English are often expressed via compounding with a hyphen, although they can also be fully lexicalized: Mary may apple-pick, horseback-ride, or babysit. As it turns out, many of the criteria for determining whether a transitive verb phrase expresses an incorporated meaning are identical to those for determining whether the construction is intensional. Most notably, the ability for an indefinite in the object-position of a transitive verb to be interpreted nonspecifically is criterial for both intensionality and incorporation, although there are several other overlapping criteria as well.

My proposal is that the nonspecific nominals in the object-positions of ITVs should be treated as verbal modifiers that combine with the verb to form new verbs. The view

I adopt for ITVs is an adaptation of the semantics for pseudo-incorporated constructions given by Dayal [2003, 2011], but incorporates important elements of a closely related semantic proposal for ITVs given by Graeme Forbes [2006]. However, the novelty of my proposal comes from giving an account of an otherwise elusive notion—heightened phrasal unity—and then showing how intensionality is best treated as a special case of one form of heightened phrasal unity: incorporation. The rest of this paper will proceed as follows. In §2, I discuss the traditional proposal, due to Quine, Goodman, and Scheffler, to treat the material in intensional contexts as helping to form a unary predicate, and discuss what their proposal looks like when taken seriously as a proposal concerning the semantics of ITVs. In §3, I lay out and clarify the notion of phrasal unity, and discuss various instances of the phenomenon that we find in natural language. §4 discusses incorporation, particularly semantic and pseudo-incorporation, at length. In §5, I argue that ITVs meet all of the criteria for expressing incorporated meanings. In §6 I make a semantic proposal for ITVs, arguing that intensional NPs act as verbal modifiers, whose semantics should be handled within derivational morphology, and in §7 I consider whether the proposal can be extended to intensional clausal verbs, including the traditional propositional attitude verbs.

3.2 Intensionality and Hyphenation

Suppose that we have a picture of Pickwick, the character from Dickens’ novel *The Pickwick Papers*. On the one hand, the picture represents Pickwick—it is a picture *of* him. But on the other hand, Pickwick does not exist—he is a fictional character, so in another sense we are inclined to say that the picture is not *of* anything at all, at least as far as being *of* something is supposed to be a relation. Goodman [1976, Ch. 1] presents this puzzle as motivation for the following view. When we say that a picture is of Pickwick, and speak truly, what we are in fact saying is that it is a Pickwick-picture. Here it is worth quoting Goodman:

...much as most pieces of furniture are readily sorted out as desks, chairs, tables, etc., so most pictures are readily sorted out as pictures of Pickwick, of Pegasus, of a unicorn, etc., without reference to anything represented. What

tends to mislead us is that such locutions as “picture of” and “represents” have the appearance of mannerly two-place predicates and can sometimes be so interpreted. But ‘picture of Pickwick’ and ‘represents a unicorn’ are better considered unbreakable one-place predicates, or class-terms, like “desk” and “table”. [Goodman, 1976, p. 21]

And Goodman continues:

Some confusion can be avoided if in the latter case we speak of a ‘Pickwick-representing-picture’ or a ‘man-representing-picture’ or, for short, of a ‘Pickwick-picture’ or ‘unicorn-picture’ or ‘man-picture’. [*ibid.*]

Goodman is here claiming that sometimes we have reason to believe that what initially seem like argument places in a relational phrase turn out, instead, to be classifying the activity or object denoted by the head of the phrase. This is in keeping with the common finding that certain NPs, particularly definite and indefinite descriptions, behave predicatively. But Goodman goes further than this. Not only should we treat the expressions in such positions as modifiers; we ought to treat these phrases as semantically unified; these phrases come to function like unary predicates. Goodman captures this proposal with a paraphrase in which he moves the noun to a position in front of the head of the larger phrase, and compounds them using a hyphen.

Scheffler [1963] and Quine [1960] make similar proposals for indirect speech reports and attitude ascriptions. Scheffler’s view is a thoroughly nominalist view, on which sentences such as

(1) Galileo said that the earth moves

are to be treated as in (2):

(2) Galileo spoke a that-the-earth-moves utterance.

According to Scheffler, “that-the-earth-moves” forms a predicate of utterances, serving to classify utterances as of one type or another. He ultimately extends this view beyond ut-

terances to the case of propositional attitudes as well. Quine's view in *Word and Object* is similar, but perhaps even more extreme. Quine originally cites Scheffler's view approvingly, but then proposes to treat (1) as in (3):

(3) Galileo said-that-the-earth-moves.

On Quine's proposal, "said that" is an operator that takes a clause and forms a unary, compound predicate. Thus, Quine's proposal imputes an extremely high degree of unity to the verb phrases of indirect speech reports and attitude ascriptions.

At first pass, you might be tempted to dismiss these proposals as *ad hoc* maneuvers to which Goodman, Scheffler, and Quine resort to avoid commitment to unwelcome intensional entities, to which most linguists are now totally accustomed. But these views were originally dismissed not because they were seen as *ad hoc*; rather, they were dismissed because Davidson [1966, 1967b, 1968] argued, forcefully, that they were non-compositional. However, as we will see below, hyphenation does not indicate that the resulting phrase is unanalyzable, and lacks logical structure altogether. Instead, it indicates that the phrase exhibits increased unity: its structure is different, and tighter, than the structure of other VPs, but this structure—at least in one sense—is still transparent to the grammar.

3.2.1 Intensional Contexts

Before I explain exactly what it is in which this additional unity consists, I want to be clear about the contexts and features that motivate Goodman's introduction of the hyphen in the first place. Goodman is pointing to what we now call intensional contexts, which are characterized by a cluster of peculiar features. First, intensional contexts seem to resist existential quantification: we can't existentially generalize into these positions, whether they contain names, definite descriptions, or indefinite descriptions, and so such contexts are free from existential import.¹ Also, when an indefinite description appears in such a context,

¹In the case of transitive verbs that take phrasal as opposed to clausal complements, the inability to quantify in and bind a variable of in object position is a strong reason to think that the object positions of these verbs do not provide arguments. This indicates that there is a deep connection between intensionality and argument structure, which is a point that is not often or fully appreciated. However, there are ways of quantifying into non-argumentative positions; for instance, we can use second-order quantification to quantify into predicate position, or we can quantify over ways, or make use of other forms of non-nominal

it can be interpreted in two ways: it can be given a non-specific interpretation or it can be interpreted specifically. Third, these contexts have traditionally been seen as resisting the substitution of coextensive expressions, so that substitution of coextensive expressions can yield a change of truth-value.

The distribution of contexts that exhibit these features is varied. Goodman provides the examples of a noun phrase that includes an intensional preposition: “of”, and a single intentional verb: “represents”. But these features are most familiar because they are characteristic of intensionality generally; we typically observe the above features in both the so-called Intensional Transitive Verbs (ITVs) such as “seeks”, “owes”, “needs”, “represents”, and “resembles”, and also in verbs that take clausal complements, such as “believes”, “desires”, “thinks”, and “knows”, among many others. For the remainder of the next two sections, I will limit my discussion to the case of ITVs, but later I will generalize my points to intensional verbs with other kinds of complements.²

The three properties just discussed, which we can call NONEXISTENCE, NONSPECIFICITY, and OPACITY, are often seen as the defining features of ITVs.³ However, while certain paradigm instances of ITVs—such as “seek”, “owe”, and “resemble”—exhibit all three features, the presence of any one of these features does not require the presence of the others. Instead, whether or not these features are present depends partly on which verb and noun phrase the verb phrase comprises. For instance, “need” exhibits NONEXISTENCE and NONSPECIFICITY, but allows for substitution, “worship” exhibits only NONEXISTENCE, while “recognize” exhibits only OPACITY. This has led to some debate as to which properties should be considered criterial for intensionality. In what follows, I will mostly treat the

quantification. See Rayo and Yablo [2001] for discussion.

²There are also a large number of intensional nouns, such as “picture”, “debt”, “search”, and “need”, many of which are nominalizations of corresponding verbs. These nouns also allow for the formation of phrases that exhibit additional unity, and this unity can likewise be captured with hyphenation. However, even though I followed Goodman in motivating the discussion with an intensional noun, in what follows I will mostly confine myself to discussion of intensional verbs.

³Given a sentence of the form **NP V NP'**, we can characterize NONEXISTENCE and NONSPECIFICITY precisely as follows:

Nonexistence: **NP Vs NP'** has a reading which fails to entail **NP' exists**, where **NP'** is upward-entailing.

Nonspecificity: **NP V NP'** has a reading that fails to entail **NP Vs a particular NP'**.

Opacity **NP V NP'** has a reading that fails to entail **NP V NP***, where **NP'**, and **NP*** are extensionally equivalent.

ability for indefinite NPs to be interpreted nonspecifically as criterial for the intensionality of a transitive verb. My reason for this is as follows. NONSPECIFICITY has clear semantic ramifications, and interacts with other aspects of our semantic machinery in ways that OPACITY and NONEXISTENCE do not. Importantly, it has effects on discourse structure, in that it licenses only certain forms of anaphora, and is not easily explained ontologically, by positing non-existent objects or senses. In light of this, it seems plausible that if we can understand how it is that indefinites relate to the verb, and receive nonspecific interpretations, this will go some distance toward helping us understand why no noun phrases within these positions are existence-entailing, and why such contexts resist substitution of co-extensive NPs. Thus, for the remainder of the paper, I'll make the plausible but defeasible assumption that the availability of a nonspecific reading for an indefinite is the most basic feature of ITVs.

The three features just discussed form just a small portion of a collection of interconnected intensional behaviors that we observe within the complements of intensional verbs. For instance, intensional NP complements can be replaced by quantifiers, but only by very particular sorts of quantifiers called special quantifiers [Moltmann, 1997, 2003b, 2004, 2008, 2013]. Special quantifiers are typically combinations of a quantificational determiner together with the morpheme “-thing”, such as “something”, “everything”, and “the same thing”. Inferentially, these quantifiers serve to replace the entire NP complement of an ITV, and they allow for modification, as in the following examples:

- (4) a. John is searching for a unicorn. \Rightarrow
 b. John is searching for something.
 c. John is searching for something magical.

Further, anaphoric reference to the complements of intensional verbs is permitted, but only with a very restricted, nonstandard set of pronouns called special pronouns: “the same thing”, “that”, “what”, and “one”.

- (5) a. John is looking for a horse.
 b. Bill is looking for *it/*him/*her too

- c. Bill is looking for one/that too.
- d. Bill is looking for what John is looking for.

Lastly, we can perfectly well talk about the identity conditions of certain intentional objects, even if they are nonspecific. If John and Mary are both looking for an assistant, there is a sense in which they are looking for the same thing, even if their searches will be satisfied by different entities. Thus there is also a notion of special sameness or special identity. This indicates that each of our ordinary logical notions has a special counterpart in intensional contexts. This parallel points to an important but underappreciated conclusion: our entire apparatus of reference, quantification, identity, and anaphora is mimicked within the complements of intensional verbs, but functions in the absence of any genuine reference or existential commitment. This non-committal logic of intensional contexts extends well beyond traditional questions of intentional identity, but remains largely unexplored.⁴

One last important feature is that constructions involving ITVs have two readings, which sometimes go by the names “specific” and “nonspecific”. For the sake of generality, I prefer to use *de re* for what would be called the specific reading, and *de dicto* for the nonspecific.⁵ The basic contrast is that between (6) and (7):

- (6) John is seeking a horse.
- (7) John is riding a horse.

While (6) and (7) are superficially similar, (6) has both a *de dicto* and a *de re* reading, while (7) has only one reading, the *de re* reading. We can bring out the *de dicto* reading of

⁴This intensional mimicry is extensive. For instance, there are even “special plurals” and “special existence”. As an illustration of the latter, consider the fact that John, when he searches for a unicorn, is not searching for something that does not exist. He is searching for an existent unicorn, not a non-existent one. What good would a nonexistent unicorn do him? Of course, we know that unicorns don’t exist, so he is searching for something that does not exist. But this simply means that there is a special sense of “existence”—a sense that patterns with the rest of the special logical machinery. Special plurals, as we will see later, are plurals that do not have any serious semantic effect: they have this status when they are found on bare plurals in the object positions of ITVs.

⁵However, this is a simplification. Fodor [1970] shows that specificity and opacity can vary independently of each other, and as a result there are in fact four readings of these sentences. In addition to the nonspecific, opaque reading and the specific, transparent reading, there are also what Fodor calls the nonspecific, transparent reading and what Szabó [2010] calls the specific, opaque reading. As I have said, I am provisionally taking nonspecificity to be the core property of ITVs, but I think my proposals provide an illuminating way of dealing with opacity as well.

(6) as follows:

(8) John is seeking a horse—but no particular one.

But this reading is not available for (7):

(9) *John is riding a horse—but no particular one.

When we substitute an empty kind term such as “unicorn” for “horse”, it is only the *de dicto* reading of (6), which is unavailable for (7), that stands a chance of being true. Additionally, while nonexistence and nonspecificity are not always present together, there is reason to think that they are connected: when an indefinite or an empty term appears in an intensional context, the NP will fail to export, and the following inference will be invalid:

(10) **NP Vs Q N** \nrightarrow
Q N is/are such that **NP Vs** it/them. Moltmann [1997]

This inference can be invalid either because exportation entails existence or because exportation entails specificity. When this inference is invalid, it is common to treat the verb as having an intensional, *de dicto* reading that does not imply its *de re* reading, and it also indicates that one reading will license only special quantification, anaphora, and identity, as noted above.⁶

3.2.2 Hyphenation

Goodman’s hyphenation strategy is an attempt to capture the important features of the *de dicto* reading of these constructions, and bring out the underlying source of their intensional behaviors. But how exactly are we supposed to understand the hyphenated paraphrases, and what exactly do they bring out? First and foremost, let’s try to understand the paraphrases themselves. Consider the some paradigm cases of intensional phrases, together with their paraphrased Goodmanian counterparts.

⁶The inference is first used as a test for the intensionality of a transitive verb by Moltmann [1997], but subsequently discarded. Richard [2000] takes the presence of the two distinct readings elicited by the test is taken as criterial for intensionality.

- (11) a. picture of a unicorn
b. unicorn-picture
- (12) a. search for a dog
b. dog-search
- (13) a. need a mechanic
b. mechanic-need

The first thing to notice is that the original phrases are syntactically complex; they have internal syntactic structure. But the paraphrases themselves are compounds, which at least have a claim to being complex morphological items. Presuming, for the moment, that this is correct, the paraphrases turn phrases that themselves have a syntactic structure into lexical items that are, arguably, syntactically simple; intuitively, “unicorn-picture” is a bare nominal and “dog-searched” is an intransitive verb. Instead of syntactic structure, however, the paraphrases have structure that needs to be accounted for by derivational morphology.

This relocation of complexity from the syntax to the morphology will have important effects on semantic interpretation. What will these effects be? First and foremost, the nouns that are compounded with the verb clearly play a different role than ordinary referential nouns or noun-phrases. Rather than contributing an argument to logical form, they serve as adjuncts or modifiers: they specify the kind of search or picture in question. Plausibly, such compounds are going to denote kinds of objects or kinds of events. But if these are lexical modifiers, then the interpretation of the entire compound will be handled by the lexicon, rather than by a syntactically driven compositional semantics.⁷ Insofar as these compounds are interpreted in the lexicon, they will receive a much more unified semantic interpretation than what would have been provided for the original phrase. Such unified interpretations for phrases are instances of what I call “heightened phrasal unity”.

However, it’s important to keep in mind that these paraphrases are just that: paraphrases, and I am not claiming that intensional transitive VPs are intransitive verbs, syntactically speaking. Rather, as I will argue below, these paraphrases capture the relevant

⁷Assuming that we view the lexicon as comprising the rules of derivational morphology. More on this below.

semantic structure of the un-compounded constructions, and express meanings of the same kind. Insofar as the phrases express unified meanings, they exhibit additional phrasal unity. The next section is devoted to spelling out this notion of additional phrasal unity, and examining the ways that it is manifested in natural language.

3.3 What is Heightened Phrasal Unity?

The simplest way of thinking of additional phrasal unity is as follows: sometimes phrases behave like words. Or, more precisely but less concisely, sometimes what seems to be a phrase with complex syntactic structure functions, for the purposes of semantic interpretation, like a single unit. Cases where we observe heightened phrasal unity—including idiomaticity, compounding, and incorporation, all of which I will discuss below—challenge the idea that there is a strict division between how syntactically complex phrases are interpreted and how morphological items are interpreted, and, as we will see below, apparent syntactic complexity will, in certain cases, be best treated as morphological complexity.⁸

Consider first some examples in which phrases actually gain morpho-syntactic unity, and as a consequence are treated as (something like) units for the purposes of semantic interpretation. Mary can go horseback-riding. Bill might be a part of a pay-to-play scheme. Terry and Suzy might be going house-hunting. Angela may be waiting to be done babysitting so she can go whitewater rafting. And Sue and Felicia might realize that this is a chicken-and-egg situation, or that a certain politician has an out-of-touch policy. These examples, at first glance, seem disparate. However, each of them has the feature that what formerly were phrases have gained additional unity and are now interpreted as (complex) morphological items. In most of these cases what was formerly either an argument or an adjunct is compounded with head of the phrase and seems to function as a modifier.

⁸It is important here to distinguish between several different notions of “word”, and clarify what I mean by “morphological item”. With DiSciullo and Williams [1988], I hold that morphology is highly productive. Morphology has rules for the formation of morphological items (what I will call “words”) from morphological atoms. Similarly, syntax has rules for the formation of phrases from syntactic atoms. DiSciullo and Williams [1988] hold that the lexicon is the place for expressions that are not derivable from either syntactic or morphological rules; it is like a jail: it is only for the lawless. These are “listemes”; they are in the lexicon because their meanings must be memorized, and they can be either morphological items or phrases. In what follows, I will use the terms “word” and “lexicon” in a way that differs from theirs; I will use “word” to mean morphological items, whether atomic or complex, and “lexicon” to be the place where such expressions reside.

There are two general ways that we can think of how this happens. On the one hand, we can see it as a syntactically driven phenomenon. It might be that the reason that these phrases seem unified from the perspective of semantic interpretation is because a syntactically complex phrase has been lexicalized, and now serves as a syntactic primitive. But we might also look at the underlying mechanism that is responsible for the compounding, incorporation, and lexicalization that we observe. This mechanism might have to do with facts like frequency, institutionalization, salience, or any number of other semantic, conceptual, and pragmatic factors. On this view, there are various semantic or conceptual factors that push phrases to gain additional unity, and these factors can ultimately result in structures like the ones above. One fact that counts in favor of the latter, semantic options is that we sometimes observe the signs of additional phrasal unity without overt morpho-syntactic manifestation. Such observations are primarily instances of two phenomena: idiomaticity and semantic incorporation. In these cases, we observe the features of additional phrasal unity without overt morphological or syntactic marking, which indicates that additional phrasal unity can be a purely semantic phenomenon. In the remainder of this section I will discuss idiomaticity, and the next section will be devoted to incorporation.

Idiomaticity is perhaps the clearest instance of additional phrasal unity in the absence of morphological marking. Consider a few examples: John might have kicked the bucket or he might just be sawing logs. After he wakes up, he might be shoot the breeze with his friend about some run of the mill happenings. In these examples, what seems to be a syntactically complex phrase gets treated, for the purposes of semantic interpretation, as a single unit. When we say that John kicked the bucket, we mean that he died, while if he is just sawing logs, then we mean that he's asleep. If John is shooting the breeze, it means he's chatting.

The traditional view on idioms is that their additional unity is ultimately due to their syntax: the traditional view was that idioms have no internal syntactic structure, and that their semantics is completely noncompositional; they have been treated, more-or-less, as syntactic primitives.⁹ On this view, what seems, on the surface, to be a phrase with a

⁹For more on this topic, see: Katz and Postal [1963], Katz [1973], Chomsky [1980]Geeraerts [1992], van der Linden [1993].

certain structure in fact is a syntactic unit, and so receives a unified meaning, as if it were a lexical item. Thus there is a mismatch between the apparent structure of the phrase and its underlying syntax and resulting interpretation. Importantly, the resulting interpretation cannot be predicted on the basis of the meanings of the words and the compositional rules that accompany them.

Quine’s view of intensional contexts, which I discussed briefly above, has been interpreted—on my view incorrectly—as an example of this idea. On this way of viewing Quine, intensional verbs, together with the material in their complements, form syntactically unanalyzable units that lack internal structure. Of course, Quine’s proposal is not, technically speaking, a syntactic proposal, since Quine didn’t believe that natural languages had syntaxes. Rather, it is a proposal concerning how we should regiment sentences of natural language in the language of first-order logic, and subsequently provide them with an interpretation. But if we were to view Quine’s proposal as a syntactic one, we could spell out his proposal in terms of syntactic *constituency*. In syntactician’s terminology, the units that are “visible” to the semantics, and interpretable, are called constituents, and constituency is the primary notion in terms of which syntacticians attempt to give an account of phrasal unity.¹⁰ Roughly speaking, a constituent is anything that has a construction history out of syntactic primitives, whatever those syntactic primitives turn out to be.¹¹ Consider the following simple example:

(14) John hit Bill.

We take it for granted that (14) has a particular history of construction: a syntax. This is what syntactic diagrams exhibit. On most syntactic theories, “hit Bill” will occupy a VP node, which is built from the DP “Bill” and the verb “hit”, together with tense. Thus, both “Bill” and “hit Bill” are considered phrases, because they have a history of construction from syntactic primitives. However, on most standard syntactic theories, “John hit” will not be

¹⁰In the generative tradition, at least. In other traditions, such as dependency grammar, constituency is not a basic notion.

¹¹Syntactic primitives are plausibly any of the morphological items discussed above, if we think that there is a strict distinction between morphological productivity and syntactic productivity. They are simply the smallest units that are visible to the syntax.

a constituent of the sentence, and will not be considered a phrase, because it does not have the appropriate kind of construction history. The correct account of constituency groups words that function together into phrasal (or clausal) units, and draws phrasal boundaries in the correct places. To use a common philosophical metaphor, to carve (14) at its joints, our account of constituency must make “hit Bill” a phrase, but keep “John hit” from being one; “hit Bill” is a unit available for semantic interpretation, while “John hit” is not.

The notion of constituency is important when we consider the interaction of syntax with semantics. Ordinarily, semantic interpretation will assign a meaning to each basic constituent, and these meanings will compose in a way that, roughly, mirrors the syntactic composition of the sentence. Each complex constituent will be assigned a semantic value that has been determined by those of its constituents.¹² Thus, ordinarily, the constituency structure—i.e. the syntax—of a sentence will guide, or at least constrain, the relations of determination between the semantic values of the constituents themselves.

Nearly every introductory syntax textbook introduces what are called “constituency tests” [Radford, 2004, Carnie, 2006]. Constituency tests are supposed to be rough-and-ready tests that identify constituent structure. The idea is that certain kinds of manipulations and substitutions show that some groups of expressions form phrases while others do not. When a particular group of expressions fails a constituency test, we can (defeasibly) conclude that the group of words has no construction history, and further, that the group of words does not serve as a self-standing, meaning-bearing part of a sentence.

Thus, if we characterize his view syntactically, Quine holds that “the earth”, as it occurs in “believes the earth moves”, is not a syntactic constituent of the sentence, and so does not receive its own interpretation. The intensional verb “seals off” the material in its complement, and forces the entire VP to function as a unit in logical form. Once this “sealing off” is effected, words inside of intensional predicates are, from the perspective of the semantics, no different from the letters inside of a word. To illustrate, on Quine’s view, “the earth”, as it occurs in “John believes the earth moves” functions analogously to “devas” as it occurs within the predicate “is devastating”: it is not “visible” from the point

¹²Of course, if we countenance a level of representation that interacts with the semantics, which we can call LF, then there may be movement or other operations present at this level that are not strictly speaking syntactic.

of view of semantic interpretation. Only certain units are “visible” to the semantics, and available for semantic interpretation. This is why Quine famously claims that you cannot quantify into intensional contexts; in order to quantify into such a position, the position must be an argument place, which can be occupied by a variable, and argument-places within a sentence must be occupied by a constituent of the appropriate type. But, on a view such as Quine’s, intensional contexts do not have constituent structure. Quantifying into them is like trying to quantify in to the middle of a word.

However, the linguistic consensus on idioms has largely shifted; idioms are now often seen as retaining some internal structure, and having a semantics that is at least partly compositional. This change was due to largely to influential arguments given by Nunberg et al. [1994], who argue that, contrary to traditional wisdom, idioms are not frozen expressions, and they still retain syntactic structure. Instead, they claim that idiomaticity is a fundamentally semantic phenomenon, and that many idioms—for instance, expressions like “pull strings” and “take advantage”—seem to retain some compositional structure, even if they receive nonstandard interpretations. They argue that such expressions are *idiomatically combining expressions*—expressions that are limited in how they combine with others—and that these idiomatic modes of combination are the source of the extra unity that idioms exhibit. On their view, in order to yield idiomatic readings, the expressions within idioms must combine in highly restricted, idiomatic ways, and within certain specific syntactic configurations. Such syntactic and sortal restrictions show that the words within idioms are, in a certain way, “meant for each other”, and that there is something important about that specific syntactic configuration for expressing the idiomatic reading. Any syntactic or selectional variation will make the idiomatic reading unavailable. This idea that some words combine in only certain highly restricted ways in order to yield a particular meaning is one form of additional phrasal unity. So while idioms still possess a history of construction, semantic factors can make certain constructions inflexible, and make it so that certain interpretations of syntactic primitives are only available within a small selectional range, and in certain syntactic configurations. In other words, semantic factors can provide additional phrasal “glue”.

We can see this more clearly if we look closely at how idioms behave when we test their

expressions for constituency. Consider the example “I shot the breeze with Alex”. The following tests are attempts to see whether “the breeze” is a constituent of the sentence.

- | | |
|--|----------------------|
| (a) *The breeze was shot with Alex. | Passivization |
| (b) *It was the breeze that I shot. | Clefting |
| (c) *The breeze is what I shot with Alex. | Pseudo-Clefting |
| (d) *There’s breeze that I shot with Alex. | Expletive |
| (e) *The breeze, I shot with Alex. | Topicalization |
| (f) *I shot it with Alex. | Proform Substitution |
| (g) *The breeze is such that we shot it. | Raising |

As the prefixed stars indicate, “the breeze” fails the tests. But what does it mean to fail such a test? Does it show that “the breeze” is not a *syntactic* constituent of the sentence? Clearly, the results of these tests are not ungrammatical. So it seems reasonably clear that “breeze” is still a syntactic constituent. Rather, in each of the starred cases, the idiomatic reading, which is the one we are interested in, becomes unavailable. That is to say, certain readings appear only in very restricted range of syntactic configurations. If a word fails the tests because the relevant reading becomes unavailable, the word is best interpreted as part of a larger unity. It is only when they are treated as part of that larger unity that the relevant reading is available. We can see these tests as detecting what I will call *semantic constituency*. When we rearrange certain phrases, and run them through constituency tests, we test for whether the word contributes to a certain reading on its own, or whether it is required to be in a certain syntactic configuration to contribute as it does. If the reading is unavailable when we rearrange sentences in this way, the expression in question is not a semantic constituent of the sentence; rather it only has the relevant meaning when it appears with certain other expressions in certain syntactic configurations.

The notion of semantic constituency does not just concern the syntactic configurations on which certain readings are available. Rather, as the idea of an idiomatically combining expression indicates, “semantic glue” can impose certain tight selectional restrictions which affect semantic constituency. For instance, “kick the bucket”, on its idiomatic read-

ing, requires that “kick” combine with “the bucket”; “smack the bucket” does not have an idiomatic reading. Thus idioms are both syntactically and selectionally restricted. This illustrates an important connection between selectional restrictions of words and the degree of unity possessed by phrases in which they appear. The tighter a word’s selectional restrictions, the higher the degree of unity possessed by phrases in which it appears. If, for instance, a reading is only available for a phrase when two specific words are co-present, it indicates that the phrase has an extremely high degree of unity, perhaps that of a frozen expression or idiom. Selectional restrictions are violations of what Gareth Evans [1982, p. 101] calls the Generality Constraint. The idea behind the Generality Constraint is best expressed by Evans himself:

If we hold that the subject’s understanding of ‘*Fa*’ and his understanding of ‘*Gb*’ are structured, we are committed to the view that the subject will also be able to understand the sentences ‘*Fb*’ and ‘*Ga*’, and a common explanation for his understanding of ‘*Fa*’ and ‘*Fb*’ [Evans, 1982, p. 101].

Evans is arguing that a phrase being structured requires something like free recombining of its parts. In order for there to be a genuine joint to carve between *F* and *a*, it must be the case that the subject who understands one can vary each of the parts freely and understand the result, within the constraints of syntax. It seems, then, that limitations on this kind of free recombining indicate a lack of structure: they indicate an increased unity between the parts that fail to freely recombine. One way to preserve Evans’ principle while accommodating such failures is to claim that such failures do not genuinely have subject-predicate structure, and so are not instances of *Fa* and *Gb*. The question then becomes one about what structural features provide the extra unity that ultimately leads to such failures of recombination.¹³

What emerges from this view is that the idiomaticity comes in degrees, and since idiomaticity is an instance in which phrases take on additional unity, phrasal unity comes in degrees. There are degrees to the selectional restrictions on idiomatic readings, and there

¹³Moltmann [2003b, 2004, 2008, 2013] seems to be getting at something similar with her discussion of the Substitution Problem and the Objectivization effect. She seems to be indicating that intensionality is somehow connected to argument structure, and that intensional contexts do not provide arguments. This is not the exact point at issue here, but it is related to the ideas of phrasal unity and semantic recombining.

are degrees of syntactic inflexibility. And moreover, these restrictions can be present even when they are not overtly marked in the syntax or morphology. What is needed is a theory of how such additional unity affects semantic interpretation; I will make a proposal for how this works in section 6 below. But we should keep in mind that neither selectional nor syntactic restrictions need result in overt morphological marking: phrases may gain unity without showing either of these overt signs.

My view is that we should interpret Quine’s proposal in this light, and once we do, it can serve as a plausible view of intensional contexts. Quine’s view is defensible if we allow that there are idiomatic modes of combination and degrees of phrasal unity. Even more specifically, the idea that intensional verbs “seal off” their complements from quantification is exactly to say that they serve as idiomatically combining expressions that serve to forge a higher degree of unity with the expressions in their complements. Davidson’s traditional criticisms of the Quinean proposal, to the effect that it would render intensional contexts noncompositional, only find their mark when we interpret Quine’s view as one which attributes the highest degree of unity to intensional phrases; it is only when a phrase has a very high degree of unity will it lack compositional structure altogether. When we allow that there can be degrees of unity that are lower than that of a lexical item with no structure, but higher than that of a freely recombinable phrase, the idea that intensionality is added phrasal unity allows us a new tool for understanding intensionality. The next section discusses one final form of additional phrasal unity—incorporation—of which, I will claim, intensionality is best treated as an instance.

3.4 Incorporation and Phrasal Unity

Incorporation is another phenomenon in which phrases gain additional unity, and it will occupy our attention for the majority of the remainder of this paper. My claim will be that intensional transitive verbs exhibit nearly all of the semantic features of incorporation structures, and thus are best treated as expressing incorporated meanings. I will then go on to show that, while intensional transitives in general do not display the formal, morpho-syntactic signs of incorporation, they are syntactically inflexible and exhibit selectional

restrictions in a way characteristic of additional phrasal unity.

3.4.1 Syntax and Morphology

Traditionally, the term “incorporation” has been reserved for the process by means of which one of a transitive verb’s internal arguments forms a morphological (or syntactic) unit with the verb [Baker, 1988, Mithun, 1984, Farkas and de Swart, 2003, Borik and Gehrke, 2015]. Most of the debate concerning incorporation has centered on whether this process is lexical or syntactic. In typical cases of incorporation, the object-position nominal undergoes a process of movement and is adjoined to the front of the verb, which in English is often indicated with hyphenation. Linguists who view incorporation as a lexical process tend to treat the result as a compound, while those who treat it as a syntactic process do not treat the result as a compound.¹⁴ Ordinarily, incorporated nominals are stripped of number and case markings, determiners, and all other morphology; that is to say, they are typically bare NPs (\mathbf{N}^0 s) as opposed to DPs. We observed these features in the examples we saw above: horseback-ride, bike-ride, duckhunt, babysit, any apple-pick, although these English examples are just one of several kinds of incorporation.¹⁵ However, more recently it has come to light that the range of nominals that can display the features of incorporated nominals is actually broader than just these bare nominals; for instance, as Dayal [2003, 2011] shows, Hindi can incorporate full NPs. This phenomenon, in which nominals are incorporated that retain some morphological markings and perhaps determiners as well is known as pseudo-incorporation. More on this terminological distinction below.

One of the core, and most stable features of incorporated nominals is that they do not refer to specific objects. Further, verb phrases that incorporate nominals typically denote a unitary activity, and display the characteristics of intransitivity.¹⁶ Here we also observe that just as there can be degrees of transitive, there also seem to be different degrees of

¹⁴Baker [1988, 1996], for instance, treats incorporation as a form of head movement, and thus as a syntactic process, rather than a lexical one. Thus he contrasts incorporation with compounding. However, DiSciullo and Williams [1988] take a lexical approach.

¹⁵See Massam [2001] for a categorization of the different kinds of incorporation. These English examples are instances of what she labels Type 1 incorporation. See also Borik and Gehrke [2015].

¹⁶Incorporation is most often associated with so-called “polysynthetic” languages [Mithun, 1984, 1986, Baker, 1988, 1996], although Mithun [2009] argues that incorporation is not necessary for a language to be polysynthetic.

incorporation; in English, for example, “babysit” is much more fully incorporated than, say, “apple-pick”.

3.4.2 Semantic Incorporation

Recently, however, semanticists have begun studying the semantics of incorporation independently of its syntax and morphology. Semantic incorporation is the name for the semantic analogue of the morpho-syntactic process of incorporation outlined above. The term “semantic incorporation” was coined in seminal work by van Geenhoven [1998] to denote the behavior of obligatory narrow-scope indefinite NPs within the complements of transitive verbs. Van Geenhoven claimed that these structures are best construed as having a semantics on which the indefinite narrows the scope of the verb, as opposed to providing it with a thematic argument. Semantic incorporation can thus occur in the absence of overt syntactic or morphological marking; we sometimes observe it in constructions whose syntax still seems to involve a transitive verb with a morphologically unincorporated direct-object. Carlson suggests that this occurs with English bare plurals that receive an existential interpretation, with English bare singulars, and with weak definites in English. There is also a wealth of cross-linguistic evidence that some transitive verbs that do not morphologically incorporate their internal arguments still express incorporated meanings. Carlson [2006] is particularly interested in the formal variation of structures that can express incorporated meanings. On his view, and in his terminology, the meaningful and the formal bounds of incorporation do not coincide; rather, the former outrun the latter significantly.

Pseudo-incorporation and semantic incorporation are closely related, and are often confused, so it is important for us to be careful with our terms. Pseudo-incorporation is a phenomenon in which a nominal displays some of the features of incorporation, typically all of its semantic features, but may retain morphological and even definiteness markings, and so is not fully morphologically incorporated. This morphological flexibility allows full NPs to be pseudo-incorporated. Pseudo-incorporated nominals also exhibit a higher degree of syntactic flexibility than fully incorporated nominals, and sometimes allow for modification Borik and Gehrke [2015]. However, pseudo-incorporated nominals are still subject to some syntactic and morphological restrictions; for instance, incorporated nominals cannot

be true definites. Thus the term “pseudo-incorporation” is distinct from both traditional, morpho-syntactic conceptions of incorporation, and the purely semantic concept of semantic incorporation. Pseudo-incorporated structures are ones that are semantically incorporated and meet certain less demanding morphological and syntactic requirements than traditional incorporation structures.

The criteria proposed in the literature for detecting incorporation vary quite significantly from author to author. Carlson [2006] explicitly attempts to distill the various formulations and proposals into six standard criteria for detecting semantic incorporation, and then proposes a seventh of his own. Because Carlson is attempting to unify disparate criteria, I will, for the most part, use his criteria in characterizing semantic incorporation. Carlson claims that for a nominal to be considered incorporated, the following six criteria must be satisfied: it must (1) be an indefinite, rather than a definite or quantified type of noun phrase, (2) receive a nonspecific interpretation, (3) be interpreted as narrow-scope only, not showing any interactions with other logical operators, (4) receive an existential, rather than a generic reading, (5) be incorporated into a verb that is a stage-level, rather than an individual-level predicate, and (6) receive a number-neutral interpretation.

Three important classes of nominals that Carlson claims are semantically incorporated are English bare singulars, weak definites, and bare plurals, although the phenomenon occurs even more frequently in other languages.¹⁷

(15) Mark attended class.

(16) Susan rode the train.

(17) Mary picked apples.

In each of these cases, the object-position nominal is interpreted as a nonspecific indefinite; there is no specific class that Mark attended, train that Susan rode, nor apples that Mary picked. Further, none of the object-position nominals display any scoping interactions. No tense or modal operators can scope between the verb and the nominal. The fourth criterion

¹⁷In the case of West Greenlandic, van Geenhoven [1998] classifies all indefinites that take obligatory narrow scope as semantically indefinite.

comes from Carlson’s work on bare plurals. He holds that when a bare plural combines with a stage-level predicate, the bare plural receives an existential interpretation. Mary did not pick apples generally, she picked some apples, and Susan did not ride trains generally; rather, she rode some train. Further, all of the three verbs are stage-level predicates: they are true of a temporal part of their subjects. And lastly, each of the NPs seems to be number neutral; none of the nominals give rise to singularity implicatures, or force any particular numeric reading. However, insofar as (17) is number neutral, it cannot be interpreted as meaning that Mary picked any particular set of apples, even though that is one of its readings. The number-neutrality of each of the examples comes from the fact that the activity is a unitary *kind* of activity.

There are two other features of incorporated nominals that Carlson does not discuss, but that I think are important to round out our set of criteria. One feature that was traditionally proposed to distinguish incorporated nominals from unincorporated ones was that incorporated nominals were “discourse opaque”, they supposedly did not introduce discourse-referents, and could not support discourse anaphora. It has turned out that with respect to discourse opacity, there is significant variation across languages and across types of anaphora, but it is worth noting that with respect to Type 1 incorporation in English, discourse anaphora does not seem to be licensed. For instance, (18) sounds bad:

(18) Mary went apple-picking; they were delicious.

This even seems to be anomalous when we use a bare plural to characterize what Mary did, as in (19):

(19) Mary picked apples; they were delicious.

Here it seems that the pronoun forces us to interpret the bare plural specifically, as indicating that Mary picked some specific apples that were delicious. This indicates that there are two separate ways to interpret statements like (17): a specific way and a nonspecific way. This is notable, since such statements are typically not taken to be ambiguous, and the specific/nonspecific ambiguity is typically taken to appear only in the case of intensional

verbs, a category into which “pick” is typically not seen as falling.

The last feature I will discuss was mentioned briefly above. Incorporation structures typically exhibit a lower degree of transitivity. For instance, consider “bikeride”, an instance of incorporation that is crosslinguistically attested. When the nominal “bike” is incorporated into “ride”, it is not that “ride” exhibits decreased transitivity. Rather, the whole construction takes on the features of intransitivity. But this is also true when transitive verbs take bare singulars, bare plurals, weak definites, and nonspecific indefinites. When Mary is picking apples, the entire VP “picking apples” exhibits decreased transitivity, indicating that it might best be classified as a single syntactic unit: an intransitive verb. Thus, one way of thinking of the extra unity exhibited by VPs that don’t fully incorporate their direct objects is that they are at a middle point between a V + NP structure and a single IV structure. Intransitivity, of course, is not a criterial feature, because the classification of a verb as intransitive in the first place depends on many of the semantic features we mentioned above. However, it is a useful diagnostic.

3.5 Intensional Constructions are Incorporated

I am now going to argue that Intensional Transitive VPs, on their intensional readings, meet all of Carlson’s criteria. This shows that intensional transitive VPs express incorporated meanings, and serve as another category of semantically incorporated construction in English. Subsequently, I will argue that intensional transitives also exhibit syntactic and selectional inflexibility, which makes it plausible that they are best construed as pseudo-incorporation structures. I will then extend my reasoning to other intensional verb phrases, such as intensional NPs like “picture of a dog”, and intensional verbs that take clausal complements, like “believes”, “desires”, and “hopes”.

3.5.1 Semantic Features

Consider several paradigmatic examples of constructions involving intensional transitive verbs:

- (20) John looking for a dog.
- (21) Bill needs the antidote.
- (22) Jill wants a new pair of shoes.
- (23) Javier believes in ghosts.
- (24) William is hunting dodos.

Recall that such constructions have two readings, a *de dicto* reading and a *de re* reading. In the case of (12), John may be looking for a particular dog, or he may be looking for a dog, but not a particular one. Above we called this feature NONSPECIFICITY, and it is present in all of the above examples. For the moment let's focus on the nonspecific reading of these constructions. I will discuss Opacity and Nonexistence below.

Consider the features of semantic incorporation listed above, and then consider, in the case of each of the above examples, whether the nonspecific readings of the examples above exhibit them. First, incorporated nominals must be indefinites; the object-position NP in (12) and (22) are clearly an indefinites, (21) is a weak definite, while (23) and (24) are bare plurals, which are definitely not definites. Second, the indefinite must receive a nonspecific interpretation. This is clearly the case: the ability to interpret the nominals in the object-position of an ITV nonspecifically is characteristic of their nonspecific reading, and is often taken as criterial for the intensionality of a transitive verb. Third, incorporated NPs take obligatory narrow scope. This is also one of the features that is nearly criterial for the nonspecific reading of intensional transitives. First, the readings of intensional transitive verbs are often individuated in terms of the different scopal readings. But the nonspecific reading is invariably the one on which the NP takes the narrowest scope. This is borne out by the fact that the nonspecific reading is unavailable if any operator scopes between the verb and the noun.

Thus far we have seen that ITVs, on their nonspecific reading, exhibit the first three features of semantic incorporation. Given that the term "semantic incorporation" was originally use by van Geenhoven [1998] to characterize the semantic behavior of indefinites that take obligatory narrow scope, it seems that we can already say that intensional transitives

are semantically incorporated. However, I want to show something stronger than just the claim that ITVs incorporate their intensional complements in van Geenhoven’s sense. Consider Carlson’s fourth criterion: on the *de dicto* readings of (12)-(24), do the nominals get existential readings? The definitely do: insofar as any of these exhibit generic behavior, they definitely do not have the ordinary universal reading of a generic. This shows that intensional indefinites meet Carlson’s fourth criterion. Additionally, each of the ITVs above are stage-level predicates, as is nearly the entire class of ITVs.¹⁸ This shows that ITVs meet the first five of Carlson’s criteria.

Additionally, as we saw at the end of section 2, ITVs also exhibit unusual behavior with respect to anaphora. As an illustration, suppose that we were to continue (12) as follows:

(25) John is looking for a dog. It is a dalmation.

Clearly this continuation forces the original sentence to take a specific reading. But on the intensional reading this continuation does not make sense. We can, however, continue it in the following, intensional way:

(26) John is looking for a dog. It has to be a dalmation.

Here the continuation with “it” is itself nonspecific, because “it” is in the scope of another intensional operator, “has to”. In general, nonspecific readings of indefinites in intensional contexts can only ground certain special kinds of anaphora, in which certain kinds of pronouns occur inside of the scope of other intensional operators. As noted above, these are called “special pronouns” To see them in action, consider (27):

(27) John is looking for a dog. Bill is looking for the same thing/one too/what Bill is looking for.

Thus “one”, “what”, and “the same thing” all function as special pronouns. But more importantly, this kind of restriction on anaphora is a form of discourse opacity, and indicates

¹⁸However, one exception is verbs of resemblance. This is in itself an interesting finding, since verbs of resemblance have so far resisted analysis within traditional approaches to ITVs.

that these indefinites do not introduce ordinary discourse referents. This indicates that there is a certain kind of semantic unity on the part of the intensional verb phrases; nonspecific indefinites do not contribute to semantic structure in a way that is fully accessible to referring expressions used later in the discourse. Another point that supports this view is that “that” seems unacceptable. Suppose that intensional complements are semantically unincorporated. Now suppose that they contribute an argument to the verb. We should perfectly well be able to say “Bill is looking for that, too”. But this sounds anomalous. This indicates that nonspecific indefinites do not contribute arguments that serve as the themes of their verbs. These behaviors are easily explained by treating such indefinites as incorporated; for instance, Dayal [2011] provides a semantics for pseudo-incorporation in Hindi that involves theme-suppression, effectively treating the indefinite like a modifier.

Additionally, intensional “transitive” verb phrases with nonspecific indefinites display significant signs of reduced transitivity. On the most widely cited account of transitivity, in Hopper and Thompson [1980], there is a multidimensional scale of transitivity, and intensional transitives with nonspecifically interpreted indefinites in their complements score extremely low on this scale, especially in terms of the individuation of O (the object), the affectedness of O, in terms of agency, as well as several other axes.

I have left the question of number-neutrality for last, partly because it is the property that intensional verb phrases do not seem to exhibit. In other words, there seem to be obvious examples of intensional states directed toward explicitly quantified NPs, such as in the following examples:

(28) John is looking for three dogs—but no particular ones.

(29) John wants three dogs—but no particular ones.

These examples seem to provide genuine cases in which the quantified NP in object position makes an important difference to the attitude, and the truth-conditions of the sentence. And moreover, in the literature on incorporation, it is a nearly universally held that quantified NPs never exhibit the signs of incorporation. Is it plausible to think that these types of intensional verb phrases actually have an incorporated semantics? I think it is plausi-

ble, for the following reasons. First, there are examples of noun-verb compounding that involve numerals, such as “three-dog search”, “three-ball juggling”, “four-hit night” (in baseball). On many views, numerals are not quantificational, but are rather adjectives. If they are adjectives, these numerical compounds come to look much more like traditional N-V compounds, such as “brown-dog search”, “federal-spending oversight”.

Number-neutrality typically concerns quantification, which in turn typically concerns how many things there need to be, for instance, of a certain kind, in order for a certain sentence to be true. But given that the sentences above are intensional, and there need not be dogs or unicorns or anything of the sort, these quantifiers do not have their ordinary force. Why do they not have their ordinary force, you might ask? The traditional answer is that they are in the scope of some kind of intensional operator, which gets rid of their quantificational force. However, my answer is simply that, when they are interpreted non-specifically, these constructions are not genuinely quantificational, but rather contribute to the specification of an activity of a certain kind. They are numerical adjectives that figure into compounds, and they ultimately serve to modify verbs. This, I take it, is a vindication of the idea that intensional ascriptions can be genuinely number-neutral, insofar as number-neutrality requires quantification.

3.5.2 Nonexistence

So far, I’ve made the case that the *de dicto* readings of intensional transitives should be treated as semantically incorporated by showing that they exhibit all of the properties laid out by Carlson. But it might be the case that even if ITVs have an incorporated semantics, incorporation might not be unique to intensional verbs; I have yet to discuss how incorporation explains NONEXISTENCE or OPACITY. For all I have said, intensionality might be a feature that goes well beyond incorporation. In this section I will discuss NONEXISTENCE, and I will discuss OPACITY below when I make a semantic proposal for intensional verbs.

It might seem problematic for my account that there seem to be verbs with an incorporated semantics that are existence-entailing, such as (17). If there are incorporated verbs that are existence-entailing, how can it be that the intensionality of ITVs is explained by

their status as incorporated? But suppose, in light of the above arguments, that (17) has a semantics similar to that of (30):

(30) Mary apple-picked.

This is the consequence of showing that “picked apples” has an incorporated reading. If it does have an incorporated reading, it seems that it must have a semantics like that of (30). However, it doesn’t seem that (30) is existence-entailing; if (30) is an apt paraphrase of (17), in the semantics, the nominal seems to play the role of a classifier or a modifier: it modifies the main verb. I take this structure to be characteristic of incorporated meanings generally. In light of this, it seems perfectly possible for Mary to go apple-picking without there being an apple that Mary picked; she might, for instance, have gone with a group and not actually picked the apples, or they might have turned up and the apples were all too green to be picked.

What we do find is that structures with incorporated semantics are existence-entailing in a different, degenerate way. For instance, (17) entails that Mary engaged in a particular kind of activity, and engaging in such an activity might entail that apples exist. This seems to be the case for constructions like the following:

(31) Joe ate salmon.

(32) Joe salmon-ate.

It might be a fact about eating that there must be some small amount of salmon that Joe ate, but the way that “Joe salmon-ate” entails the existence of salmon is different than the way that “John caught a salmon” does. If there is existential commitment on incorporated constructions, the nonspecificity of the nominal only allows for existential commitment in the way that a mass noun is existentially committing, since the nominal is plausibly bare, and must be interpreted nonspecifically. Consider an analogy. If I babysit, does it follow, in virtue of the form of the sentence, that there is some baby whom I sat? Clearly not. It seems much more plausible that the existential entailments of structures that express incorporated meanings have their existential commitments in an altogether different way, perhaps via a

kind of lexical presupposition. This leads us to the view that no semantically incorporated constructions are existence-entailing, at least in a way with which semantic theory should be concerned. An analogy may be helpful here: perhaps apple-picking necessitates the existence of apples in the same way that drinking water necessitates the existence of H₂O. Semantics has very little to say about this form of existential commitment; such commitment is a metaphysical fact, perhaps to be captured lexically.¹⁹ If such a story is correct, then sentences like (17) are on a par with sentences like “John duck-hunted” in terms of their existential entailments, as long as both express incorporated meanings.

3.5.3 Opacity and Restriction

Carlson’s last criterion is what he calls “restriction”. Restriction is the phenomenon that incorporated meanings can only be expressed by particular kinds of verbs and nominals; not just any transitive verb + noun combination can express an incorporated meaning. Some languages, for instance restrict incorporation to a small number of distinctive verbs. For instance, Pawnee only allows incorporation with three verbs: hit, grab, and burn, while allowing all sorts of variation in the type of nominal incorporated. But on the other hand, there can also be restriction in the type of nominals that can be incorporated, without restriction on the types of verbs that can incorporate them.

Restriction on the nominal is most apparent from the case of near synonyms. Consider the following examples:

- (33) a. Mike went to prison/*penitentiary. Carlson [2006]
 b. The ship is at sea/**ocean/*lake... Carlson [2006]
 c. Mike went to jail/to court/on death row.
 d. The ship is at anchor/at harbor/at shore.²⁰

Carlson’s proposal is that this kind of restriction is due to a form of *semantic enrichment*;

¹⁹This idea is closely related to Bloomfield’s problem, which poses the question of how a semantic theory can provide a genuine theory of meaning without also being required to be a theory of everything. We need to set limits on the aspirations of a semantic theory, if we genuinely hope to achieve its goals. I’m grateful to Zoltán Gendler Szabó for the analogy.

²⁰I’m again grateful to Zoltán Gendler Szabó for the examples in (33-c-d).

the activity in question is a *characteristic* kind of activity, one that may play an important social role, or for which speakers have a well-developed schema. Semantic enrichment, in this sense, should remind us of idiomaticity: there are particular forms of words that are used to describe certain activities, and these forms are used even when there is a synonym in the area that would do just as well from the standpoint of literal meaning. Semantic enrichment binds certain words together for the purposes of describing a notable activity or object. Thus, semantic enrichment is similar to collocation, and may be the result of frequent use, familiarity, or some other conceptual or pragmatic mechanism.

As I understand Carlson, restriction is just a form of additional phrasal unity. Restriction in Carlson's sense is a special kind of selectional restriction in which certain words must be used in order to express a particular kind of meaning, in this case, an incorporated meaning. This phenomenon—that certain words must be used in order to express incorporated meanings—is a phenomenon that we observed above with idioms. In order to express the thought that we ordinarily express with (34):

(34) Justin and Alex shot the breeze,

the very words “shot” and “breeze” are required. We can't, for instance, express the same meaning with a synonym; (35) does not have an idiomatic reading, and sounds nonsensical:

(35) Justin and Alex shot the gentle wind.

This is the same behavior exhibited by the examples in (33). There is a way in which the specific choice of words are essential for expressing the intended meaning, and that meaning can only be expressed when the words co-occur. Thus, the words are “meant for each other”. This is additional phrasal unity at its clearest.

As with idioms, it is helpful to consider this form of added phrasal unity as a violation of the generality constraint. Ordinarily, the generality constraint is seen as connected to selectional restrictions; violations of the generality constraint are supposedly cases where selectional restrictions are not met. But the generality constraint shows something more than this—it shows that certain thoughts can only be expressed with certain forms of

words. It is not that violations of the generality constraint are unintelligible, or result in category mistakes. Rather, they yield meanings that are not the intended meanings. They yield different thoughts than the one that would be expressed had the correct words been chosen. In other words, violations of the generality constraint can show that words combine idiomatically.

Intensional constructions exhibit restriction, and violate the generality constraint. This shows that they both meet Carlson's criterion of restriction, and exhibit the phrasal unity characteristic of such constructions. Consider the following examples:

(36) Bill resembles a dog.

(37) Bill is petting a dog.

Clearly, (37) does not have a nonspecific reading; "pet" is an extensional verb, and as such, there is no reading on which "a dog" is interpreted nonspecifically. On the other hand, (36) does seem to have a nonspecific reading, which we can bring out with "but no particular one". The point that we can take from these kinds of substitution failures is that there is a kind of interaction between indefinites and intensional verbs that yields the nonspecific reading. Neither an intensional verb nor an indefinite on their own are sufficient to yield a nonspecific reading, and fail to do so when they combine with other kinds of expressions. Thus it seems that the nonspecific reading is not predictable on the basis of the ordinary uses of these expressions. Both an intensional verb and an indefinite description are necessary to generate the relevant reading. Further, this reading is fragile. We do not get nonspecific readings for sentences like "Bill resembles some dog", "Bill resembles at least one dog", or "Bill resembles exactly one dog". Thus, substitutions either for the verb or within the complement can make the intensional reading unavailable. This shows that there is a more intimate relationship between an intensional verb and its complement than there is when a verb is extensional.

Moreover, intensional verbs themselves form a small subset of the class of all the verbs in the language. The fact that we can only get incorporated, nonspecific readings with a small class of verbs is itself a form of restriction. In order to get a truly intensional reading,

we must combine both an intensional verb and an NP of the right sort; thus, intensional verbs exhibit restrictions on both the side of the verb and the side of the nominal. This indicates that there are certain kinds of readings that are available only in virtue of the choice of the words themselves; substitution of a synonym is will not even yield a sentence with the same meaning.

3.5.4 Syntactic Features

Constructions involving intensional transitive verbs exhibit syntactic inflexibility, and their direct objects fail many of the constituency tests above. Consider how we might apply these tests to the intensional reading of constructions involving ITV, such as the following: “Socrates resembles a pig”.

- | | |
|--|----------------------|
| (a) *A pig is resembled by Socrates. | Passivization |
| (b) It is a pig that Socrates resembles. | Clefting |
| (c) *There is a pig that Socrates resembles. | Clefting 2 |
| (d) A pig is what Socrates resembles. | Pseudo-clefting |
| (e) *A pig, Socrates resembles. | Topicalization |
| (f) *Socrates resembles it/that/him/her. | Proform Substitution |
| (g) *A pig is such that Socrates resembles it. | Raising |

The issue here is whether the tests listed distort the non-specific reading, or make it unavailable. In many cases, they do. Each time we move the intensional NP from the object position to another syntactic position, it becomes difficult to hear its nonspecific reading. This sounds even worse when we add “but no particular one” to the end of the tests.

What is also notable is the fact that of these tests, the ones that seem to most closely preserve the reading are the ones that make use of special pronouns. For instance, in the case of Pseudo-clefting, the new construction uses the pronoun “what”, which is a special pronoun. Similarly, Moltmann categorizes expletive uses of “it” as special pronouns. The reason this counts in favor of syntactic inflexibility is that special pronouns serve the function of preserving the intensional readings of complements to which they refer—that is

what makes them special. But now note that in the tests which do not make use of special pronouns, the nonspecific reading is almost completely unavailable.

This indicates that the nonspecific complements of intensional verbs have a decreased degree of constituency: the meanings that they help to express are only available within a certain small syntactic range. Similarly to how such meanings are available within small selectional ranges, these tests show that there is some heightened degree of phrasal unity that is due to the particular syntactic configuration of a phrase. This is a form of failure of free recombability. Expressions that freely recombine are ones that have a lower degree of unity—there is nothing idiosyncratic about them, and they are not outliers in terms of the frequency of their use. They lack all semantic and syntactic glue.

3.6 A Semantic Proposal: Lexical Modifiers

There is broad convergence between the semantic approaches to ITVs and the semantic approaches to incorporation. This, in itself, should provide some evidence that the two phenomena are semantically related. In fact, van Geenhoven and McNally [2005] even treats intensional transitive verbs as incorporated, adapting and refining the property account of intensionality proposed by Zimmermann [1993] to overcome several obstacles. Further, both Moltmann [1997, 2008, 2013] and Forbes [2006] propose views of intensional transitive verbs that are very closely related to views proposed for the semantics of incorporation structures; both views can be seen as proposals concerning how an object position nominal comes to form a semantic unit with an intensional transitive verb. More on both of these views below.

My proposal is that we treat the semantics of incorporated constructions, and so also the semantics of intensional transitives, just as we treat the interpretation of noun-verb compounds. Apple-picking is a particular kind of picking. Horseback riding is a particular kind of riding. And searching for a dog is to engage in a particular kind of search. We can create such noun-verb compounds almost at will, but they are only lexicalized in certain cases, when the activity in question becomes notable. But many constructions that are not explicitly compounded can express the same kind of meaning. The correct way to treat these constructions semantically is in terms of verbal modification, and, in the case of intensional

NPs, in terms of noun modification. Thus, nonspecific indefinites in intensional contexts will be modifiers, or adjuncts. But they are not just ordinary modifiers, like “quickly”. Rather, the modifiers in compounds bear a much more intimate relation to the verb than ordinary adverbs; the modifiers used in compounds are lexical modifiers, and help to form a new, unified kind of activity.

Most semantic approaches to incorporation treat incorporated nominals as having denoting properties. For instance, van Geenhoven [1998], Espinal and McNally [2011], and DiSciullo and Williams [1988] all treat incorporated nominals in this way. The question is just what role the property denoted by such a nominal plays with respect to the verb. On several views of incorporation, the property restricts the range of a variable that plays the role of the theme of the activity. For instance, van Geenhoven [1998] treats the incorporated nominal as a property that restricts the range of an existentially bound variable. Chung and Ladusaw [2003] make use of a special predicate modification rule; however, the truth-conditions they derive are identical to those provided by van Geenhoven. The problem with these views, as I see it, is that they fail to do justice to the fact that the readings we seem to be getting are the same readings that we get for fully incorporated structures; if “Mary picked apples” has an incorporated reading, its semantics should be the same as that of “Mary apple-picked”. But the semantics of the latter explicitly involves verbal modification.

The only account that treats the incorporated nominal as a verbal modifier is Dayal [2003, 2011]. Her view is that the relation between verbs like “pick” and “apple-pick” is like that between “cook” and “boil”. Just as every event of cooking involves some manner of cooking, a particular manner-of-cooking verb may suppress the manner argument. I propose that this is exactly how intensional transitives work. Consider a new example. John might be hunting. Just as with cooking, every event of hunting may require a way in which John hunted. Now suppose John is hunting ducks. Its semantics might be exactly like that of “duckhunting”, which now no longer has an argument-place for the manner in which John hunted, nor does it have a theme argument.

On Dayal’s view, the property denoted by the incorporated nominal serves as a verbal modifier, saturating a special, non-thematic argument place reserved for the manner in which the activity is pursued or carried out. In other words, the property allows the

incorporated construction to denote a new type of event that is a subtype of the old. Dayal represents this formally as follows:

- (38) a. $\text{catch}_{TV} = \lambda x \lambda y \lambda e [\text{catch}(e) \ \& \ \text{Agent}(e) = y \ \& \ \text{Theme}(e) = x]$
 b. $\text{catch}_{INC-V} = \lambda P \lambda y \lambda e [P\text{-catch}(e) \ \& \ \text{Agent}(e) = y]$
 where $\exists e[P\text{-catch}(e)] = 1$ iff $\exists e'[\text{catch}(e') \ \& \ \exists x[P(x) \ \& \ \text{Theme}(e') = x]]$

Dayal’s example is that of mouse-catching, but her example can be transposed easily to the examples that we’ve been using. The important thing to notice is that the variable P is in the position of a verbal modifier, and so $P - \text{catch}$ is really a unified lexical item. Dayal then gives her semantics for this lexical item, and it is here that she distinguishes verbs that are existence-entailing from those that are not.

Recall earlier our discussion of whether incorporated verbs were existence-entailing, and consider a few examples:

- (39) a. Mary apple-picked.
 b. Anu mouse-caught.
 c. Mary bike-rode.
 d. Bill babysat.

It is my view that none of these constructions are, strictly speaking, existence-entailing. Strictly speaking, existence-entailingness is a matter of logical form. But consider the logical form that Dayal provides for such incorporating verbs in (39-b). It does not have an argument-place into which an existential quantifier can generalize. Nor, in fact, is there any quantificational material, which is exactly what we would expect from an account of semantic incorporation, since, in receiving an incorporated semantics, the semantic effects of determiners and other morphological markings on the nominal are undone. Thus, if there are existential entailments here, they are a lexical matter. For instance, it might be a fact that events of certain kinds, such as mouse-catchings, must involve mice, but that is a piece of lexical or world-knowledge, and not part of the structure of the semantics for incorporating verbs. This is why Dayal gives the existence condition for “catch” separately, as an

additional piece of information in the lexicon. Structures that are semantically incorporated do not entail existence structurally, for there are no genuine quantifiers or entity-type arguments to bind in logical form.

Thus, distinguishing existence-entailing from non-existence entailing incorporating verbs merely becomes a matter of providing a lexical addition for these verbs, but it is not a fact about the compositional semantics of intensionality or of incorporation. Further, on the standard neo-Davidsonian assumption that a verb denotes a set of events, verbal modification of this sort will yield a subset of those events. This means that monotonicity inferences like “John is searching for a dog, therefore John is searching for a mammal” can be treated as material inferences concerning classes of events. That is to say, the semantic treatment that validates certain inference patterns for intensional transitives is going to be given in terms of the algebra of events.

The final proposal, then, is as follows. Given constructions like

(40) John seeks a dog

(41) Bill needs a massage, and

(42) That’s a picture of grandma,

we ought to construe the nominals inside the scope of the verb as pseudo-incorporated. Most importantly, we ought to construe their semantics on the model of Dayal’s semantics for pseudo-incorporation, where the object-position nominal in fact acts as a modifier of the verb or noun, and gets a unified semantic interpretation. Thus, (40) should be given a semantics like the following:

(43) $\text{seek}_{INC-V} = \lambda P \lambda y \lambda e [\text{P-seek}(e) \ \& \ \text{Agent}(e) = y]$

The question is, what exactly is a P-seeking? And what does it take for an event to be a P-seeking? One way to give a more precise answer, and derive specific truth-conditions is by invoking a suggestion due to Graeme Forbes [2006]. Forbes’s view is in many ways similar to Dayal’s. Forbes also works within an event-semantic framework. Like Dayal, Forbes

holds that on the *de dicto* reading of an ITV, the NP in its object position is non-thematic. Instead, Forbes introduces a special thematic role, called “CHAR”, that serves as a relation between the event and the denotation of the quantificational NP. His core idea is that nonspecific nominals in the object positions of ITVs characterize the activity in question as one of particular sort. But more specifically, Forbes gives an account of characterization in terms of *satisfaction conditions*. The NP in object position specifies what is required for the search (or need, or want, or desire, etc) to be satisfied, and we can use this idea of satisfaction conditions to spell out what it means for an activity to be a P-seeking.

However, giving an account of P-seeking, or of intensional transitive verbs more generally, in terms of success or satisfaction conditions cannot be the correct account generally. Satisfaction conditions, in terms of which nearly all accounts of ITVs have been formulated, seem to only be relevant to a small number of intensional verbs. Consider how could we provide satisfaction conditions for verbs such as: “resemble”, “admire”, “scorn”, “respect”, “sculpt”, “draw”, or many other verbs. What is it for a resemblance to be satisfied? or to be successful? Or admiration? What would it be for a state of admiration to be satisfied? What about a state of scorn? Could we explain scorn in terms of its success or satisfaction conditions? Could we explain what it is to sculpt a bust of Beethoven in terms of success conditions in a way that is informative?

My own approach to intensionality differs in the following way. Dog-searches are different from domesticated-canine-searches. Why? Because they are different kinds of events. Why are they different kinds of events? Well, partly because the agent of the event has different intentions in the two cases. Insofar as the agent has different intentions, this is sufficient to classify the events as distinct. It is not the success conditions that differentiate two intensionally equivalent searches from one another, it is their intrinsic features: features of the events themselves.

3.7 Intensional Clausal Verbs

I have argued that we ought to provide the *de dicto* reading of ITVs with a semantics identical to that of verb phrases that incorporate their nominals. On this view, the inten-

sional NP in the object position of an ITV serves to modify the verb, just as is the case with noun-verb compounds. But since the nominal is not morphologically incorporated, I proposed that this phenomenon is an instance of increased phrasal unity. However, above I also claimed that intensionality generally is an instance of increased phrasal unity, but have said little about intensional clausal verbs: intensional verbs that combine with questions, infinitival clauses, or whole sentences to form verb phrases. What reason do we have to think that they should likewise be treated as instances of additional phrasal unity?

One main reason to think that the intensionality of verbs with clausal complements is relevantly similar to that of ITVs is that all intensional verbs allow special quantification within their complements, but special quantification is not admissible within the complements of extensional verbs. Examples will be helpful here:

(44) John is looking for a red Cadillac. \Rightarrow
John is looking for something.

(45) Jerry expects to become famous. \Rightarrow
Jerry expects something.

(46) Orcutt knows how to tie his shoes. \Rightarrow
Orcutt knows something.

(47) Mary believes that a comet will hit the earth tomorrow. \Rightarrow Mary believes something.

In addition to allowing for special quantification, clausal intensional verbs allow for the whole special apparatus within their complements: they allow special anaphora and identity, but disallow their ordinary counterparts. Consider the following examples:

(48) John is looking for a Cadillac. Bill is looking for one too.

(49) Jerry expects to become famous. Bill expects that too/the same thing.

(50) Mary believes that a comet will hit the earth tomorrow. Jill believes that too/what Mary believes/the same thing.

Moreover, special quantification over clauses, just like special quantification over NPs, patterns with predicative uses of indefinites, allowing modified special quantification:

- (51) Mary is a professor. \Rightarrow
Mary is something impressive.
- (52) John became a nurse. \Rightarrow
John became something noble.
- (53) Jerry expects to become famous. \Rightarrow
Jerry expects something unlikely.

This indicates that there is a connection between certain non-standard argument structures—namely ones where NPs are used predicatively—and the admissibility of special quantification in the whole range of intensional verbs.[Moltmann, 2003b, 2004, 2008].

The fact that special quantifiers are only admissible in the complements of intensional verbs and in cases where descriptions are used predicatively gives us reason to think that there is an underlying similarity that accounts for why such nonstandard quantification is appropriate, and that this similarity is structural. Additionally, the very same constituency tests that serve to show that ITVs are syntactically inflexible serve to show that propositional attitudes are syntactically inflexible. I won't go through all of the tests here, but the conclusion is what you'd expect from the above discussion. In the case of a sentential complement, if there is an indefinite in the complement position of the sentence, the indefinite reading is distorted when it is moved to the subject position, passivized, or manipulated in a way that does not involve a special pronoun. These manipulations, just as with ITVs, manage to leave only the *de re* reading available. Similarly to ITVs, the test for whether quantifiers export from clausal complements serves, in surprising fashion, as a test for syntactic inflexibility.

In light of this underlying similarity, I propose to treat intensional clausal VPs as having the same semantics as phrasal compounds, in which an entire phrase or clause serves as a single modifier. Consider the following examples of phrasal compounds from the British National Corpus:

- (54) a. She also knows that the media tendency to lump together women singersong-writers in a “gee whiz, gosh, women are now making it” syndrome is patronising, if not pernicious. (BNC, A7S190)
- b. Bombay-based Anil put India’s failure to exploit its manpower and mind power and its lack of excellence in sport, economics and the arts down to a “Learn what is there and don’t question it” attitude (BNC, HAE4088)

Carrying these kinds of examples over to the case of intensional clausal verbs, consider the case of belief, and suppose that John believes that the earth moves. I propose to treat the semantics of (55) as given by (56):

- (55) a. John believes that the earth moves
- b. John has a that-the-earth-moves belief.

- (56) a. John wants to walk on the moon.
- b. Jon has a PRO-walk-on-the-moon desire.

Thus, John has a particular kind of belief, or believes in a particular way. Further, just as in the case of intensional transitives, there is additional phrasal unity present, underscored by the fact that the semantics of such constructions is like that of the above kinds of compounds. Insofar as ordinary compositional semantics is not responsible for the semantics of various kinds of compounds, it will not be responsible for deriving the internal semantics of compounds such as the ones in (55) or (56).

We can give a schema for the semantics of ICVs as follows, using belief as a model:

- (57) $\text{believe}_{INC-V} = \lambda P \lambda y \lambda e [\text{P-belief}(s) \ \& \ \text{Agent}(s) = y]$

Exactly analogously to the case of ITVs above, the intensional material within the complement plays the role of a modifier, modifying the verb “believes”, which is a predicate of a state. Roughly speaking, this state treats the material within intensional clauses as categorizing the state as one of a particular kind. The question, then is to provide a semantics for P-belief, which will be a matter for the semantics of compounding.

This view of propositional attitudes is almost exactly the view suggested by Goodman [1951] and defended by Israel Scheffler [1963], although both of their views treat the material in the complement of a propositional attitude verb as a unified predicate of inscriptions. Consider the following propositional attitude ascription:

(58) Smith believes that the earth moves.

Their idea is to treat “that the earth moves” as a predicate of the things that Smith believes-true, namely inscriptions, yielding an analysis like the following:

(59) $\exists x(\text{That}(\text{the-earth-moves})x \ \& \ \text{believes-true}(\text{Smith},x))$

This analysis, of course, is outmoded, insofar as it treats belief as a relation to an inscription. But the core insight is that the material within the complement of the PA verb plays a predicative, rather than a referential role, and forms a tighter unit with the PA verb than do groups of expressions in ordinary contexts.

3.8 Conclusion

I take myself to have begun the process of providing a theory of phrasal unity. However, a full theory of phrasal unity, on which we examine the different ways, and degrees to which, phrases take on additional unity would require a book-length examination. I merely take myself to have pointed to, and described instances of the phenomenon, and then categorized intensionality as special case of one such instance. The idea that intensional phrases exhibit extra unity, and that such phrases should be generally categorized as instances of semantic incorporation is, to my knowledge, a new idea, as is the view that we should treat intensional phrases as having a unified interpretation that is determined compositionally through modification. These ideas, I believe, make it plausible that Quine, Goodman, and Scheffler detected something deep in their initial attempts to account for intensional contexts without resorting to intensional entities; this vindicates their attempt to connect intensionality with argument-structure. Further, in line with their initial aims, this also casts doubt on

the idea that we should pursue the semantics of intensional phrases in the way that we have since Carnap; if intensionality is a matter of phrasal unity and argument-structure rather than a matter of intensional types, it should redirect our efforts in attempting to account for the puzzling features of these structures.

Bibliography

- G. E. M. Anscombe. The intentionality of sensation: A grammatical feature. In Ronald J. Butler, editor, *Analytic Philosophy*, pages 158–180. Blackwell Publishers, 1965.
- D. M. Armstrong. *A World of States of Affairs*. Cambridge University Press, 1997.
- D. M. Armstrong. *Truth and Truth-makers*. Cambridge University Press, 2004.
- Andrew Bacon. Quantificational logic and empty names. *Philosopher's Imprint*, 13(24), 2013.
- Mark Baker. *Incorporation: A Theory of Grammatical Function Changing*. University of Chicago Press, Chicago, 1988.
- Mark Baker. *The Polysynthesis Parameter*. Oxford University Press, New York, 1996.
- George Bealer and Uwe Mönnich. Property theories. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 10, pages 143–248. Kluwer Academic Publishers, 2003.
- George Berkeley. *A Treatise Concerning the Principles of Human Knowledge*. 1710.
- Ned Block. Inverted earth. *Philosophical Perspectives*, 4:53–79, 1990.
- Ned Block. Mental paint and mental latex. In Enrique Villanueva, editor, *Philosophical Issues*, volume 7. Atascadero, CA, Ridgeview, 1996.
- Lawrence BonJour. Epistemological problems of perception. *Stanford Encyclopedia of Philosophy*, 2016.

- Olga Borik and Berit Gehrke, editors. *The Syntax and Semantics of Pseudo-Incorporation*. Syntax and Semantics. Brill, Leiden and Boston, 2015.
- David Bourget. Intensional perceptual ascriptions. *Erkenntnis*, pages 1–18, 2016.
- Franz Brentano. *Psychology from an Empirical Standpoint*. Routledge, London and New York, 1973.
- Miles Burnyeat. Aquinas on “spiritual change” in perception. In Dominik Perler, editor, *Ancient and Medieval Theories of Intentionality*, pages 129–154. Brill, Leiden, Boston, Köln, 2001.
- Alex Byrne. Intentionalism defended. *Philosophical Review*, 110(2):199–240, 2001.
- Greg Carlson. The meaningful bounds of incorporation. In Svetlana Vogeleer and Liliane Tasmowski, editors, *Non-definiteness and plurality*, Linguistics Today 95. John Benjamins Publishing Company, 2006.
- Andrew Carnie. *Syntax: A Generative Introduction*. Blackwell Publishers, Malden, MA, second edition edition, 2006.
- Richard Cartwright. Macbeth’s dagger. In *Philosophical Essays*, pages 13–20. The MIT Press, 1987.
- Victor Caston. Connecting traditions: Augustine and the greeks on intentionality. In Dominik Perler, editor, *Ancient and Medieval Theories of Intentionality*, pages 23–48. Brill, Leiden, Boston, Köln, 2001.
- Gennaro Chierchia and Raymond Turner. Semantics and property theory. *Linguistics and Philosophy*, 11:261–302, 1988.
- Roderick Chisholm. Perceiving: A philosophical study. In David Rosenthal, editor, *The Nature of Mind*, chapter 11. Oxford University Press, 1956.
- Noam Chomsky. *Essays on Form and Interpretation*. North Holland, New York, 1977.
- Noam Chomsky. *Rules and Representations*. Columbia University Press, New York, 1980.

- Noam Chomsky. Language and nature. *Mind*, 104(413):1–61, January 1995.
- Noam Chomsky. *New Horizons in the Study of Language and Mind*. Cambridge University Press, Cambridge, UK, 2000.
- Sandra Chung and William A. Ladusaw. *Restriction and Saturation*. Linguistic Inquiry Monographs. The MIT Press, 2003.
- John Collins. Knowledge of language redux. *Croatian Journal of Philosophy*, 8:3–44, 2008.
- John Collins. Methodology, not metaphysics: Against semantic externalism. *Proceedings of the Aristotelian Society Supplementary Volume*, LXXXIII:53–69, 2009.
- John Collins. Representations without representa. In Piotr Stalmaszczyk, editor, *Semantics and Beyond: Philosophical and Linguistic Inquiry*, Philosophical Analysis. De Gruyter, 2014.
- Tim Crane. *The Mechanical Mind*. Routledge, London and New York, 1995.
- Tim Crane. Intentionality as the mark of the mental. In A. O’Hear, editor, *Contemporary Issues in the Philosophy of Mind*, pages 228–251. Cambridge University Press, 1998.
- Tim Crane. What is the problem of perception? *Synthesis Philosophica*, 20:237–264, 2005.
- Tim Crane. Is there a perceptual relation? In T. Gendler and J. Hawthorne, editors, *Perceptual Experience*, pages 126–146. Oxford University Press, Oxford, 2006.
- Tim Crane. Intentionalism. In Ansgar Beckerman and Brian McLaughlin, editors, *The Oxford Handbook to the Philosophy of Mind*, pages 474–493. Oxford University Press, 2009.
- Tim Crane. *The Objects of Thought*. Oxford University Press, Oxford and New York, 2012.
- Tim Crane. *Aspects of Psychologism*. Harvard University Press, Cambridge, MA and London, England, 2014a.
- Tim Crane. Introduction: In defense of psychologism. In *Aspects of Psychologism* Crane [2014a], pages 1–19.

- Tim Crane. Brentano's concept of intentional inexistence. In *Aspects of Psychologism*, pages 25–39. Harvard University Press, Cambridge, MA, 2014c.
- Tim Crane and Craig French. The problem of perception. *Stanford Encyclopedia of Philosophy*, 2016.
- Donald Davidson. Theories of meaning and learnable languages. In *Inquiries into Truth and Interpretation*. Clarendon Press, 1966.
- Donald Davidson. Truth and meaning. *Synthese*, 17(3):304–323, September 1967a. ISSN 00397857. URL <http://www.jstor.org/stable/20114563>. ArticleType: research-article / Issue Title: Language in Use Including Wittgenstein's Comments on Frazer and a Symposium on Mood and Language-Games / Full publication date: Sep., 1967 / Copyright © 1967 Springer.
- Donald Davidson. The logical form of action sentences. In Nicholas Rescher, editor, *The Logic of Decision and Action*. University of Pittsburgh Press, 1967b.
- Donald Davidson. On saying that. *Synthese*, 19(1-2):130–146, 1968.
- Donald Davidson. Reply to foster. In Gareth Evans and John McDowell, editors, *Truth and Meaning: Essays in Semantics*, pages 33–41. Clarendon Press, 1976.
- Donald Davidson. The inscrutability of reference. *Southwestern Journal of Philosophy*, 10(2):7–19, 1979.
- Wayne A. Davis. *Meaning, Expression, and Thought*. Cambridge University Press, Cambridge, UK, 2003.
- Veneeta Dayal. *A Semantics for Pseudo-Incorporation*. Ms., Rutgers University, <http://www.semanticsarchive.net/Archive/WI2ZjZkM/pseudo-incorporation.pdf>, 2003.
- Veneeta Dayal. Hindi pseudo-incorporation. *Natural Language and Linguistic Theory*, 29(1):123–167, 2011.
- Anna Maria DiSciullo and Edwin Williams. *On the Definition of Word*. Number 14 in Linguistic Inquiry Monographs. MIT Press, Cambridge, MA, 2nd edition, 1988.

- M. Teresa Espinal and Louise McNally. Bare nominals and incorporating verbs in spanish and catalan. *Journal of Linguistics*, 47(1):87–128, March 2011.
- Gareth Evans. *The Varieties of Reference*. Clarendon Press, Oxford, 1982.
- Anthony Everett and Thomas Hofweber, editors. *Empty Names, Fiction, and the Puzzles of Non-Existence*. CSLI Publications, Stanford, CA, 2000.
- Delia Graff Fara. Descriptions as predicates. *Philosophical Studies*, 102(1):1–42, January 2001. Originally published under the name “Delia Graff”.
- Delia Graff Fara. Names are predicates. *Philosophical Review*, 24(1):59–117, 2015.
- Donka Farkas and Henriëtte de Swart. *The Semantics of Incorporation*. CSLI Publications, Stanford, CA, 2003.
- Kit Fine. The non-identity of a material thing and its matter. *Mind*, 112(446):195–234, April 2003.
- Janet Dean Fodor. *The Linguistic Description of Opaque Contexts*. PhD thesis, Massachusetts Institute of Technology, 1970.
- Graeme Forbes. *Attitude Problems*. Oxford University Press, Oxford, 2006.
- J. A. Foster. Meaning and truth theory. In Gareth Evans and John McDowell, editors, *Truth and Meaning: Essays in Semantics*, pages 1–32. Oxford University Press, 1976.
- Dirk Geeraerts. Specialization and reinterpretation in idioms. 1992. Paper presented at the Tilburg Idioms Conference.
- Nelson Goodman. *The Structure of Appearance*. Harvard University Press, Cambridge, MA, first edition, 1951.
- Nelson Goodman. *Languages of Art*. Hackett Publishers, 2nd edition, 1976.
- Stephen Gross. Can empirical theories of semantic competence really help limn the structure of reality? *Noûs*, 40(1):43–81, 2006.

- Ian Hacking. *Why Does Language Matter to Philosophy?* Cambridge University Press, Cambridge, UK, 1975.
- Gilbert Harman. The intrinsic quality of experience. *Philosophical Perspectives*, 4: Action Theory and the Philosophy of Mind:31–52, 1990.
- John Hawthorne and Juhani Yli-Vakkuri. What is narrow content? In John Hawthorne and Juhani Yli-Vakkuri, editors, *Narrow Content*, pages 1–37. Oxford University Press, Oxford, forthcoming 2017.
- Paul J. Hopper and Sandra A. Thompson. Transitivity in grammar and discourse. *Language*, 56(2):251–299, June 1980.
- Ray Jackendoff. *Semantics and Cognition*. MIT Press, Cambridge, 1983.
- Mark Johnston. The obscure object of hallucination. *Philosophical Studies*, 120(1-3):113–183, 2004.
- David Kaplan. Quantifying in. *Synthese*, 19(1/2):178–214, December 1968.
- Jerrold J. Katz. Compositionality, idiomaticity, and lexical substitution. In Stephen R. Anderson and Paul Kiparsky, editors, *A Festschrift for Mirris Halle*, pages 357–376. Holt, Rinehart, and Winston, 1973.
- Jerrold J. Katz and Paul Postal. Semantic interpretation of idioms and sentences containing them. *Quarterly Progress Report of the MIT Research Laboratory of Electronics*, 70:275–282, 1963.
- Ezra Keshet. *Good Intensions: Paving Two Roads to a theory of the De re/De dicto Distinction*. PhD thesis, Massachusetts Institute of Technology, 2008.
- Ezra Keshet. Split intensionality: a new scope theory of *de re* and *de dicto*. *Linguistics and Philosophy*, 33:251–283, 2011.
- Saul Kripke. Vacuous names and fictional entities. In *Philosophical Troubles*. Oxford University Press, 1973.

- Saul Kripke. Speaker's reference and semantic reference. *Midwest Studies in Philosophy*, 2: 255–276, 1977.
- Karel Lambert. *Free Logic: Selected Essays*. Cambridge University Press, Cambridge, 2003.
- Fred Landman. Groups i. *Linguistics and Philosophy*, 12:559–605, 1989a.
- Fred Landman. Groups 2. *Linguistics and Philosophy*, 12, 1989b.
- Richard Larson, Marcel den Dikken, and Peter Ludlow. Intensional transitive verbs and abstract clausal complementation. SUNY at Stony Brook, Vrije Universiteit Amsterdam, 1997.
- Richard K. Larson. The grammar of intensionality. In Georg Peter and Gerhard Preyer, editors, *Logical Form and Language*, pages 228–262. Oxford University Press, Oxford and New York, 2002.
- Richard K. Larson and Gabriel M. A. Segal. *Knowledge of Meaning: An Introduction to Semantic Theory*. MIT Press, Cambridge, MA, 1995.
- David Lewis. General semantics. *Synthese*, 22:18–67, 1970.
- David Lewis. 'tensions. In *Philosophical Papers*, volume 1. Oxford University Press, New York and Oxford, 1983.
- David Lewis. Putnam's paradox. *Australasian Journal of Philosophy*, 62(3):221–236, 1984.
- David Lewis. *On the Plurality of Worlds*. Blackwell Publishers, Oxford, 1986.
- Fiona Macpherson and Dimitris Platchias, editors. *Hallucination: Philosophy and Psychology*. MIT Press, Cambridge, MA, 2013.
- Genoveva Martí. Aboutness and substitutivity. *Midwest Studies in Philosophy*, XIV:127–139, 1989.
- M. G. F. Martin. The limits of self awareness. *Philosophical Studies*, (120):37–89, 2004.
- M. G. F. Martin. On being alienated. In T. Gendler and J. Hawthorne, editors, *Perceptual Experience*, pages 354–410. Oxford University Press, Oxford, 2006.

- Diane Massam. Pseudo noun incorporation in niuean. *Natural Language and Linguistic Theory*, 19(1):153–197, 2001.
- John McDowell. *Mind and World*. Harvard University Press, Cambridge, MA, 1994.
- Alexius Meinong. The theory of objects. In Roderick Chisholm, editor, *Realism and the Background of Phenomenology*, pages 1–51. Free Press, Glencoe, IL, 1904.
- Ruth Millikan. *Varieties of Meaning*. MIT Press, Cambridge, MA, 2004.
- Marianne Mithun. The evolution of noun incorporation. *Language*, 40(4):847–895, 1984.
- Marianne Mithun. On the nature of noun incorporation. *Language*, 62(1):32–38, 1986.
- Marianne Mithun. Polysynthesis in the arctic. In M. Mahieu and N. Tersis, editors, *Variations on Polysynthesis, the Eskaleut Languages*, pages 3–18. Benjamins, Amsterdam, 2009.
- Friederike Moltmann. Intensional verbs and quantifiers. *Natural Language Semantics*, 5(1):1–51, 1997.
- Friederike Moltmann. Propositional attitudes without propositions. *Synthese*, 135:77–118, 2003a.
- Friederike Moltmann. Nominalizing quantifiers. *Journal of Philosophical Logic*, 32:445–481, 2003b.
- Friederike Moltmann. Nonreferential complements, nominalizations, and derived objects. *Journal of Semantics*, 21:1–43, 2004.
- Friederike Moltmann. Intensional verbs and their intentional objects. *Natural Language Semantics*, 16:239–270, 2008.
- Friederike Moltmann. Quantification with intentional and intensional verbs. *draft*, 2013.
- Richard Montague. *Formal Philosophy: Selected Papers of Richard Montague*. Yale University Press, 1974a.

- Richard Montague. The proper treatment of quantification in ordinary english. In Montague [1974a], pages 247–270.
- Richard Montague. On the nature of certain philosophical entities. In Montague [1974a].
- G. E. Moore. *Some Main Problems of Philosophy*. George Allen and Unwin, London, 1953.
- Geoffrey Nunberg, Ivan A. Sag, and Thomas Wasow. Idioms. *Language*, 70(3):491–538, September 1994.
- Terence Parsons. Referring to nonexistent objects. *Theory and Decision*, 11:95–110, 1979.
- Terence Parsons. *Nonexistent Objects*. Yale University Press, New Haven and London, 1980.
- Terence Parsons. *Events in the Semantics of English*. MIT Press, 1990.
- Robert Pasnau. *Theories of Cognition in the Later Middle Ages*. Cambridge University Press, Cambridge, UK, 1997.
- Dominik Perler, editor. *Ancient and Medieval Theories of Intentionality*. Brill, Leiden, Boston, Köln, 2001a.
- Dominik Perler. What are intentional objects? a controversy among early scotists. In Dominik Perler, editor, *Ancient and Medieval Theories of Intentionality*, pages 203–226. Brill, Leiden, Boston, Köln, 2001b.
- John Perry. Possible worlds and subject matter: A discussion of barbara h. partee’s paper ”possible worlds in semantics: A linguistic approach”. In Allén Sture, editor, *Possible Worlds in Humanities, Arts, and Sciences: Proceedings of Nobel Symposium 65*. Walter de Gruyter, 1989.
- Stanley Peters and Dag Westerstahl. *Quantifiers in Language and Logic*. Oxford University Press, 2008.
- Paul Pietroski. On explaining that. *Journal of Philosophy*, 97(12):665–672, December 2000.

- Paul Pietroski. The character of natural language semantics. In Alex Barber, editor, *Epistemology of Language*. Oxford University Press, New York and Oxford, 2003.
- Paul Pietroski. Meaning before truth. In G. Preyer and G. Peters, editors, *Contextualism in Philosophy*, chapter 10, pages 253–300. Oxford University Press, 2005.
- Paul Pietroski. Semantic internalism. to appear in the next edition of the Cambridge Companion to Chomsky, forthcoming.
- Paul M. Pietroski. Character before content. In Judith Thompson and Alex Byrne, editors, *Content and Modality: Themes from the Philosophy of Robert Stalnaker*. Oxford University Press, Oxford and New York, 2006.
- Paul M. Pietroski. Minimalist meaning, internalist interpretation. *Biolinguistics*, 2(4): 317–341, 2008.
- Graham Priest. Objects of thought. *Australasian Journal of Philosophy*, 78(4):494–502, December 2000.
- Graham Priest and Stephen Read. Meinongianism and the medievals. *Australasian Journal of Philosophy*, 82(3):421–442, September 2004.
- W. V. Quine. *Mathematical Logic*. Harvard University Press, Cambridge, MA and London, England, 1940.
- W. V. Quine. *Methods of Logic*. Harvard University Press, Cambridge, MA and London, England, 1950.
- W. V. Quine. Natural kinds. In *Ontological Relativity and Other Essays*, pages 114–138. Columbia University Press, 1969.
- W.V. Quine. *From a Logical Point of View*. Harvard University Press, 1953a.
- W.V. Quine. The problem of meaning in linguistics. In *From a Logical Point of View* Quine [1953a], chapter III, pages 47–64.
- W.V. Quine. On what there is. In *From a Logical Point of View* Quine [1953a].

- W.V. Quine. Quantifiers and propositional attitudes. *The Journal of Philosophy*, 53(5): 177–187, March 1956.
- W.V. Quine. *Word and Object*. Cambridge University Press, 1960.
- Andrew Radford. *Minimalist Syntax: Exploring the Structure of English*. Cambridge Textbooks in Linguistics. Cambridge University Press, Cambridge, UK, 2004.
- Agustín Rayo and Stephen Yablo. Nominalism through de-nominalization. *Noûs*, 35(1): 74–92, 2001.
- Mark Richard. Semantic pretense. In Anthony Everett and Thomas Hofweber, editors, *Empty Names, Fiction, and the Puzzles of Non-Existence*, pages 205–232. CSLI Publications, 2000.
- Mark Richard. Seeking a centaur, adoring adonis: Intensional transitives and empty names. In *Context and the Attitudes: Meaning in Context*, volume 1. Oxford University Press, 2013.
- Richard Routley. Exploring meinong’s jungle and beyond. *Journal of Philosophy*, 80(3): 173–179, 1983.
- Bertrand Russell. On denoting. *Mind*, 14(56):479–493, October 1905. ISSN 00264423. URL <http://www.jstor.org/stable/2248381>. ArticleType: research-article / Full publication date: Oct., 1905 / Copyright © 1905 Oxford University Press.
- Bertrand Russell. *The Analysis of Mind*. Muirhead Library of Philosophy. Project Gutenberg, project gutenber ebook edition, 1921.
- Bertrand Russell. The relation of sense-data to physics. In *Mysticism and Logic*, pages 108–131. Barnes and Noble Books, Totowa, NJ, 1951.
- R. M. Sainsbury. *Reference Without Referents*. Oxford University Press, Oxford and New York, 2005.
- Nathan Salmon. Nonexistence. *Noûs*, 32:277–319, 1998.

- Israel Scheffler. *The Anatomy of Inquiry*. Bobbs-Merrill, Indianapolis, 1963.
- J. R. Searle. The logical status of fictional discourse. In P. A. French et. al., editor, *Contemporary Perspectives in the Philosophy of Language*, pages 233–243. University of Minnesota Press, 1979.
- Gabriel M. A. Segal. *A Slim Book about Narrow Content*. The MIT Press, Cambridge, MA, 2000.
- Wilfrid Sellars. Empiricism and the philosophy of mind. In Herbert Feigl and Michael Scriven, editors, *Minnesota Studies in the Philosophy of Science*, volume 1. University of Minnesota Press, Minneapolis, 1956.
- Ted Sider. *Writing the Book of the World*. Oxford University Press, Oxford and New York, 2011.
- Susanna Siegel. Do experiences have content? In B. Nanay, editor, *Perceiving the World*. Oxford University Press, New York, 2010a.
- Susanna Siegel. *The Contents of Visual Experience*. Oxford University Press, New York, 2010b.
- Scott Soames. Direct reference, propositional attitudes, and semantic content. *Philosophical Topics*, XV(1), Spring 1987.
- Scott Soames. Semantics and semantic competence. *Philosophical Perspectives*, 3 (Philosophy of Mind and Action Theory):575–596, 1989.
- Scott Soames. Truth, meaning, and understanding. *Philosophical Studies*, 65:17–35, 1992.
- Scott Soames. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford University Press, Oxford and New York, 2002.
- Scott Soames. *Reference and Description*. Princeton University Press, Princeton, NJ, 2005.
- Zoltán Szabó. On qualification. *Philosophical Perspectives*, Language and Philosophical Linguistics, 2003a.

- Zoltán Gendler Szabó. Book review - knowledge of meaning: An introduction to semantic theory. *The Philosophical Review*, 106(1):122–124, January 1997.
- Zoltán Gendler Szabó. Believing in things. *Philosophy and Phenomenological Research*, LXVI(3), May 2003b.
- Zoltán Gendler Szabó. The determination of content. *Philosophical Studies*, 148:253–272, 2008.
- Zoltán Gendler Szabó. Specific, yet opaque. In Maria Aloni, Harald Bastiaanse, Tikiu Jager, and Katrin Schulz, editors, *Logic, Language and Meaning: 17th Amsterdam Colloquium, Amsterdam, The Netherlands, December 16-18, 2009, Revised Selected Papers*, pages 32–41. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- Zoltán Gendler Szabó. Major parts of speech. *Erkenntnis*, 80(1):3–29, 2015.
- Zoltán Gendler Szabó and Richmond H. Thomason. *Philosophy of Language*. Oxford University Press, Fortchoming.
- Alfred Tarski. The concept of truth in the languages of the deductive sciences. In John Corcoran, editor, *Logic, Semantics, Metamathematics: papers from 1923 to 1938*. Hackett Publishers, Indianapolis, 1983.
- A. L. Thomasson. *Fiction and Metaphysics*. Cambridge University Press, Cambridge, 1999.
- Charles Travis. The silence of the senses. *Mind*, 113(449):57–94, 2004.
- Michael Tye. The adverbial theory: A defence of sellars against jackson. *Metaphilosophy*, 6(2):136–143, 1975.
- Michael Tye. The adverbial approach to visual experience. *The Philosophical Review*, 93(2): 195–225, 1984.
- Michael Tye. *Ten Problems of Consciousness*. The MIT press, Cambridge, MA, 1995.
- Erik-Jan van der Linden. *A Categorical, computational theory of idioms*. PhD thesis, Utrecht: OTS Dissertation Series, 1993.

- Veerle van Geenhoven. *Semantic Incorporation and Indefinite Descriptions*. CSLI Publications, Stanford, CA, 1998.
- Veerle van Geenhoven and Louise McNally. On the property analysis of opaque complements. *Lingua*, 115(6):885–914, 2005.
- Peter van Inwagen. Creatures of fiction. *American Philosophical Quarterly*, 14(4):299–308, 1977.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford and New York, 2000.
- Timothy Williamson. *Modal Logic as Metaphysics*. Oxford University Press, Oxford, 2013.
- Stephen Yablo. *Aboutness*. Princeton University Press, Princeton, NJ, 2014.
- Edward Zalta. *Intensional Logic and the Metaphysics of Intentionality*. MIT Press, 1988.
- Thomas Ede Zimmermann. On the proper treatment of opacity in certain verbs. *Natural Language Semantics*, 1:149–179, 1993.
- Thomas Ede Zimmermann. Unspecificity and intensionality. In C. Féry and W. Sternefeld, editors, *Audiatur Vox Sapientia. A Festschrift for Arnim von Stechow*. Akademie Verlag, Berlin, 2001.
- Thomas Ede Zimmermann. Monotonicity in opaque verbs. *Linguistics and Philosophy*, 29: 715–761, December 2006.