

**Reference and Grammar:
Reasserting Wittgenstein's Picture**

by

Andrew Tweedie

**A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs
in partial fulfillment of the requirements for the degree of**

Master of Arts

in

Philosophy

**Carleton University
Ottawa, Ontario**

**© 2015
Andrew Tweedie**

Abstract

The aim of this essay is to present a novel line of argument in support of Ludwig Wittgenstein's view of philosophy of language. It consists in three parts. First, it explores a family of referential theories of meaning belonging to the tradition of Millian semantics – a family dubbed 'Millian externalism.' Second, it draws on the work of both Wittgenstein and Noam Chomsky to undermine the tenability of Millian externalism, all the while generating insights into the nature of an adequate semantic theory as it proceeds. And third, from the fact that Millian externalism fails insofar as it purports to be a scientific theory of meaning, the essay concludes that philosophy of language must give in to a theory of meaning which endorses Wittgensteinian quietism. If it refuses, then it must conform to the methodology of linguistics.

Table of Contents

Abstract	ii
Introduction	1
Chapter 1: The Millian Externalists	6
Chapter 2: Relata, Grammars, and Explanations	27
Chapter 3: The Quietist Picture	60
Afterword	91
References	92

Introduction

i. It would be understated, to say the least, if one were to remark that there are evident contrasts between Ludwig Wittgenstein and Noam Chomsky.¹ Whereas Wittgenstein steadfastly rejected the idea that language could be subjected to scientific analysis, Chomsky has arguably been the single most influential individual in modern science of language; and whereas Wittgenstein's hostility for mentalist theories of meaning led him to become the father of semantic externalism, Chomsky is recognized as having been the first to lay a conceptual foundation for a cognitivist approach to linguistics. On the face of it then, there seem to be no two individuals more polarized in their ideas than Ludwig Wittgenstein and Noam Chomsky. It therefore comes as a surprise to be told that followers of Wittgenstein find an unlikely ally in Chomsky in their pursuit of a quietist philosophy of language.²

¹ Unless otherwise stated, references to Ludwig Wittgenstein in this paper will always be made with regards to his later work, beginning in the 1930's with *The Blue and Brown Books* (1958), and culminating in his *Philosophical Investigations* (1953).

² Quietism is the metaphilosophical position which states philosophy has no specialized subject matter, and no positive theses to contribute to our knowledge. It urges that philosophy is not a body of facts or theory, but rather an *activity* whereby intellectual discomforts are resolved by examining our use of concepts and expressions.

ii. In the following pages, we have but one goal: to discern the features of a satisfactory theory of meaning through a critique of referential semantics, and to thereafter propose a theory of meaning which embodies these features. In the course of this project, the ideas of both Wittgenstein and Chomsky feature prominently – not as contrasting counterparts, but rather as complimentary lines of attack against a family of referential theories of meaning which belong to the tradition of Millian semantics. This family of theories, dubbed ‘Millian externalism,’ serves as a foil to the theory of meaning which is proposed at the project’s end. It is therefore through an extensive criticism of Millian externalism that we arrive at this proposal, since an examination of Millian externalism’s faults reveals to us the many requirements which a semantic theory *must* satisfy if it is to be a genuine, explanatory theory of meaning. Incidentally, it turns out that such a theory of meaning must be one which is reconciled to Wittgenstein’s quietist picture of philosophy of language. To understand why this must be the case, we need to discuss the project’s overall structure.

There are three chapters. In the first chapter, we begin with a detailed exploration of Millian externalism, tracing its historical roots and identifying the central characteristics of its approach to semantic content. To put it generally, a referential theory of meaning is ‘Millian externalist’ if its approach to semantic content is extensional, externalist, and anti-individualist. Upon discussing these features at length, we then turn to Millian externalism’s metasemantic foundations, where it soon becomes evident that Millian externalism aspires to be an empirical theory which is continuous with scientific investigation.

It is not until the second chapter that Millian externalism is subjected to scrutiny. Here, we begin with a pair of Chomskian arguments, as presented by Robert Stainton (2006), which argue against the possibility of a scientific theory of meaning constructed from word-world relations. The arguments achieve this conclusion by advancing the claim that referential theories' objects of study are not 'real' in the way which is required for scientific intelligibility, and thus there can be no science which investigates the relations of reference which occur among them. This 'Chomskian critique' effectively opens the door for a second line of argument, one which proceeds from an entirely different direction; issuing from the perspective of common sense and ordinary discourse, the 'Wittgensteinian critique' explores what it is for a theory to be genuinely explanatory, and thereafter concludes that referential theories of meaning fail to explain the phenomenon of meaning. The second chapter then concludes with the observation that whatever theory of meaning we propose as an alternative to Millian externalism, it must be what Michael Dummett (1996) called a 'theory of understanding.' That is, it must explain what it is speakers *know* when they are competent with a language.

The third chapter concludes the project. Proceeding from insights gleaned from the previous chapter's attack on Millian externalism, we begin the third chapter by noting that a satisfactory theory of meaning must reverse the Millian externalist approach to semantic content all the way down. In a word, a satisfactory theory of meaning must be intensional, internalist, and individualist. Upon acknowledging this, it is proposed that a version of conceptual role semantics (i.e., a family of theories which explicate the meanings of expressions

in terms of expressions' use in thought) might satisfy all these requirements. However, before we can even endorse a particular version of conceptual role semantics, a problem arises. Specifically, a more general examination of the causes for Millian externalism's vulnerability to the Chomskian critique (as presented in the second chapter) reveals that whatever version of conceptual role semantics we elect to adopt, it must be one which is stripped of all normative concepts and objects. Simply put, it must be a theory which does not resemble a traditional philosophical theory of meaning in any way. This in turn raises the question of why we should not adopt Chomsky's approach to semantics – a question for which we have no answer. And in lacking an answer, we cannot help but observe that the Chomskian critique we employed against Millian externalism has an unsettling conclusion: philosophy of language ought to become indistinguishable from linguistics.

The third chapter ends with an alternative to this proposed sad state of affairs. Upon noting that Millian externalism was vulnerable to the Chomskian critique only because it simultaneously *i*) attempted to approximate to a scientific theory, and *ii*) invoked the normative categories and concepts of common sense, we arrive at the observation that a philosophical theory of language can evade the Chomskian critique so long as it gives up one of these two aspirations. At this point, I urge the reader to give up the former, since this would allow a philosophical theory of meaning to retain philosophy of language's historic subject matter. This victory, however, comes at a cost: because a philosophical theory of meaning can only dodge the Chomskian critique by confining itself to the framework of common sense and ordinary discourse, it follows that

philosophy of language must reconcile itself to Wittgensteinian quietism.

Namely, it cannot expect to produce discoveries comparable to those of science, but rather can only aspire to make contributions to our understanding. The remainder of the third chapter then ends with a general sketch of a version of conceptual role semantics which would accommodate Wittgenstein's quietism.

iii. In the upcoming pages, I take care not to forecast the conclusions of future chapters, since I wish for the reader to arrive at the project's final conclusion by way of incremental steps. In this sense then, the project proceeds much like an investigation which gradually builds upon itself as time passes, slowly amassing insights and observations until it constructs something worthwhile. I will confess, however, that the project's conclusion was conceived well before it was undertaken.

To elaborate, at the beginning of this introduction I noted various ways in which the ideas of Ludwig Wittgenstein and Noam Chomsky strike hard contrasts. Yet, I failed to mention one other way in which they differ: Wittgenstein was of the opinion that philosophy of language is too scientific, while Chomsky has consistently criticized philosophy for not being scientific enough. On the surface, these appear to be diametrically opposed opinions, and yet the attentive reader must note that although these are mutually-contradicting views, they both proceed from a common assumption. Specifically, they both assume that whatever direction philosophy of language is taking, it is taking the *wrong* direction. And when we come to realize this, the idea that Wittgenstein and Chomsky could be allies suddenly doesn't seem so peculiar after all.

1

The Millian Externalists

1. The term ‘Millian externalism,’ which I have here borrowed from Arthur Sullivan’s essay of the same name (2010), does not denote a strikingly homogenous set of theories of meaning. Undeniably, there are pronounced differences among the following authors, all of whom have advocated for Millian externalist positions: Stephen Boër and William Lycan (1974), Michael Devitt (1981), Howard Wettstein (1986, 2004), Kent Bach (1987), Nathan Salmon (1989, 2004), Scott Soames (1989, 2002), Marleen Rozemond (1993), and Michael Thau (2002). These differences, however, are not our concern. In the present chapter, we will occupy ourselves with an examination of the essential characteristics of Millian externalism – characteristics which definitively mark such theories as opponents to the conception of meaning I will argue for in later chapters. For now, it suffices to say that one is a Millian externalist if they advance a referentialist theory of meaning with an approach to semantic content that is *i*) extensional, *ii*) anti-individualist, and *iii*) externalist (Sullivan, 2010, p. 264). Each of these three characteristics will be explored in greater detail below, along with the metasemantic commitments which motivate them. Before that,

however, it will help us to situate Millian externalism within its proper historical context, and to identify the sources from which it inherits its key features.

2. As its name would suggest, Millian externalism belongs to two closely related movements in contemporary philosophy of language: Millian semantics and semantic externalism.

Millian semantics, also referred to as ‘direct reference semantics,’ originated in the 1970’s following the work of Keith Donnellan, David Kaplan, Saul Kripke, Ruth Marcus, and Hilary Putnam. Summarized generally, it was a countermovement against the Fregean descriptivism which dominated the field at the time, engendering a set of ideas which led to a referentialist turn in semantics. Its initial, central claim was the thesis that certain expressions belonging to natural language – most noticeably proper names, indexicals, and demonstratives – had an exclusively referential function, contributing to the truth-conditions of the sentences containing them nothing more than what they refer to, relative to a context of use (Bianchi, 2012, p. 79). The thesis was then quickly extended to include other types of categorematic words, the first being natural kind terms (e.g., Kripke, 1972; Putnam, 1975); it was argued that although natural kind terms are not technically singular terms like proper names, they nonetheless function analogously to proper names and therefore deserve analogous semantic treatments. Moreover, since this line of argument was generally regarded as plausible, the fact that a referentialist treatment *could* be given to natural kind terms – that is, to terms not traditionally regarded as referential devices – had the effect of sparking interest in the related (but

stronger) thesis that the truth-conditions of any sentence belonging to a natural language could be determined exclusively by its syntactic structure and the references of the terms belonging to it (Bianchi, 2012, p. 80).³ This idea, naturally, was in plain contrast to the traditional Fregean insight that more is required to determine the state of affairs picked out by a sentence.

Now although it is perhaps self-evident, it is worth noting that Millian semantics retained philosophy of language's historic obsession with representation. Millian semantics encompasses a body of relational theories of meaning (in which Millian externalism is included) wherein the semantic value of an expression is determined by the expression's relation to an object or state of affairs in the world; if a word is a purely referential device, such as a proper name, then its meaning *just is* the object it represents. Consequently, while there were obvious, deep-seated differences between Fregean descriptivism and the newly-conceived Millian semantics of the 1970's, there were also obvious ways in which Millian semantics had not significantly departed from the traditional norms of philosophy of language. For one, it was inclined (like most pre-Wittgenstein philosophies) to conflate representation with meaning; it proceeded from the assumption that the question of how a sentence comes to represent something other than itself would have the same answer as the question of how a sentence comes to have meaning. Secondly, as is suggested by the previous point, it left unchallenged the assumption that meaning is a relational phenomenon. As

³ It is perhaps worth noting that although there has been much optimism for the prospects of giving Millian semantics a broader scope of application, Millian semantics continues to struggle with adjectives, verbs, and common nouns other than natural kind terms.

such, it was arguably Millian semantics' sister movement, semantic externalism, which marked a more radical departure from orthodoxy.

3. Like Millian semantics, semantic externalism entered mainstream philosophy in the 1970's. Simply put, it advanced the thesis that meaning is not determined by the contents of a speaker's mental state, but rather is determined by some extrinsic relation (Sullivan, 2010, p. 261). Though he did not coin the term, Putnam is widely recognized as the originator of semantic externalism, having been the first to popularize the idea that meanings "just ain't in the head" (1973, p. 704). At the center of his argument was a thought-experiment in which there is a world – a so-called 'Twin Earth' – which is identical to Earth in every way, up to and including containing molecule-for-molecule duplicates of the inhabitants of Earth (p. 700). Yet, within this thought experiment, Twin Earth differs from Earth in one important regard: in the place of water, it has a substance which is virtually indistinguishable from water, but is nonetheless *not* water since it has a different chemical composition (p. 701). Even so, the inhabitants of Twin Earth still refer to this substance as water, resulting in different extensions of the word 'water' between the two planets. Given all this, from the fact that an Earth speaker who was suddenly and unknowingly transported to Twin Earth would be in no position to determine the extension of his use of the word 'water' (p. 703), Putnam urged that it follows that the speaker's mental state has no bearing on the meaning of the word (p. 704); the speaker would have entered a new environment in which the word 'water' is bound up in new historical and causal relations to items belonging to that

environment. A similar point was argued by Tyler Burge (1979), in which he said the intentional mental content of two molecule-for-molecule identical individuals could differ solely on account of differences between their respective social environments and linguistic communities.

4. Now although most contemporary philosophers are quick to cite Putnam in a discussion of semantic externalism's origins, the truth is otherwise: semantic externalism in contemporary philosophy of language predates Putnam's 'Meaning and Reference' (1973) by at least several decades. In his book *The Magic Prism* (2004), Howard Wettstein rightly identifies Wittgenstein as the father of the 1970's revolution in philosophy of language, since it was Wittgenstein who first championed the externalist claim that language is a public phenomenon which must be understood as a set of shared social practices (p. 95). Additionally, it was Wittgenstein who was the first to earnestly attempt a divorce between meaning and the mental. This much was evident as early as the 1930's, as shown in the so-called 'blue and brown' books (1958).

Wittgenstein's views of language in particular deserve close examination here, since whether Millian externalists know it or not, much of their metasemantic commitments are inherited from Wittgenstein's work. The Wittgensteinian perspective, often unfairly summed up with the simple slogan 'meaning is use,' is incredibly multifaceted – enough that philosophers both past and present have found Wittgenstein extremely difficult to pin down. Even so, we can get at the relevant ideas while sparing ourselves excessive detail.

At the heart of the Wittgensteinian view is the idea that language is an *activity* which must be appropriately situated in broader patterns of human activity if it is to be correctly investigated. Moreover, it recognizes that language is *not* a homogenous body of activity; rather, it is a conglomerate of different activities, some of which share overwhelming similarities, while others exhibit radical differences (Wittgenstein, 1953, §§23, 65, 67). These activities, which Wittgenstein refers to as ‘language-games,’ are inextricably bound up in other, non-linguistic activities, and are consequently embedded in what Wittgenstein called a community’s ‘form of life’ (1953, §§19, 23, 241) – that is, the broader collection of cultural, sociological, and historical activities of the community that arise, in part, from universal human behaviors.⁴ They are therefore neither contrived nor invented: language-games naturally arise from (or in conjunction with, but never independently of) the social activities they are bound up in, and are themselves intrinsically social activities.⁵

⁴ The expression ‘form of life’ is used a mere three times (§§19, 23, 241) in the *Philosophical Investigations* (1953), and only a handful of times in other works. Unsurprisingly then, the scarcity of Wittgenstein’s comments on the matter has led to much speculation as to what exactly he meant by the term – and hence my reluctance to use anything but broad terms to articulate the expression here.

⁵ This is not to imply that Wittgenstein denies the fact that new discourses are, on occasion, invented. What he does deny, however, is the possibility of inventing language-games or discourses in total isolation from pre-existing ‘grammatical’ norms – where ‘grammar’ is to be understood as the cumulative body of rules for determining what moves may be made within a language-game. And since a significant component of a language-game (as will be discussed) is the set of grammatical rules it encapsulates, it follows that language-games are never *truly* invented, even when they come into existence by human design. Rather, they merely trace a lineage to other language-games; a ‘new’ language-game merely introduces a few new expression types and/or modifies the grammars of pre-existing language-games.

Now having gotten a handle on the idea of language-games, it remains to be said how they figure in Wittgenstein's non-relational conception of meaning. Inspired by the mathematical formalists who (contrary to Frege) argued that arithmetic is a game in which contentless signs are used in accordance to combinatorial rules (Baker & Hacker, 2009b, p. 46), Wittgenstein began to toy with the idea that linguistic signs only have meaning in virtue of performing a function within a rule-governed context. Thus he stumbled upon the game analogy: in much the same way that a chess piece cannot be defined by its physical features, so too a word token cannot be defined by its mode of signification (Wittgenstein, 1953, §31). Instead, a chess piece is like a word token insofar as it has significance in virtue of the function it may perform – which, in turn, depends on the ways in which it may be used. Of course, not any given use of a chess piece or word token is indicative of its function. Only *correct* uses are indicative of function, and so one who makes use of a chess piece or a word token does so with meaning *only* if the use occurs within a rule-governed context, and is in accordance with the rules. After all, what is a chess piece taken independently of the game in which it is situated, if not a piece of material completely lacking in any significance? So it is with word tokens: a sound is imbued with meaning only because it performs a function given the rules of the activity in which it is situated – and that activity is the language-game.

So, the slogan 'meaning is use' falls a little short: meaning is not so much the use of an expression as it is the function, where function is determined by the use one makes of it in accordance with the rules of the language-game being played. From this, we can generate a number of theses which characterize the

Wittgensteinian perspective – many of which have either been directly adopted or modified by Millian externalists. They are as follows:

1. The meaning of a speaker's expression is determined by pre-existing linguistic institutions (since the expression's meaning just is its function, and there is no function absent proper usage within a language-game – the linguistic institution).
2. Given (1), semantic intentions cannot alter the established meaning of an expression. (If an individual uses an expression in perfect accordance with established practices surrounding that word – i.e., uses it correctly within a language-game – then the expression means what it standardly means within that game, regardless of whether or not the speaker intends to mean it differently by way of some peculiar mental ceremony he performs when he speaks the word.)
3. It is not necessary that a speaker should be aware of the meaning of an expression in order to use that expression meaningfully, so long as he uses the expression correctly (since a listener who is participating within the same language-game will nonetheless understand the expression, even if the speaker does not).⁶

I will revisit these theses later, at which point I will touch upon their relevance to Millian externalism. For the time-being, it is enough for us to acknowledge that although Millian externalism encompasses a set of theories

⁶ Such situations might arise, for example, when a student feigns understanding in a technical discourse merely by adopting the jargon of the discourse and repeating the forms of expression common to it, without quite understanding what it is he speaks.

which belong to Millian semantics, these theories equally belong to the tradition of semantic externalism, of which Wittgenstein – the great enemy of representation – is the father. Now having painted Millian externalism in broad strokes, we now turn to a discussion of the specific characteristics of Millian externalists’ approach to semantic content, thereby bringing Millian externalism into better focus as a distinctive set of positions within their broader theoretical traditions.

5. In his essay ‘Millian externalism’ (2010), Sullivan’s main point of contention with Millian externalists is their underlying assumption that “a rejection of an individualist or internalist approach to reference entails a weak sufficient condition for singular thought” (p. 246).⁷ Here, however, we are concerned neither with the conditions for singular thought, nor with Millian externalist approaches to the subject. For our purposes, it is sufficient to note that while Millian externalists are genuinely concerned with mental processes outside of semantics, Millian externalist conceptions of meaning are wholly devoid of any serious considerations for the properties of speakers – let alone for speakers’ mental states.

We begin to get a sense of this when we examine the Millian externalist’s approach to a well-known Fregean principle on cognitive significance: if two statements differ in cognitive significance, then the two statements merit different semantic treatments. This is the case since, as Sullivan puts it, a

⁷ One such example is Nathan Salmon’s claim that merely looking at a class-enrollment list is sufficient to enable singular thoughts about enrollees with whom one is unfamiliar (Salmon, 2004, p. 254).

“difference in cognitive significance is sufficient proof of a difference in [semantic] content” (2010, p. 249). The principle is so central to Frege’s work that it is outright presupposed in the opening lines of ‘Sense and Reference’ (1892):

Identity gives rise to challenging questions which are not altogether easy to answer. Is it a relation? A relation between objects, or between names or signs of objects? In my *Begriffsschrift*, I assumed the latter. The reasons which seem to favor this are the following: $a=a$ and $a=b$ are obviously statements of differing cognitive value; $a=a$ holds *a priori* and, according to Kant, is to be labelled analytic, while statements of the form $a=b$ often contain valuable extensions of our knowledge and cannot always be established *a priori*. (p. 209)

The Millian externalist denies the principle. While the Fregean would not regard the statements ‘the morning star is the morning star’ and ‘the morning star is the evening star’ as semantically equivalent, since the former is a logical tautology which cannot be doubted, while the latter is a truth of fact which might plausibly be denied by a rational agent, Millian externalists *do* regard the two statements as semantically equivalent. As will be discussed further on, this has the effect that otherwise competent language users are often not in a position to recognize synonymy between statements, even when the statements are used by the same speaker with different semantic intentions. For now, that Millian externalism would regard the expressions ‘the morning star’ and ‘the evening star’ as semantically equivalent is indicative of at least one obvious feature of its semantic program: it endorses an *extensional* approach to semantic content.

For the Millian externalist, what constitute the contents of our thoughts and utterances are not abstract intensions, as Frege would suggest, but rather real-world extensions. (Thus, with respect to the classic debate between Frege and Russell, the Millian externalist sides with Russell in believing Mont Blanc *is* a constituent of the proposition ‘Mont Blanc is over 4000 meters high.’) Moreover, as was already indicated by Millian externalism’s stance on the ‘morning/evening star’ example, this extensional approach to content forces the Millian externalist to endorse a weak notion of semantic equivalence; since ‘the morning star’ is a referential expression, it is semantically equivalent to any referential expression which likewise picks out the same entity – in this case, the planet Venus. This weak notion of semantic equivalence is found, for example, in Boër and Lycan (1974), who proceed from the assumption that two sentences are semantically equivalent if they have *i*) the same logico-grammatical structure, *ii*) predicates which denote the same properties, and *iii*) singular terms (including definite descriptions)⁸ which refer to the same objects.

⁸ While Boër and Lycan are inclined to treat definite descriptions as singular terms, a Millian externalist can coherently deny that definite descriptions and singular terms deserve analogous semantic treatments. (For example, one can follow Russell in thinking definite descriptions are to be analyzed as sets of quantified general statements.) In fact, it is worth noting that there is widespread disagreement as to how definite descriptions should be approached, not merely among Millian externalists, but among Millians generally. One such issue that divides Millians concerns Donnellan’s distinction between referential and attributive uses of definite descriptions (such that a definite description can be used to designate an object either referentially or denotatively, depending on the speaker’s mental state) (1966). Millians disagree on whether or not such a distinction properly belongs to semantics.

In similar fashion, Millian externalism's extensional approach to semantic content also commits it to *anti-individualism*⁹ – the thesis that speakers do *not* have first-person authority over the references of their expressions. Because Millian externalism takes an extensional approach to content, there arise cases in which a speaker may unknowingly refer to an object despite having the intention to refer to something else altogether. For example, if a Millian externalist were inclined to treat definite descriptions as singular terms (though many of them don't), and therefore regard definite descriptions as referential devices, then it is possible that although a speaker may intend to refer to Christopher Columbus with the definite description 'the man who discovered North America,' his definite description would nonetheless refer to someone else entirely. After all, the extension of 'the man who discovered North America' is not Christopher Columbus, but rather (in all likelihood) some tenth-century Viking. Similarly, it is possible on the Millian externalist view that a speaker may express the same thought (or 'state of affairs,' depending on the preferred terminology) on two separate occasions, without being in a position to realize that she has done so

⁹ For Millian externalists, their extensional approach to semantic content commits them to anti-individualism *only* because of their referential realism – that is, their belief that the referential expressions of natural language can, and do, refer to the ordinary objects which populate our daily lives. When referential realism is coupled with an extensional approach to content, anti-individualism *must* be the case, since no individual has an infallible knowledge of the things to which they refer. Here, I do not mean to suggest that anti-individualism is *logically* entailed by an extensional approach to semantic content, since there are ways in which individualism can be reconciled to an extensional approach. For example, if one were inclined to believe that our referential expressions do not refer to the ordinary objects of daily life, but rather refer exclusively to sense-data, then one will have advanced an extensional theory of meaning in which speakers *do* have first-person authority over their references, since sense-data are private objects about which one can never be mistaken (in principle).

(e.g., Wettstein, 1986, p. 197). For Millian externalists, it is not merely possible, but rather quite plausible, that one should be capable of using an expression meaningfully without actually knowing the meaning of that expression (i.e., the expression's extension).

Now having dealt with its extensional and anti-individualist characteristics, we are left with only one remaining characteristic of Millian externalism to consider – namely, its *externalism*. As was already stated earlier, externalism is the thesis that meaning is not determined by anything belonging to a speaker's mental state, but instead is determined by some extrinsic relation. Typically, this extrinsic relation has been articulated as a causal-historical chain, following Kripke's *Naming and Necessity* (1972). One such example is Kaplan's claim that "language itself carries meanings" through chains of transmission, so that when a listener encounters a token of a referring expression, the listener is thereby put into sufficient contact with the referent (through the chain) so as to enable her to refer to that referent and even entertain thoughts about it *through* the language (1989, p. 603).¹⁰ In a nutshell, externalism holds that a speaker is capable of referring to an object not because of any internal state or act on the part of the speaker, but rather because there are institutions of natural language (e.g., linguistic conventions) whereby the speaker's terms refer.

¹⁰ Though Millian externalists have made extensive use of the sort of idea expressed here, Kaplan himself is not a Millian externalist. Moreover, it should be pointed out that Kaplan has since moved on from the idea that linguistic expressions carry meanings, and is now inclined to adopt a form of subjectivist semantics, following Donnellan. See: Kaplan, D. (2012.) An idea of Donnellan. In J. Almog & P. Leonardi (Eds.), *Having in mind* (pp. 122-175). Oxford: Oxford University Press.

Here, the connection between Millian externalism and Wittgenstein is self-evident. For Millian externalists, it is typically a causal-historical chain which is the linguistic institution in virtue of which a referring expression is made meaningful – an institution made possible by conventions whereby speakers resolve to use referring expressions in keeping with past uses. For Wittgenstein, the linguistic institution is the language-game: words have meanings because a language-game is constituted by conventions governing appropriate usage, and speakers render their utterances meaningful only in virtue of abiding by these conventions, whether or not there is a deliberate intention on their part to do. Both in the Millian externalist and Wittgensteinian picture, no excessive theoretical baggage is required; meaning is a linguistic phenomenon which can be exhaustively explained by way of pointing to other properties and features of language, without any mention of speakers and their inner mental episodes. This leads to some impressive similarities between Millian externalism and the Wittgensteinian perspective. Whatever the differences between them (and God knows there are many), Millian externalists nonetheless find themselves endorsing deeply Wittgensteinian claims: meaning is determined by institutions which are independent of individual speakers, even if said institutions are generated from the social behaviors of speakers; individual semantic intentions do not alter the established meanings of expressions; and lastly, a speaker may use a word meaningfully even if she has no knowledge as to what the meaning is. So on the one hand, Millian externalism has been defined by its refusal to follow Wittgenstein in his opposition to relational, representational theories of meaning – but on the other, Millian externalism has been profoundly shaped (albeit,

probably inadvertently in most cases) by Wittgenstein's contempt for mentalist theorizing, as well as by his social picture of language.

This brings us to a complete account of the three defining characteristics of Millian externalism's approach to semantic content: its extensionalism, anti-individualism, and externalism. Of course, *why* Millian externalists adopts this triad is another matter entirely. Since one's approach to semantics critically depends on the expectations one places upon a semantic theory, answering the question of why Millian externalism takes on these characteristics will require a discussion of its metasemantic views. This is the issue to which we now turn.

6. As was noted at the beginning of this chapter, there are pronounced differences among the theorists who have been collected here under the label 'Millian externalism,' and so it comes as no surprise that views on metasemantic issues vary among them. But even so, a small sample can prove strongly indicative of general metasemantic commitments among Millian externalists. In what follows, an exploration of two Millian externalist approaches to two separate metasemantic issues will shed light on the underlying considerations which motivate Millian externalists to adopt the extensional, anti-individualist, and externalist approach to semantics which was discussed above. These metasemantic issues concern the places (if any) that must be afforded to *i*) cognitive significance and *ii*) semantic competence within a theory of meaning.

Earlier, we made a point of noting that Millian externalists shrug off Frege's principle that a difference in cognitive significance between two statements is proof of a difference in semantic content, and therefore should

result in different semantic treatments. In ‘Has Semantics Rested on a Mistake?’ (1986), Wettstein provides a reason as to why Millian externalists are entitled to reject Frege’s principle. To put the matter generally, his argument is that much of the attack levelled against Fregean descriptivism by the ‘new theorists’ (Wettstein’s term for theorists belonging to the tradition of Millian semantics) has consisted in showing that although Frege’s picture of language better accounts for problems concerning cognitive significance, his picture simply cannot be correct since it is incompatible with actual linguistic practices (Wettstein, 1986, p. 201). This suggests a simple point that many of the new theorists have failed to properly acknowledge in Wettstein’s view: Frege and the new theorists have proceeded from entirely different conceptions of how semantics ought to be done.

Frege was principally concerned with the language-thought relation. As such, he considered it a testament to the plausibility of any theory of language that it could avoid generating the sort of epistemic puzzles which were created by referential approaches to meaning. In his paper, Wettstein provides three examples of such puzzles: *i*) the problem of non-denoting terms; *ii*) the problem concerning the informational content of co-referring expressions; and *iii*) the problem of action. Problem (*i*) arises in the case of empty names such as ‘Vulcan.’ If an individual asserts, “Vulcan is a large, heavenly body,” the new theorist seems incapable of accounting for the presence of the determinate belief which corresponds to the statement, since ‘Vulcan’ refers to nothing, and so there is no complete proposition which serves as the object of the belief (p. 187). Problem (*ii*) was already demonstrated earlier in our discussion of the morning

star and evening star. Because ‘the morning star’ and ‘the evening star’ are co-referring expressions, the new theorist is committed to saying they express the same proposition – despite the fact that the two sentences obviously correspond to two different beliefs and therefore express different thought-contents. And lastly, problem (iii) concerns the inability of singular propositions to satisfactorily figure in explanations of human action and behavior. To use Wettstein’s example, let us imagine that I see a man about to be mugged – but unbeknownst to me, I am witnessing the scene by means of a series of mirrors, and the intended mugging victim is actually myself (p. 188). If I were then to say, “He is about to be mugged,” I would be expressing the same proposition as if I had said, “I am about to be mugged.” But, the problem here is the beliefs expressed by these two statements typically lead to dramatically different actions, which gives us the unsavory conclusion that differences in human behavior and action are not reflected at the level of the objects of belief – that is, at the level of propositions.

These are the sort of epistemic troubles which new theorists, on Wettstein’s view, have been duped into trying to accommodate within the confines of their anti-Fregeanism – an unfortunate state of affairs which Wettstein credits to the new theorists’ Fregean upbringing. For Wettstein, what the new theorists ought to have realized is that Frege’s picture of language is incompatible with actual linguistic practices *because* his picture focuses on the language-thought relation. The new theorists, on the other hand, have been generally more successful in describing actual linguistic practices because they strive for a direct picture of the language-*reality* relation. They strive to uncover the “institutionalized conventions in accordance with which our terms refer” (p.

201). Consequently, there was never any reason for believing that in their endeavours to explore our linguistic practices, new theorists were supposed to have also captured the ways in which speakers think about their referents (p. 201). This much should already have been evident in the ways new theorists have approached contemporary semantic issues. For example, if one looks at the three competing accounts for the word ‘that,’ one finds a debate about linguistic practices instead of a debate about how speakers cognize the referents of ‘that’: *i*) the reference of ‘that’ is the unique individual that stands in the appropriate causal relation to the utterance; *ii*) the reference of ‘that’ is the unique individual the speaker has in mind when using the utterance; and *iii*) the reference of ‘that’ is the unique individual that is indicated by the cues that are available to a competent and attentive addressee (p. 202). Wettstein therefore concludes that “an account of linguistic meaning is no longer to be seen as an account of anything like what the competent speaker understands by his terms, but rather as an account of the practices he has mastered” (p. 204). As such, Frege’s principle concerning cognitive significance simply has no place within the province of semantics.

In much the same way that Wettstein responds to issues of cognitive significance by jettisoning them from semantics, so too does Soames (1989) argue for the exclusion of issues pertaining to semantic competence – that is, issues concerning the things in virtue of which speakers can match sentences to their contents. The argument is relatively straightforward. To begin with, Soames assumes that at the very least, a semantic theory must “specify the principles by which sentences represent the world” – and that is just to say that regardless of

whatever else we might expect a semantic theory to do, it must at the very least characterize truth-conditions (pp. 576-577). This follows from *i*) the “central semantic fact about language” is that it carries information about the world, and *ii*) that there is no information without truth-conditions, even if information is not limited to them (p. 575). This latter point is explained as follows: a sentence represents a world as being a particular way, imposing conditions which must be satisfied if the world is to conform to the way it is represented, and so the information encoded in a sentence determines the conditions under which it is true (p. 576).

Once he establishes that a semantic theory must invariably include theories of truth-conditions to account for language’s representational nature, the argument then takes an extremely peculiar turn. Having just assumed the antithesis of the Wittgensteinian view of language (i.e., he assumes that language’s principal function is to represent), he then proceeds in Wittgensteinian fashion to articulate and criticize the ‘Augustinian picture’ of semantic competence.¹¹ He claims that within the Augustinian picture of language, semantic competence arises *from* semantic knowledge (i.e., knowledge of sentences’ truth-conditions); individuals are capable of forming beliefs and entertaining propositions prior to language-acquisition (Soames, 1989, p. 589), and it is in virtue of this capability that language-acquisition is even possible at

¹¹ The expression ‘Augustinian picture’ is in reference to the opening passage of Wittgenstein’s *Philosophical Investigations* (1953), in which Wittgenstein takes a quotation from St. Augustine to be indicative of the sort of conception of language which he hopes to dissuade us from adopting.

all. Simply put, one learns a language by having epistemic access to a set of objects or properties, and then learns which sentence to pair with that set.

Yet, this picture is hopelessly false. The difficulty with it, according to Soames, is that certain objects are only apprehended via linguistic mediation; we do not epistemically access an arithmetical function, for example, prior to language-acquisition (p. 590). Instead, the two are simultaneous. As such, it follows that the correct direction of explanation is from semantic competence to semantic knowledge. Soames therefore arrives at the conclusion which Wittgenstein advanced nearly fifty years earlier: knowledge of truth-conditions is neither sufficient nor necessary for understanding a language.

What does this mean for semantic theory? In sum, there is no obligation upon semantics to explain semantic competence. Since the purpose of semantics is to pair sentences with the propositions they express, it *cannot* explain semantic competence because semantic competence is, in the order of things, *prior* to the knowledge provided by semantic theory. Soames adds that this is no defect of semantics, since the expectation that semantics could explain semantic competence by way of semantic knowledge belongs to an inadequate conception of the role that language plays in our cognitive lives (p. 591) – that conception being the Augustinian picture.

7. To conclude then, both Wittgenstein and Soames begin with a central assumption concerning the purpose of semantics (respectively: to uncover the institutionalized conventions whereby our terms refer, and to characterize the representational information encoded in sentences), and then proceed to show

why this purpose is incompatible with any additional expectation placed upon semantic theory. I will have more to say about this strategy in later pages. For now, both Wettstein and Soames provide us with an idea of how Millian externalism looks to immunize itself from attacks against its semantic program, and perhaps provides us with an idea as to why its approach to semantic content is vigorously extensional, anti-individualist, and externalist. Additionally, their respective metasemantic positions reflect a feature of Millian externalism which I have so far left unmentioned, though it might have been obvious to anyone who has had prior familiarity with the scientism which saturates contemporary philosophy of language: Millian externalism regards itself as an empirical enterprise which is continuous with natural science. Although it is perhaps an obvious and trivial detail, it will be helpful to remember this in the following chapter when we investigate possible avenues of criticism against Millian externalism, particularly ones which target its referentialist foundation.

2

Relata, Grammars, and Explanations

8. Perhaps the most notable arguments against Millian externalists are ones which do not explicitly target Millian externalism, but rather are broad-sweeping criticisms of contemporary philosophy of language's endeavour to articulate word-world relations. These arguments derive most famously from the work of Noam Chomsky, though a minority of philosophers of language have advanced similar views. One such example is Robert Stainton (2006) who argues for the broad conclusion that a theory of meaning should not, and cannot, simply consist in "assigning external, real-world objects to names, nor sets of external objects to predicates, nor truth values or world-bound thoughts to sentences" (p. 913). In what follows, we will provide a brief examination of some of these Chomsky-inspired arguments, since they level legitimate criticisms which a scientific theory of meaning should earnestly strive to avoid. (Such criticisms will be taken into consideration in the upcoming chapter when I motivate an alternative theory of meaning to Millian externalism.) Two arguments in particular, advanced by Stainton, will give us the general flavor of the Chomskian critique – though we shall only endorse one of the two possible, mutually-exclusive conclusions which he offers.

This chapter, however, will not be explicitly focused on the Chomskian critique. It features merely as a part of a single, broader line of argument. Once it has been presented, our principal focus will lie with a Wittgenstein-inspired critique of my own making – one which shows that the value of articulating relations of semantic reference consists solely in elucidating the ‘grammars’ of words and, by extension, ways in which speakers think about the objects which correspond to referential expressions. (The word ‘grammar’ refers here to something distinctive from what we would ordinarily regard as grammar. It will be explained further below.) When coupled with a certain naturalistic principle concerning what it is for a description of x to be an explanation of x , the argument has the effect that a semantic theory is obligated to explain the determinate ways in which speakers conceptualize objects. I will articulate this principle at a later time. For now, we turn to Stainton’s arguments, both of which target referential theories’ objects of study.

9. A few philosophers have taken it as self-evident that natural languages such as English, German, and French simply do not exist, since all there *really* is are people and their various written and acoustical products.¹² These philosophers, however, are a minority. Such a view has struck a hard contrast against the general consensus within philosophy of language that ordinary forms of expression used to describe language are representative of legitimate categories for a theoretical and technical analysis of language.

¹² For example, Donald Davidson (1992).

Chomsky and like-minded philosophers do not fall within this consensus. As Chomsky says, “It is not that ordinary discourse fails to talk about the world, or that the particulars it describes do not exist, or that accounts are too imprecise. Rather, the categories used and principles invoked need not have even loose counterparts in naturalistic inquiry” (2000, p. 138). The idea here is that concepts employed in common sense explanations of natural phenomena need not carry over into scientific investigations of the same phenomena, since many of these concepts cease to be of any use once we move to a deeper level of analysis. For example, it is of no concern to chemistry how people decide whether something is tea or merely water, no concern to biology how people decide at which point in the transition from simple gasses to bacteria that something can be deemed alive, and no concern to the earth sciences whether a river which has changed its direction can still be regarded as the same river (Chomsky, 2000, p. 139). In short, the Chomskian takes it as self-evident that in the investigation of mind and language, one should not feel constrained by the suggestions of ‘folk theorizing’ on the subject (Stainton, 2006, p. 915). Instead, one can – and should expect to – put aside common conceptions, and use specialized terms which may not be faithful to their ordinary discourse counterparts (if they even have any). This is, after all, simply what scientific investigation does. Additionally, one should not entertain *a priori* restrictions on the evidence that can be brought to bear in naturalistic inquiry. As Stainton says, anything that “can be got is fair game” (2006, p. 915).

The notion of semantic reference is among those concepts which belong to ordinary discourse and folk theorizing on language, and one which Chomskians

are inclined to believe has no place in a naturalistic investigation. Stainton provides us with several arguments against the validity of reference as a scientific category, though here we shall only consider two. The first works to undermine reference by bringing into doubt the idea that reference is a relation which transpires between *real* relata. Specifically, the argument contends that if reference is a relation between public linguistic items (e.g., names, predicates, sentences) and worldly items (e.g., external objects, sets, world-bound thoughts), then there can be no such thing as reference if one of these two sets of relata simply don't exist. Here, it is the reality of public linguistic items which is brought into question. Stainton begins by articulating the thesis that the way ordinary discourse individuates natural languages from each other does not actually correspond to real empirical patterns (p. 918). One such example concerns Mandarin and Cantonese: though they are regarded in ordinary discourse as being two different dialects of Chinese, in reality they are more distinct from each other than Swedish and Danish, despite the fact that Swedish and Danish are regarded as distinct languages. Similar considerations apply to dialects. In short, we find that common sense notions of language do not carve nature at its joints: what we call 'natural languages' and 'dialects' are not so much determined by robust linguistic divides as they are by "colonial history, similar writing systems, canonical works of literature, present military might, arbitrary national boundaries, religious differences, and so on" (p. 918). Instead, the Chomskian urges, the only real entities worthy of study are individual idiolects.

This poses a problem for reference. If our common forms of expression about natural languages and dialects do not correspond to objective differences in

kind, then how does one go about individuating words among languages? If the English language is not objectively real, on what grounds can one say ‘hammer’ and ‘nail’ are both English words when they have no property in common except being words which are commonly regarded as belonging to English? It will not help to respond that English is a symbolic system held in common by a community, since here one has no means of individuating the community from others except on the grounds that the community speaks English – the existence of which is the very thing that has been brought into question (p. 919). Once we forego the idea that there are such things as natural languages, nothing permits us to say that ‘foTAHgrafer’ and ‘fotoGRAFer’ are *really* the same word, albeit different pronunciations of the word ‘photographer’ (p. 918-919) – not if we are striving to make our determinations reflect objective reality. Because determinations of the above sort come down to matters of mere decision, the Chomskian is led to believe that there simply are no such things as public linguistic items. And if there are no public linguistic items, then there is nothing which stands in the determinate relations to worldly items that were supposed to constitute reference.

Stainton has other points in service of the above argument, though we need not discuss them here. It is enough to give the sense of how such an argument proceeds, and how it threatens the sort of semantic projects discussed in the previous chapter. If the above argument is correct, then Wettstein’s conviction that semantics should strive to uncover the institutions of natural language whereby our words refer is obviously precluded since *i)* the idea of an ‘institution of natural language’ is rendered unintelligible once it becomes clear

that our intuitions about natural languages do not reflect objective differences in kind, and *ii*) there can be no investigation into the means by which our words refer if, to put it crudely, there are no such thing as words as common sense describes them. Similarly, Soames's argument that semantics must at the very least articulate truth-conditions also collapses, since the Chomskian argument also casts doubt on the idea that there is anything real (i.e., the public linguistic items) to which objects, sets of objects, and predicates are supposed to be assigned.¹³

10. Stainton's second argument takes an altogether different approach. Rather than argue for the non-reality of things without which there can be no relation of reference, he instead argues that while such things as public linguistic items *are* real, they aren't real in the way scientific investigation requires. The argument crucially relies on the claim that common sense and science reveal the world to us in different ways, a claim for which we have two observations in support. First, this is indicated by the differences between the types of concepts deployed between common sense and science (p. 923). To gloss over the matter, the concepts deployed in common sense are not learned, but rather innate; are universally possessed by every non-pathological human; have complex internal structures which suit them for practical application in day-to-day matters; and

¹³ There are other popular assumptions within philosophy of language which are destabilized once we give up on the idea that our common sense concept of a natural language corresponds to any determinate entity. For example, Putnam's notion of the division of linguistic labor (in which the meanings of terms are fixed given a community's deference to the authority of experts who alone know the criteria whereby a particular term's reference is determined) critically relies on the notion of a public language.

lastly, are laden with implicit references to normative and intentional notions, and are therefore not well-suited to describing the world as it ‘really’ is independently of human interests (p. 923). In sharp contrast, scientific concepts are learned (if at all), rather than acquired; not universally shared, even among non-pathological humans; have contents which are austere, rather than rich; and lastly, they are intended for description of the world as it is, independent of human concerns (pp. 923-924). The second observation concerns differences in standards of evaluation – that is, norms which determine what may count as evidence or justification. Stainton notes that within science, such standards are social artifacts which are collectively determined over time, whereas in common sense, standards of evaluation are innately specified (p. 924).¹⁴

Assuming that concepts are the means by which the world is rendered intelligible, the discrepancy between the conceptual apparatus of common sense and science has the effect that each is *blind* to various objects which might be seen from the perspective of the other. So for example, while science can make sense of “quarks, tectonic plates, genomes, and many other things,” it cannot make sense of normative categories made intelligible within common sense – things such as “good wine, liveable cities, or well-prepared osso buco” (p. 924). Even more notably, science cannot make sense of “not-explicitly normative yet mind-dependent things like clouds, tea, desks, sunsets, breakfast cereal, and hockey scores” (p. 924). The argument therefore runs something like this: the objects which are made intelligible exclusively from the perspective of common

¹⁴ This point also speaks in favor of Stainton’s earlier claim that science is entitled to disregard common sense and to progress beyond the categories of common sense.

sense *do* in fact exist, but they exist as social constructs which are so highly dependent on human perspectives and interests that they *cannot be made to be intelligible to science*. Since both ordinary objects and public linguistic items belong to this category of things which cannot be made intelligible to science, it therefore follows that a scientific reference-based semantics is impossible. The relata of semantic reference are real, yes, but “nevertheless not real *in the right sort of way*” (p. 926).

Here again, I should hope, it is obvious how the sort of metasemantic commitments explored in the previous chapter have much to lose in the face of the above argument. As we discussed, both Wettstein and Soames employ a strategy in which they identify a minimal commitment which a semantic theory *must* meet, regardless of whatever else we might expect it to do. Thereafter, they proceed to show how rival metasemantic positions put expectations on semantic theory which are incompatible with this minimal commitment. But the Chomskian position advanced by Stainton undercuts such a strategy by outright denying that there even can be such a minimal commitment of the sort identified by Wettstein and Soames. It is not to say that one cannot identify the institutions whereby our expressions refer, or that one cannot articulate the truth-conditions of sentences; rather, Stainton’s point is that such endeavours can never amount to a genuine scientific theory of meaning.

11. It is important to recognize that Stainton’s two arguments provide us with conclusions which are not mutually compatible: either public linguistic items aren’t real, or they aren’t real in the way which renders them intelligible to

science. And yet, the fact that Stainton even bothers to present two arguments with two separate conclusions should strike us as peculiar. After all, the observations brought to bear in the first line of argument can actually be subsumed under the second. For example, in service of the first argument (which concludes that public linguistic items aren't real), Stainton observes that the way we individuate and identify public linguistic items in ordinary discourse does not correspond to any robust linguistic divide. However, it does not follow from this that public linguistic items are not real. Rather, all that seems to follow from this observation is that public linguistic items do not correspond to anything real *on the level of observation which is of interest to science* – and this, very obviously, is a point in favor of the second conclusion that public linguistic items *are* real, albeit not real in the way which makes them visible to scientific analysis. It seems then that Stainton ought to have presented a single unified argument in favor of the second conclusion.

Additionally, there is one other consideration which speaks in favor of endorsing the second conclusion over the first. Recall Chomsky's following claim: "It is *not* that ordinary discourse fails to talk about the world, or that the particulars it describes do not exist, or that accounts are too imprecise. Rather, the categories used and principles invoked need not have even loose counterparts in naturalistic inquiry" (2000, p. 138). If Stainton agrees with Chomsky on this matter (and presumably he does), then surely he believes ordinary descriptions of language are indeed about *something*. But if this is the case, an awkward question arises: how can he even make sense of the idea that ordinary discourse is about anything at all if there are no such things as public linguistic items?

What would ordinary descriptions of language be about if not about names, sentences, and so on? So, while I have presented both of Stainton's arguments, I have done so only because *i)* I have had an aim towards textual fidelity, and *ii)* the observations employed in support of the first conclusion actually lend support to the second conclusion. From here on out, however, it is the second conclusion which we will hold as true: public linguistic items are real, albeit not real for science.

12. For the remainder of this chapter, I would like to discuss an altogether very different line of criticism against Millian externalism, but one which is nonetheless continuous with the Chomskian critique presented above. This second line of criticism is comparable to the Chomskian critique insofar as it targets referential theories of meaning generally. However, it does differ in one important regard: while the Chomskian critique attacked referential theories' attempt to approximate to scientific theories by criticizing the validity of their objects of study, this new line of argument attacks referential theories of meaning from within the conceptual framework of ordinary discourse and common sense. It argues that articulating relations of reference *does* have value – albeit, not of the sort which referentialists think it has. This in turn leads to the conclusion that reference is only of interest to semantics insofar as the study of reference sheds light on what Wittgenstein called 'grammars.'

Let us consider a statement of the following form: the word x refers to object y . What is the meaning of such a statement? Assuming that a listener comprehends a sentence by understanding its meaning, one way in which we

might answer the question is to discern what exactly it is that is understood when someone is confronted by a statement of the above form. Here, I think, it will not suffice to say that the listener understands that the word x picks out, designates, denotes, or represents the object y , since such an explanation simply transfers the burden of explanation from one term to another, and gets us no further; if we want to know what it is that an individual understands when they learn that a particular word refers to a particular object, it does not advance our knowledge of the situation to simply substitute the word 'refer' for any one of the many synonyms available to us. Even worse than this sort of explanation would be something like the following: the listener understands that the word x contributes to the truth-conditions of the sentences containing it nothing more than what it refers to, relative to a context of use. Such an explanation, after all, makes explicit appeal to the concept of reference – a concept which presumably plays a part in the listener's understanding of statements of the form 'the word x refers to object y ,' which is precisely the thing we were trying to explain. A more plausible explanation would be this: the listener understands that there is a relation that obtains between the word x and the object y . The explanation is obviously deficient, since the fact that there is a relation between x and y is made obvious by the statement's logical form alone – and what we are after is the listener's understanding of the statement's content. But even so, we might flesh the explanation out a bit, and articulate what kind of relation reference consists in. We might, for example, provide the sort of causal-historical explanation of reference that Millian externalists are fond of. And yet, this too poses a problem. When an individual is told 'the word x refers to object y ,' it plainly is not the case

that what she comprehends is the fact that there are causal-historical relations between x and y (assuming she has not been afflicted by prior philosophical reflection on the matter). Ordinary language speakers simply do not attach the kind of theoretical baggage which philosophers do to such ordinary concepts as reference.

The question of what is understood (and therefore what is meant) by statements of the form ‘the word x refers to object y ’ is perhaps better approached if we consider what kind of effect such a statement would have on a competent listener. And here, I think, it is obvious that the effect of telling an individual ‘the word x refers to object y ’ is that the individual thereby acquires a competence with the word x , so long as the individual has a sapient awareness of the object y (i.e., she is aware of that object as being what it is).¹⁵ And given that this is the effect of telling someone ‘the word x refers to object y ,’ I would urge that what is understood by such statements – and therefore, what is meant by them – is a rule concerning the appropriate use of the word x .

Now, the attentive reader might believe that the above line of thought cannot constitute any sort of knockdown argument against Millian externalism, since it presupposes that in the investigation of meaning, one ought to take into

¹⁵ The concept of sapience is aptly illustrated by Robert Brandom (1994), who distinguishes sapience from sentience (p. 5). Sentience is what we share with nonverbal animals such as cats. It simply concerns the capacity to be aware, and therefore consists in reliable differential responsiveness to an environment, perhaps through arousal or irritability. To be sapient, on the other hand, is to be subject to normative considerations and the authority of reasons. As opposed to arousal or irritability, it concerns understanding and intelligence. Unsurprisingly then, Brandom connects sapient awareness with conceptual awareness, and I have followed suit. Also, I will have more to say about the role which sapience plays in apprehending meaning.

account speakers' mental states. (In this case, we have taken into consideration what it is they understand when they encounter a statement of the form 'the word x refers to object y .') As such, it may seem that I have begged the question if not against referential theories of meaning generally, then at least against Millian externalism. But this is not the case. In light of the Chomskian critique presented in the preceding sections, there is no longer any reason whatsoever to think that Millian externalists can study meaning without consideration for individual speakers. After all, that idea was only intelligible if we took seriously the idea that public linguistic items could be legitimate objects of scientific study – but that idea is jettisoned if we accept either of the conclusions of Stainton's two arguments (i.e., public linguistic items aren't real, or they aren't real in the way required for scientific analysis). Because we have accepted the second conclusion, and because Millian externalism does strive to engage in an empirical project which is continuous with natural science, then Millian externalists *must* accept that individual speakers are to be taken into account, for otherwise it is not entirely clear what else they should study. Given this, it should now be apparent why I have been so bold as to mention the names 'Chomsky' and 'Wittgenstein' on the same page: the Chomskian critique creates the gap through which we may thrust the point of the Wittgensteinian critique.

Now the idea which I am attempting to bring into focus is illustrated in Wittgenstein's view of the function of ostensive definition within language. On the Augustinian picture of language to which Wittgenstein was opposed,¹⁶

¹⁶ We discussed the Augustinian picture of language in the previous chapter under our discussion of Scott Soames's (1989) approach to semantics.

ostensive definition serves to connect language up with reality: it is responsible for “injecting content” into language (Baker & Hacker, 2009a, p. 7), thereby imbuing the lifeless signs of language with intentionality (p. 10). Furthermore, it holds that ostensive definition may occur in one of several ways,¹⁷ but however it happens, what is ostensively pointed to just is the meaning of the word for which the ostensive definition is given. Now on the radically different view proposed by Wittgenstein, ostensive definition does not serve to connect reality with language. Rather, an ostensive definition remains within language, since the object designated in ostensive definition is a sample which belongs to the means of representation, and not to what is represented (Baker & Hacker, 2009a, p. 7).

Naturally, this claim requires some unpacking. To begin with, Wittgenstein noted that pairing words with the objects they represent could not constitute any sort of foundation for language (i.e., could not be the source of language’s content), because it is always possible for an ostensive definition to be variously interpreted, even though they rarely are (Wittgenstein, 1953, §28).¹⁸

¹⁷ In the case of public ostension, one points to an object and says, “That is a z,” thereby connecting the object up with the word z. This differs from private ostension, in which one – by way of some mental act or process – connects an inner, mental entity to a word.

¹⁸ To get a sense of how it might be possible to misinterpret an ostensive definition, consider how one might teach a child what the word ‘two’ means by way of ostensive definition (Wittgenstein, 1953, §28). For example, imagine pointing to two nuts and saying, “That is called ‘two.’” Wittgenstein notes that while this is a perfectly appropriate way of defining the number two, there is a way in which it can be made questionable as to how the number two can be defined this way at all. After all, the individual to whom the ostensive definition is given doesn’t know what it is that the definition-giver wants to call ‘two.’ It is perfectly possible that the child may, for example, assume that the word ‘two’ is the name given to that particular set of nuts. Similarly, it is not logically impossible for the child to take the proper name of a person,

Wittgenstein took this observation to be quite telling. To illustrate what he was after, consider the following: if I point to a wall which has been painted over with many shades of blue, none of which have determinate boundaries since each shade blends into the next, you will nonetheless grasp my meaning if I then say, “This is called ‘blue.’” You will *not* interpret me to mean a particular shade to be distinguished from the rest. What this suggests is something the Augustinian picture cannot account for: there is something *prior* to naming and ostension in language which constrains listeners’ interpretations of ostensive definitions, thereby ensuring the success of said definitions. This prior ‘something’ is the set of rules surrounding the usage of the words for which we give the definitions, and it is in virtue of our antecedent comprehension of these rules that we never radically misinterpret ostensive definitions, even though it would be possible to do so in principle.

These rules are encapsulated in what Wittgenstein called the ‘grammars’ of words. In Wittgenstein’s peculiar, idiosyncratic use of the term, ‘grammar’ does not refer to the rules of syntax. Instead, it has much a broader connotation: the grammar of a word is the much wider network of rules which surround the appropriate use of the word, setting out when it may or may not be employed within a language-game. In short, a word’s grammar is the set of norms which determine how the word may be used meaningfully. So, while a sentence such as Chomsky’s famous ‘colorless green ideas sleep furiously’ is well-formed and does not violate any specific syntactic rule, the sentence is nonetheless nonsensical

of which an ostensive definition is given, to mean a color, race, or even a point on a compass.

since it violates what Wittgenstein regarded as the grammatical rules of the words contained within. To say of an object that it is both green and colorless is to commit a ‘grammatical mistake.’

Let us return to the subject of ostensive definitions. If ostensive definitions do not connect language with reality, do not inject semantic content into words, then we can only make sense of their function if we regard them as having a normative role within language. Specifically, ostensive definitions serve to “guide linguistic behavior by providing standards of correctness for the use of the expressions whose meanings they explain” (Baker & Hacker, 2009c, p. 88). In a word, they *show* words’ grammars. It therefore helps us to think of the objects pointed at in ostensive definitions not as non-linguistic items outside of language, but rather as linguistic samples; if I point to a rock and say, “That’s a rock,” the object itself is a part of the means by which I show the rules governing the appropriate usage of the word ‘rock.’ The crucial upshot of this view is that the statement ‘that’s a rock’ is not any sort of empirical or descriptive proposition, since it is not truly about any given state of affairs. Rather, it’s a metalinguistic proposition. Or, as Wittgenstein would have preferred to say, it is a *grammatical* proposition.

Hence the thesis which seemed incomprehensible earlier: an ostensive definition remains within language, since the object designated in ostensive definition is a sample which belongs to the means of representation, and not to what is represented. What is represented is the grammar of the word that has been explained by way of ostensive definition. And moreover, not merely the linguistic statement occurring in the ostensive definition, but also the object

designated, are the means by which that grammar is represented – both of which ought to be considered as belonging to language.

Now earlier I asked the reader to consider the meaning of a statement of the following form: the word x refers to object y . I asked that the reader do so by considering what it is that a competent listener comes to understand through such a statement. And in turn, I asked the reader to determine what the listener comes to understand by further considering what kind of practical effect the statement would have on her. From the fact that a listener who is told ‘the word x refers to object y ’ will have thereby acquired a competence with the word x (assuming she possesses sapient awareness of object y), I suggested that the meaning of such a statement is a rule concerning the use of the word x . Given that, it will perhaps now be obvious as to what I have in mind. The function of articulating relations of reference is analogous to the function of ostensive definition in language. In much the same way that the statement ‘that is a rock,’ when accompanied by an ostensive pointing, is not an empirical proposition but rather a grammatical proposition, so too does a statement of the form ‘the word x refers to object y ’ express a proposition which is about the ways in which one ought to use the word x . It is *not* a proposition about any sort of tangible relation that obtains between x and y , for if it were, it would be impossible to make sense of the fact that competent listeners apprehend no such relation when they hear the statement, either tacitly or otherwise. The purpose of articulating relations of reference consists in functionally classifying the use of the linguistic items which stand on the one side of the relation. In a word, a sentence which identifies a relation of reference merely serves to elucidate the grammar of a word.

13. Does the above argument have devastating ramifications for referentialist theories of meaning? I do not think so. It has the consequence that referential theories of meaning are of use only insofar as they shed light on bodies of linguistic rules governing the uses of words (i.e., grammars). In turn, this may have the unsavory conclusion for referentialists that their semantic program merely serves to elucidate something which they have regarded as properly belonging to pragmatics. (This may even suggest that the distinction between pragmatics and semantics is not as fine-grained as Millian externalists have assumed.)¹⁹ Whatever the case, I believe the argument in the previous section acquires teeth against referential theories of meaning when we couple it with a relatively innocent, uncontentious principle concerning what it is for a description of x to be an explanation of x . The principle is this: a description is only an explanation of a natural phenomenon insofar as it allows for the direct prediction and/or control of that phenomenon, at least in theory. If a description does not allow for the direct prediction and/or control of the phenomenon which it purports to be about, at least in theory, then it is not an explanatory description with respect to that phenomenon. Moreover, though it is not explicitly a part of the principle which I am articulating, it would follow from this principle that we cannot regard a theory as explanatory if it consists merely in non-explanatory statements.

What immediately follows from the adoption of this principle? It would suggest that even if a theory appears to coherently describe a phenomenon

¹⁹ It will perhaps come as no surprise to the reader that for Wittgenstein, there was no discernable distinction between pragmatics and semantics.

through a variety of novel concepts, perhaps with liberal appeal to all sorts of invented postulates, it nonetheless fails to explain the phenomenon if the capability of an individual to directly predict and/or control said phenomenon is not improved by way of that theory. At best, such a theory – if it can be called that at all – merely broadens the set of possible descriptions which an individual can use to characterize the phenomenon. (But as we shall see, I do not even think this is the case.) With respect to language, a description which purports to be about an expression’s meaning cannot be called an explanation of that expression’s meaning if it does not allow for at least theoretical, direct prediction and/or control of that expression’s meaning – *where prediction and/or control of meaning would be detectable as semantic competence with the relevant expression*. After all, it is only in virtue of being competent with a word that a speaker can wield it with desired effect, making herself understood as she pleases, and effecting the sort of changes she wishes to produce in a listener’s actions or beliefs. She is certainly controlling meaning if she is reliably capable of all these things and more via the use of her expressions. So, with respect to language, our principle tells us this: a description which purports to be about an expression’s meaning is only an explanation of that expression’s meaning if, upon providing that description to a listener, the description is directly responsible for causing semantic competence with the relevant expression. It will be essential to remember this hereafter.

Now at a glance, it might seem that the principle is one of pragmatism; it seems to be about what we *ought* to call an explanation of a natural phenomenon, and not what really is an explanation. Proceeding from this sentiment, one might

object to the principle on the grounds that we ought to be concerned with truth for truth's sake. Thus, if a theory describes things as they really are, then it is a theory worth having – regardless of whether or not the descriptions the theory consists in allow for the theoretical, direct prediction and/or control of the thing they purport to describe.

But such an objection is misguided. The principle is not one of pragmatism, for it genuinely is a means of distinguishing explanations from non-explanations, relative to particular phenomena.²⁰ If a statement purports to describe a natural phenomenon, but does not allow for the direct prediction and/or control of that phenomenon even in theory, then there are no grounds upon which one can even sensibly say that the statement actually describes the phenomenon it claims to describe. It no longer becomes sensible to even regard that statement as being *about* the phenomenon in question. This is the case since it is impossible for an explanation (or a description) of a phenomenon to not allow for direct prediction and/or control, in theory, so long as it genuinely is about that particular phenomenon.

The key qualifier here is 'in theory.' Given this qualifier, the principle only demands that an explanation improve one's ability to predict and/or control the relevant phenomenon *were it the case* the individual using the explanation was in circumstances where prediction and/or control was possible. For example, while

²⁰ As will become evident, the principle determines whether or not a statement is an explanation of *x*. It does not, however, determine whether or not a statement is an explanation generally (since it may be the case that the statement does explain something, albeit not what it purports to explain). The same is true of the principle's application to descriptions of natural phenomena. Hence my emphasis here on 'relative to particular phenomena.'

technological limitations may prevent us from exerting any sort of control over black holes, this does not preclude the possibility of control in theory. We would be capable of exerting control over black holes were it the case we were in circumstances where we could exploit what it is we know about black holes (i.e., were we in circumstances where we had the means of exerting control).

Descriptions of black holes provided by scientific theories are therefore genuinely explanatory. They contrast sharply with what I would consider to be statements which merely purport to explain or describe a phenomenon, but actually fail to do so. Such statements, I think, can be found in metaphysical theories such as trope theories – that is, theories which postulate the reality of abstract particulars (i.e., tropes) to ‘explain’ the properties of objects. To put the matter bluntly, no individual in the history of anything has ever improved her ability to manipulate the structural integrity of an object by having studied trope theory. That is because, as I have suggested, statements which issue from trope theories are not genuinely about objects and their properties, and this is evident since they do not allow for the direct prediction and/or control of the phenomena they purport to be about – not even in theory.²¹ In summary then, our principle is not one of pragmatism. It is not vulnerable to such objections. Moreover, the ‘in theory’ clause guards it against counterexamples which might appeal to situations in which control over a phenomenon was beyond the limits of nomological

²¹ If I were to hazard an explanation as to what trope theories are really about, if anything at all, I would say that they are about ways in which human agents may conceptualize objects and properties. This is perhaps the one thing for which they provide explanations. Thus, while I am happy to leave it open to debate as to whether or not trope theories are worthwhile theories in their own right, I nonetheless maintain my conviction that they are not theories where explanations of objects and properties are concerned.

possibility. Additionally, the scope of the phenomena which can be appealed to in such counterexamples is narrowed by the fact that the principle does not purport to be true of all things which are called ‘explanations’ in ordinary usage. Instead, it is a principle intended for naturalistic inquiry. It is true only of explanations of natural phenomena.

Before we move on, there is one other thing to be noted. I have already said, more or less, that the principle dictates that a theory is not to be deemed explanatory (relative to a phenomenon) if it consists in statements which do not explain the phenomenon they purport to describe. But this claim doesn’t go far enough. Given the sort of considerations we explored above, I would argue that there are no sensible grounds upon which to say a theory is even about x , if it consists merely in statements which are not themselves about x (since they neither explain nor describe x). It will be important to remember this in later sections. Additionally, the reader should not forget my earlier claim that prediction and/or control over the meaning of an expression occurs in the form of semantic competence with that expression.

14. So how does this principle have any bearing on theories of meaning? If one is inclined to adopt the principle, as I am, then one must accept Michael Dummett’s claim that a theory of meaning must be a *theory of understanding* – that is, what a theory of meaning has to provide is an account of what it is someone *knows* when she understands a language (Dummett, 1996, p. 3). It is not sufficient to merely provide an account of which concepts are expressed by which words (p. 5), nor is it permissible to construct a theory which simply

consists in issuing direct ascriptions of meaning (i.e., propositions of the form ‘the sentence/word *x* means...’ or ‘the meaning of the word/sentence *x* is...’) (p. 3). This is because such theories – what we will eventually call ‘modest theories’ of meaning – encounter situations in which the descriptions they provide do not effect an improvement in one’s ability to predict and/or control the meaning of an expression, which is the thing those descriptions purport to describe. That is, they do not produce semantic competence with the relevant expression. And moreover, even when they do provide descriptions which allow for prediction and/or control over an expression’s meaning, the prediction and/or control they produce is nevertheless not of a direct sort. Therefore, while modest theories do consist in descriptions (we shall say of what in later sections), the descriptions they provide are not explanations of what they purport to be about – namely, the meanings of expressions. And as such, it follows that such theories are not really theories of meaning at all.

There are a few steps to be taken in articulating what it is I am suggesting here. First, for the sake of simplicity and clearer prose, I will assume hereafter that the reader knows that to be able to control and/or predict the meaning of an expression, at least in theory, is to be semantically competent with the relevant expression. Second, let us borrow Dummett’s distinction between *full-blooded* theories of meaning and *modest* theories of meaning (Dummett, 1996, p. 5). A full-blooded theory is a theory of understanding: it explains not only which concepts are expressed by which words, but also the concepts themselves which are expressible within the language. A modest theory of meaning, at most, just explains the former. It is *not* a theory of understanding. In a manner of

speaking, modest theories of meaning merely construct taxonomies in which meanings are paired up with sentences.

Now if modest theories of meaning are to be of any use whatsoever to a competent speaker, then that speaker must be in possession of a kind of antecedent knowledge or understanding. Dummett's point is that modest theories of meaning function much like translation manuals; a translation manual "leads to an understanding of the translated language only via an understanding of the language into which the translation is made, an understanding which [the manual] does not itself supply" (p. 6). This is the case since a translation manual obviously presupposes that the speaker has a mastery of at least one of the languages which it uses. Something similar applies to modest theories of meaning: they allow a speaker to comprehend an expression's meaning (and they therefore effect semantic competence with the relevant expression) only if the speaker has a prior understanding of the concepts expressed therein – concepts which the theory does not itself explain (p. 6).

This is the reason for my earlier claim that theories of meaning which do not fall in with Dummett's conception of a theory of understanding (i.e., modest theories of meaning) are guilty of violating the principle I presented in the previous section. If readers will recall, that principle is the following: a description is only an explanation of a natural phenomenon insofar as it allows for the direct prediction and/or control of that phenomenon, at least in theory. Modest theories of meaning violate this principle whenever there come about situations in which the descriptions they afford – ones which purport to be about

particular expressions' meanings – do not result in a positive change to an individual's semantic competence with the relevant expression.

Such situations arise whenever it is the case that a speaker encounters an alleged description of meaning from a modest theory (e.g., an ascription of meaning), but is not in possession of the antecedent understanding required for the description to be of any use. Or, in other words: the speaker is not in possession of the antecedent understanding required for the description to produce semantic competence with the relevant expression. Given this, our principle dictates that such descriptions – as provided by modest theories – are not genuinely about meanings of expressions as they claim to be, since they are not explanations of expressions' meanings. And furthermore, it would follow from this that modest theories are not theories of meaning at all, since they consist in such descriptions.

Why should it matter to us that modest theories of meaning are not genuine theories of meaning, since the descriptions they consist in violate our principle? For one simple reason: all referential theories of meaning, by their very nature, are modest theories. It therefore follows that we ought to reject them as being insufficient or inadequate for the purpose of explaining linguistic meaning, assuming the correctness of our principle. Of course, the onus is upon us to show that they are in fact modest theories of meaning. The first step to be taken towards this goal is to revisit the conclusion which we arrived at in the previous section – namely, that the function of articulating relations of reference is analogous to the function of ostensive definition, insofar as it elucidates grammars.

Articulating relations of reference may have a function within language which is comparable to that of ostensive definition, but they are nonetheless very different. After all, it was no mistake that I said the function of articulating a relation of reference consists in elucidating a grammar *only if the listener possesses sapient awareness of the object of reference* – the obvious implication being that it is more than possible for such awareness to be absent. Therein lies the difference between articulating relations of reference and providing ostensive definitions. In the case of ostensive definition, it can never be the case that the recipient of the definition lacks sapient awareness of the object for which the definition is given, since it is always the case in ostensive definitions that the recipient is in perceptual contact with the object – and to be in perceptual contact with an object is to be conceptually aware of that object, and therefore to be aware of that object as being what it is (i.e., to have sapient awareness).²² It is only in virtue of the recipient's sapient awareness of the object that the ostensive definition of a word is successful, insofar as it produces semantic competence by elucidating the word's grammar. Why? Because a speaker's comprehension of the broad network of rules and norms which dictate how the word may be used meaningfully depends on the speaker being able to think about the object in determinate ways, and these modes of thought are only afforded via conceptual awareness / sapience. In short, to comprehend the meaningful usage of the word

²² Such a view obviously presupposes the sort of account that John McDowell advances in *Mind and World* (1994), in which conceptual capacities are drawn on in experience, thereby permeating experience with conceptuality. Unfortunately, it is beyond the scope of this paper to present a full-fledged presentation and defence of McDowell's view. Nonetheless, the reader should be aware that I presuppose it, and hence why I use the terms 'sapience' and 'conceptual awareness' interchangeably.

is to be able to entertain conceptual thoughts about the word's corresponding object. Thus, if I were to provide an ostensive definition of the word x to an individual who was perceptually cut off from the relevant object (let us pretend that such a situation can even be called an ostensive definition, for it is not), and she had no prior understanding of the word, then no elucidation of the word's grammar would be achieved, and therefore no semantic competence with that word would result. The recipient would not be able to reliably designate objects as being instances of x , nor would she be able to grasp the sort of inferential relations that the sentence 'this is an x ' could stand in, relative to other sentences. She couldn't do anything with the word x . In short, she would be completely epistemically barred from the patterns of usage surrounding the word x .

Now compare ostensive definition to articulating relations of reference. While it is never the case that an individual lacks sapient awareness of the object for which an ostensive definition is given, it is the case that an individual may lack sapient awareness of the object of reference, whenever a relation of reference is articulated to her. In such cases, having a relation of reference articulated to a speaker does nothing to elucidate the relevant word's grammar (and therefore does nothing to produce semantic competence with that word). Here, we might show what we mean by appeal to analogy. Telling a listener that 'the word "Recnik" refers to Recnik' is comparable to telling a monolingual English-speaker 'that (Norwegian word 'skole') means 'Schule' in German' (assuming the listener lacks sapient awareness of Recnik). Such statements explain nothing, absent the antecedent understanding and knowledge required to make use of them. And here, I should note, the analogy is perfectly appropriate since *i*) to elucidate a

word's grammar (that is, the network of rules and norms which determine how the word is to be used meaningfully) just is to classify that word's function, and *ii*) the purpose of translation is also to classify the function of words. Moreover, I should also add that it is no accident that the case of translation lends itself so readily to our analogy, since this was Dummett's earlier point that a modest theory of meaning operates not unlike the way a translation manual does. And as I have been suggesting, a referential theory of meaning is undoubtedly a modest theory of meaning. Modest theories of meaning encounter situations in which the descriptions they provide do not effect an improvement in one's ability to predict and/or control the meaning of an expression, which is the very thing those descriptions purport to describe. This is precisely the sort of situation which we have been describing above. It occurs whenever a referential theory of meaning provides a description which articulates a relation of reference, but the recipient of the description lacks the sapient awareness of the relevant object – awareness which is needed for the description to produce semantic competence with the referring expression.

Of course, even if the reader endorses my principle, and agrees with me insofar as there are cases in which articulating a relation of reference does not produce semantic competence with the relevant referring expression, they might have concerns. Such concerns, I imagine, would take the following shape: while it is certainly true that articulating relations of reference might fail to produce semantic competence, such occasions are the exception, and on the whole we find that articulating relations of reference *usually* produces semantic competence, and so a referential theory of meaning is generally explanatory, even if not

infallibly so. To fend off such an objection, let us reconsider our principle: a description is only an explanation of a natural phenomenon insofar as it allows for the direct prediction and/or control of that phenomenon, at least in theory. Here, it is essential to note that the prediction and/or control afforded must be *direct*. If a statement does not allow for direct prediction and/or control of a particular phenomenon, at least in theory, then it simply is not about that phenomenon – and therefore especially not an explanation of that phenomenon. The fact that articulating relations of reference *can* fail to improve one's ability to predict and/or control the meaning of an expression (insofar as it fails to produce semantic competence with that expression), even if it only happens on the rarest of occasions, is proof enough that even when it does succeed, it does so only indirectly. But this should not come as any surprise. We have already discussed, at length, that the function of articulating relations of reference is to elucidate grammars – and that it is *only* by way of elucidating these grammars that the descriptions provided by referential theories can indirectly produce semantic competence with referring expressions (i.e., can produce the ability to predict and/or control the meanings of referring expressions). In sum, descriptions of reference are not about expressions' meanings. They are about grammars. However we go about it, referential theories of meaning are only modest theories of meaning, and modest theories of meaning must be rejected if we take seriously the sort of principle I have been pushing. Because they consist of descriptions which neither describe nor explain the meanings of expressions, there is no ground upon which to say referential theories are even theories of meaning at all.

15. Let us now cobble together the insights of the present chapter with an aim towards a summary. We began with a pair of Chomskian arguments, as presented by Stainton, directed against theories of meaning which merely consist in articulating word-world relations. These arguments left us with one of two unpleasant conclusions for referential theories of meaning: either public linguistic items aren't real, or they aren't real in the way required for rigorous theoretical analysis. For several reasons, it was the latter conclusion which we chose to endorse. But even so, both made short work of the idea that an empirical study of meaning can afford to ignore individual speakers and their properties, since such an idea was dependent on the assumption that public linguistic items are legitimate objects of scientific study.

As such, the Chomskian critique opened the door for a Wittgenstein-inspired line of argument against referential theories of meaning. This argument urged that what is understood by a speaker when she comprehends a relation of reference is not any sort of empirical description, but rather a set of rules concerning the meaningful use of the referring expression. What she comprehends is the expression's grammar. The point was further illustrated in a discussion of ostensive definition, in which I suggested the function of articulating relations of reference is analogous to that of ostensive definition: both show – or better yet, elucidate – the grammars of words.

By itself, this conclusion was unthreatening. However, when paired with a principle which decreed that a description is only an explanation of a phenomenon insofar as it allows for direct prediction and/or control of that phenomenon (at least in theory), the conclusion would eventually have damaging

implications for referential theories of meaning. Because it followed from this principle that a theory of meaning is not *really* a theory of meaning if it consists in descriptions which fail to explain the meanings of expressions, it followed that what Michael Dummett called ‘modest theories of meaning’ had to be rejected. This was the case since modest theories provide descriptions which purport to be about the meanings of expressions, but these descriptions either fail to produce semantic competence with the relevant expressions, or produce semantic competence only indirectly. This posed a problem for referential theories of meaning since they are undoubtedly modest theories. To illustrate as much, I pointed out that articulating a relation of reference between word x and object y only produces semantic competence with the word x by elucidating the word’s grammar, which in turn depends on the listener being able to think in determinate ways about the object y . These determinate ways of thinking are object-directed modes of thought which are only acquired through sapient awareness (and are therefore conceptual). As such, because articulating a relation of reference only ‘explains’ meaning insofar as a listener has a prior possession of these object-directed modes of thought afforded through sapient awareness, referential theories of meaning are modest by definition, since they explain nothing without tacitly taking into account listeners’ antecedent understanding of objects as being what they are (i.e., speakers’ sapient awareness of said objects). This conclusion was further strengthened by our discussion of the ways in which referential theories of meaning can either fail to allow for prediction and/or control of meaning, or merely produce prediction and/or control of an indirect sort – which, as we noted earlier, are the faults of modest

theories. As such, all of this is to say that we must side with Dummett: a theory of meaning must be a theory of understanding.

16. A few things deserve mentioning before we move on. First, since I endorsed the Chomskian thesis that public linguistic items are not real in the way which would render them intelligible to the scientific perspective, it might seem as if I ought to have been barred from invoking public linguistic items in the Wittgenstein-inspired line of argument which followed the Chomskian critique. This is not the case. As I said earlier, the Wittgensteinian critique belongs to the framework of ordinary discourse and common sense; it does not purport to issue from a scientific perspective or theory. As such, it is not vulnerable to the conclusion of the Chomskian critique in the way Millian externalism and other referential theories find themselves vulnerable. Millian externalism finds itself susceptible to criticism from both directions because it simultaneously *i)* purports to be an empirical enterprise which is continuous with science, and *ii)* invokes the categories and concepts of ordinary discourse and common sense. The Chomskian and Wittgensteinian critiques are therefore not mutually-exclusive lines of argument, but rather two pieces of a general attack.

Secondly, it may strike the reader as peculiar that I have focused not on arguments which target the fine-grained features of Millian externalists' approach to semantic content, but have instead employed 'shotgun arguments' that hit much more than just Millian externalism. As I have said, I have advanced arguments which tend to target referential theories of meaning generally. But even so, Millian externalism features as the principal opponent of

the conception of meaning which I will argue for in the following chapter. This is the case since, if we are attentive to the implications of the arguments given in the present chapter, it must be evident to us that a theory of meaning must take an approach to semantic content which rejects the Millian externalist approach all the way down: it cannot be extensional, anti-individualist, or externalist. Again, I will have more to say about this in the following chapter, and how our present considerations ought to inform our selection of a semantic theory. I raise the issue here only to let the reader know that Millian externalism is still the enemy – even if it is only one among many.

3

The Quietist Picture

17. What follows in this chapter is a largely an attempt to work through the consequences of the Chomskian and Wittgensteinian critiques. We will therefore begin with a re-evaluation of the particular features of Millian externalism's approach to semantic content, and arrive at the conclusion that a satisfactory theory of meaning must be intensional, internalist, and individualist. Upon determining as much, we will then look towards conceptual role semantics (i.e., a family of theories which explicate the meanings of expressions in terms of expressions' use in thought) for a semantic theory which can satisfy the various theoretical requirements we have assembled over the course of the entire paper.

But of particular note in this chapter are the ramifications of the Chomskian critique. In short, while it proved satisfyingly useful in our attack against Millian externalism, it has the unfortunate effect that whatever theory of meaning we select, it must be completely devoid of normative concepts. In turn, it compels us to wonder why we do not simply embrace the methodology of linguistics, if it is the case we are intent on producing a scientific theory of meaning.

Unsurprisingly then, the latter sections of the chapter are a response to this felt need to make linguistics and philosophy of language indistinguishable. Here, we note that a philosophical theory of meaning can retain its traditional subject matter so long as it gives up any aspiration to providing a scientific theory of meaning. And yet, this alternative to merging philosophy of language with linguistics comes with a heavy price: it demands that philosophy of language be reconciled to Wittgensteinian quietism. In the end, it's the Wittgensteinian picture which we endorse, and the semantic theory we sketch in place of Millian externalism is one which takes seriously Wittgenstein's claim that the role of philosophy is *only* to make contributions to our understanding.

18. If the reader will recall, the Millian externalist claims that the contents of our expressions and thoughts are real-world extensions. If one were to say, "The Reichstag was shelled in the battle for Berlin," then the Reichstag itself is a constituent of the proposition expressed by the sentence – not some abstract intension which the speaker uses to single out the Reichstag as the proposition's subject. On the extensional approach to semantic content, the meaning of a singular term like 'Reichstag' just is the object for which the term stands.

Now of the three defining characteristics of Millian externalism which we explored in the very beginning of this paper, it is the extensional approach to semantic content which we find to be the least credible, given the considerations raised in the previous chapter. When we discussed the comparable functions of providing ostensive definitions and articulating relations of reference, we noted several ways in which knowing the extension of an expression is not directly

responsible for producing semantic competence with that expression. To reiterate our findings: to be competent with a term which has a real-world extension requires antecedent possession of object-directed modes of thought, relative to the object which is the term's extension. These modes of thought – ways of thinking about objects through concept-use – are only afforded through conceptual awareness of that object. (I have on occasion referred to this awareness as 'sapient awareness' or 'awareness of an object as being what it is.')

Now, since these modes of thought are invariably required for semantic competence with an expression which has a real-world extension, and because we have sided with Dummett in his belief that a theory of meaning must be a theory of understanding, it follows that an extensional approach to semantic content cannot be endorsed. Why? Because a theory which identifies expressions' meanings with their extensions does not explain the modes of thought required for competency with said expressions, nor does it explain the concepts which these modes of thought consist in. Given this, such a theory cannot be a theory of understanding – that is, a theory of meaning which explains what it is speakers must know or understand when they are competent with a language. Such theories are therefore modest theories, and must be rejected for all the reasons discussed in the previous chapter.

Moving on, it is interesting to note that without extensionalism, Millian externalists lose much of the motivation they had for embracing semantic externalism – the claim that meaning is not determined by anything within a speaker's mental state, but is instead determined by some manner of extrinsic relation. After all, certain arguments which are the mainstay of semantic

externalism are explicitly dependent on an extensional approach to content. Consider Putnam's Twin Earth, for example. Putnam's rejection of semantic internalism was supposed to follow from the fact that an Earthian would be unable to identify the extension of the word 'water' on Twin Earth, were it the case she was transported to Twin Earth without her knowledge. But nothing follows from this observation now. If an extensional approach to semantic content is untenable, then it is of no consequence to a theory of meaning whether or not individuals are capable of finding themselves in situations where their words have unanticipated extensions. It is of interest to a theory of reference, certainly, but Putnam's observation does not illuminate anything worthwhile about the phenomenon of meaning.

Of course, Putnam's brand of semantic externalism is not the only kind out there; troubles were taken to present Wittgenstein as the father of semantic externalism earlier in this paper. Now while Millian externalists would never adopt a Wittgensteinian conception of meaning (for that would more or less amount to unconditional surrender on their part), even so, there is the open question as to whether or not *we* should opt for externalism. And here it seems only natural that we should follow Wittgenstein, given that much of our argument up to the present has drawn upon his insights. Yet, things have turned out otherwise. Between semantic externalism and semantic internalism, I am inclined to think the latter is the better motivated option.

Here then, I suppose, is where we discuss how Wittgenstein went wrong. No one can accuse Wittgenstein of having been ignorant of normativity, or having failed to account for it in his view of language. But in his vehement hostility and

contempt for mentalist theorizing, Wittgenstein seems to have turned a blind eye towards questions of what must be true of a language-user's mental state for her to be responsive to the normativity which allows for successful acts of meaning and understanding. Consequently, while he was able to recognize that ostensive definitions are successful only because there are pre-existing norms which govern their use, he did not entertain the sort of idea we invoked earlier – that is, an ostensive definition produces semantic competence with a word by elucidating the word's grammar *only* because the recipient of the definition must (by ostensive definition's very nature) be in possession of the object-directed modes of thought which are afforded through sapient awareness of the relevant object. Wittgenstein recognized that the function of ostensive definitions was to elucidate grammars, and that this was the means by which they effected semantic competence with the words for which they were given. However, that seems to have been as far as he went. Here, I like to think, it is self-evident that to be competent with a referring expression, one must be in possession of a set of concepts which enable the speaker to think about the corresponding object in all the ways which are reflected in her various usages of that expression. And given that we are committed to avoiding modest theories of meaning, I am inclined to think that a theory of meaning ought to assign a primary importance to these concepts, if not outright construing the acts of meaning and understanding as consisting in a specialized form of their use. Naturally, I have just such a theory in mind. But for now, it is enough to say that while there is no knockdown argument immediately available to us against semantic externalism, we should keep two things in mind: first, there are overwhelming reasons for us to endorse

semantic internalism; and secondly, there are significantly fewer reasons for Millian externalists to endorse semantic externalism, given the untenability of an extensional approach to semantic content.

That leaves us with only Millian externalism's anti-individualism to consider. What are we to make of its claim that speakers do not have first-person authority over the meanings of their sentences? It is worth remembering that Millian externalists were committed to anti-individualism on account of their extensionalism. This was the case because, like the rest of the tradition of Millian semantics, Millian externalism endorses referential realism – the idea that referring expressions can, and do, designate the objects which populate our day-to-day lives. Anti-individualism was therefore a practical necessity for Millian externalists, since it is never the case that an individual has an infallible knowledge of the things to which she refers.

Additionally, anti-individualism within Millian externalism was motivated by externalist considerations. If meaning is determined by some extrinsic relation (e.g., a causal-historical chain), then there is no reason to suppose that speakers have first-person authority over their meanings, since it is implausible to believe that speakers would have a comprehensive knowledge of the extrinsic relations whereby their expressions mean what they do. However, it is now questionable how much motivation such considerations can even afford anti-individualism. After all, semantic externalism itself has lost much of its motivation (for Millian externalism, at least) if I am correct in my claim that an extensional approach to semantic content is not viable. And since the adoption of anti-individualism is no longer a practical necessity for Millian externalists, given

that it followed from their extensionalism and referential realism, one cannot help but question why Millian externalists would continue retain it. But then again, at this point it is not so much a question of whether or not anti-individualism is still tenable for Millian externalism, but rather a much larger question of whether or not Millian externalism is still salvageable as a whole.

Obviously, it is a separate matter as to whether or not we should adopt an anti-individualist approach to semantic content. And here, in much the same way that externalist considerations motivated Millian externalists to adopt anti-individualism, I would suggest that internalist considerations should incline us to do the reverse. Since we have shunned externalism, there seems to be no ground upon which to say either social or environmental factors are the determinants of meaning, over and above the intentions of individuals. Moreover, if we are set upon giving a primary importance to concept-use in an account of meaning, perhaps going so far as to say that meaning is constituted by a form of concept-use, then we seem to be left with individualism as a matter of default. After all, what better authority could there be than the individual whose concepts *are* the meanings of the expressions she uses? One might object on the grounds that no individual has infallible epistemic access to her inner states, but this should not concern us – and nor should it concern us that on rare occasions, individuals use overt language without even realizing they are doing so (when they are talking in their sleep, for example). If the criteria for individualism are made to be so stringent that one cannot have authority over her meanings unless she has unrelenting, infallible knowledge of her mental states, then the distinction between individualism and anti-individualism is of no sensible use. To warrant

the label 'individualist,' it is enough that a theory of meaning regard *no better* authority on the meaning of a speaker's expression than the speaker herself.

All of this seems to point towards one conclusion: we require a theory of meaning which is intensional, internalist, and individualist. And yet, even this deserves some qualification. I would point out that the rejection of an extensional approach to semantic content does not necessitate positing intensions, where intensions are understood as Frege described them. That is, we do not need to posit entities whose function is to act as mediating representations between speakers and objects, entities which belong to a 'third realm' which is independent of human minds. That would be the case only if we were committed to a relational theory of meaning, one which faithfully preserved philosophy's perennial fascination with language's ability to represent, and if we were unsettled by metaphysical questions of how speakers 'connect up' with their external realities. But none of that is the case here. It is true that by positing these entities, Frege dodged the sort of epistemic problems generated by referential theories of meaning, but any intensional theory of meaning achieves the same results without needing to reify intensions. In the sort of theory we have only just begun to sketch (where meaning is a form of concept-use), Frege's insight that two statements deserve different semantic treatments if they differ in cognitive significance is easily accommodated, since a difference in cognitive significance would just be a difference in concept-use and therefore a difference in meaning. So, to say that we must take an intensional approach to semantic content is only to say that our theory will concern itself with the contents of concepts.

Critiquing the specific features of Millian externalism's approach to semantic content has allowed us to glean a minimal set of expectations for a semantic theory. Of course, having acquired these expectations, it now remains to be said what kind of theory will satisfy these demands. To this end, I am inclined to point to *conceptual role semantics*.

19. Not all versions of conceptual role semantics (hereafter CRS) would satisfy the demands we've arrived at, but a few would. CRS is much like Millian semantics insofar as it is a heterogeneous body of theories, wherein one can find radically different positions with only a handful of commitments in common among them. Given that, it is an opportune time to note that it is not within the scope of our present discussion to provide a complete survey of CRS, nor is it even within our means to present an extensive description of a few versions. (I cannot help but mention here that Robert Brandom [1994] uses no less than 700 pages to detail and defend his version of CRS.) Having said that, the present section shall be interested only in providing a general account of what CRS consists of, and what it has to offer. Furthermore, the reader should be warned that we shall withhold our endorsement of any one version of CRS until such time as all the criteria for an adequate semantic theory are made known to us.

Having placed much emphasis on the connection between patterns of usage and semantic competence in the previous chapter, it may come as no surprise that I have elected to endorse a semantic theory which was originally developed from use theories of meaning. In general, CRS construes the meanings of linguistic expressions in terms of the roles those expressions play in thought

(Greenberg & Harman, 2006, p. 295) – or rather, in terms of their *uses* in thought.²³ This means different things to different theorists, but all advocates of CRS agree that an expression has no intrinsic semantic content. There is only content with use.

Typically, the specific kind of use CRS theorists are concerned with is an expression's use in inference. The tradition of CRS, therefore, begins with Wilfrid Sellars (1953; 1954). Following his disenchantment with relational theories of meaning, and the shortcomings he perceived in Wittgenstein's anti-mentalist approach, he was the first to identify linguistic meaning with inferential use.²⁴ To put it *very* simply, he argued that meaning statements (e.g., 'the German word "*Schule*" means "school" in English') merely serve to classify the functions of expressions of object languages, and they do this by revealing their uses to speakers through speakers' antecedent comprehension of a metalanguage. Given that this is a means of producing semantic competence with a translated expression, realizing as much should incline us to think of the meaning of an expression as a sort of functional role it has within a wider system.²⁵ And to investigate this functional role of an expression, we need only to observe three

²³ Theories belonging to CRS are equally theories of the contents of mental states, though here our discussion will be generally limited to CRS as a set of semantic theories.

²⁴ Not all advocates of CRS agree that the meaning of a symbol should be *identified* with its conceptual role. Greenberg and Harman (2006) are one such example.

²⁵ If this line of argument seems reminiscent of my earlier claim that articulating a relation of reference only produces semantic competence insofar as it elucidates an expression's grammar (which, as I said, just is to classify that expression's use), there is good reason as to why this is the case: there are strong parallels between the Sellarsian perspective and the Wittgensteinian perspective.

types of pattern-governed linguistic behavior surrounding the expression's use.

These patterns are as follows:

1. Language Entry Transitions: The speaker responds with appropriate linguistic activity to objects in perceptual situations and in certain states of herself.
2. Intra-Linguistic Moves: The speaker's linguistic conceptual episodes reliably occur in patterns of valid inference, both theoretical and practical, and do not tend to occur in patterns which violate logical principles.
3. Language Departure Transitions: The speaker responds with appropriate behavior to such linguistic conceptual episodes as 'I will now raise my hand' (Sellars, 1974, pp. 87-88).²⁶

Of course, not all contemporary proponents follow Sellars on this matter. Since his original formulation of the inferential approach to meaning, advocates of CRS have preserved some of his insights, but have otherwise developed the idea in different directions. For example, all agree that a speaker understands the sentence 'my sister has the flu' if she can make transitions from it to other sentences such as 'my sister has an ailment.' And yet, *how* CRS ought to understand transitions of this sort is another matter altogether, and one which divides its supporters. I will speak to this issue in a moment. Before moving on, it is worth noting that although CRS has progressed beyond the original

²⁶ That is, her statement 'I will now raise my hand' is accompanied with an upward motion of the hand, and so on.

Sellarsian picture, Sellars's idea that meaning is a functional role *within a wider system* has carried on and generally achieved widespread adoption. CRS is self-evidently a holistic theory of linguistic meaning. On virtually any account of CRS, an expression cannot have meaning in isolation from other expressions, since it only has a meaning in virtue of standing in inferential relations to other expressions. This wider system Sellars spoke of is the total body of language available to an individual – or, more accurately, it is the entire conceptual apparatus in which one employs language.

Now this question of how CRS ought to understand speakers' inferences from one expression to another is typically approached in one of two ways. The first approach is more in keeping with the Sellarsian picture (and, more distantly, the Wittgensteinian picture), since it contends that the meaning of an expression is constituted by the transitions that one *ought* to make. Call this the normative approach to CRS.²⁷ On this approach, a speaker understands an expression only if she grasps the proprieties of inferential moves that connect it to other expressions. Consequently, while most versions of CRS account for normativity in explaining an expression's meaning, this approach *outright explicates* meaning in terms of normativity. To illustrate, it has been said of the approach's most recognizable proponent, Robert Brandom (1994), that he constructs meanings out of rules. But naturally, Brandom is not alone. Eike von Savigny (1988), John Skorupski (1997), and Charles Travis (2000) all take normative approaches to inferential transitions.

²⁷ It is also, on occasion, called the 'conventionalist approach.'

The second approach is admittedly less homogenous than the first. Some, like Ned Block (1987), are motivated by issues belonging to cognitive science, and therefore find themselves concerned not merely with the use of linguistic expressions, but also with that of mental representations. In Block's case, the kind of uses which are of interest are not merely actual, but also counter-factual; it is of interest to his brand of CRS what effect a thought *would have* were it the case stimuli or other mental states had been different. Given all this, it might not surprise the reader to be told that there is a smattering of computational models of CRS in which it is not so much the inferential role of an expression's use (or a thought's use) which is investigated, but rather the causal role. In general, we might characterize this second approach merely as the 'non-normative approach,' since nothing seems to be held in common among its proponents other than the refusal to identify an expression's meaning with the inferential transitions a speaker ought to make from that expression. Unlike normative approaches, this second approach usually employs an explicitly scientific and naturalist methodology.

I have said little about this second approach and its proponents for a reason which I will address towards the end of this chapter. For now, it would seem that we ought to point the way towards a particular version of CRS which will satisfy the expectations we derived from our critique of Millian externalism. (Namely, we need a version of CRS which is intensional, internalist, and individualist.) Unfortunately, we cannot yet do that. There is another matter which requires our attention, since this hitherto overlooked issue has significant implications for our search for an appropriate version of CRS. Specifically, this

issue gives reason to think that the restrictions on what kind of theory we may endorse are much, much more stringent than previously expected. It can be brought into view by considering why Millian externalism, in the broadest possible sense, failed in light of the critiques presented in the previous chapter.

20. A convenient way towards this issue is to consider the manner in which Millian externalism was vulnerable to the Chomskian critique. The critique came to the conclusion that while the relata of referential theories (i.e., public linguistic items, normative categories, not-explicitly normative yet mind-dependent objects, etc.) are indeed real, they aren't real in the way required to render them intelligible to the scientific perspective – and thus, it is impossible to construct a scientific theory of meaning from word-world relations. Evidence produced in support of this conclusion included the following: the ways we individuate public linguistic items in ordinary discourse do not correspond to objective differences in kind; common sense and science deploy concepts which are categorically different from one another, and thereby reveal the world to us in different ways; and lastly, the relata of referential theories are intrinsically normative, and scientific concepts fail to capture normativity.

Millian externalism took the brunt of the Chomskian critique because it simultaneously *i*) purports to be an empirical project which is continuous with natural science, and therefore attempts to approximate to a scientific theory, and *ii*) it aims to explicate (and therefore invokes) the categories and concepts which belong to ordinary discourse and common sense. It is essential to note here that the Chomskian critique would be deprived of its edge against Millian externalism

if Millian externalism gave up either of these two aspirations. Millian externalism was vulnerable to both the Chomskian critique *and* to the Wittgensteinian critique (two very different lines of attack) because it occupies a place which is awkwardly situated halfway between science and folk theorizing. Or to borrow an image from Sellars, Millian externalism failed because it finds itself in a kind of ‘purgatory’ between a scientific image of man in the world, and a manifest image. (I leave it to the reader to decide which one is heaven, and the other hell.) It is extremely telling, with regard to what we can and cannot endorse among semantic theories, that Millian externalism is vulnerable to the Chomskian critique only because it refuses to give up the concepts of common sense while still aspiring to be more than just a common sense theory.

Where are we going with this observation? The obvious point here is that the Chomskian critique cuts both ways: inasmuch as it strikes Millian externalism down, it can also harm us if we are not careful to avoid the mistakes of Millian externalists. Thus, we cannot choose a theory which – like Millian externalism – purports to be a scientific theory while issuing from the framework of ordinary discourse and common sense. If we are interested in the phenomenon of meaning as it really is (and are therefore engaged in a scientific enterprise), then we cannot make public linguistic items our objects of study, nor can we appeal to normative concepts in our explications. Even the notion of intentionality will have no place within our theory. As Chomsky says, “Naturalistic inquiry will always fall short of intentionality” (1992, p. 229). To regard intentions as being ‘real’ is about as meaningful as declaring oneself to be a realist with respect to cats or desks (Chomsky, 1992, p. 207).

All of this is to say that there are massive restrictions on what sort of version of CRS we may select – restrictions which go far beyond our need to produce a theory which is intensional, internalist, and individualist. If we are aiming for a scientific theory, one which describes language as it really is – independent of notions which are bound up in human interests and concerns – then that theory cannot resemble anything like a typical philosophical theory of language. It cannot explicate common sense notions, nor explicate by way of them. In short, while the Chomskian critique was immensely useful for rejecting Millian externalism, it has now put us in an immensely awkward place by demanding a philosophical theory of meaning which is wholly unlike anything one has come to expect of philosophy.

In fact, the Chomskian critique causes problems for us which go beyond merely imposing stringent criteria for an adequate scientific theory of meaning. Specifically, it brings into question why we have even bothered to begin a search for an adequate version of CRS. To get a sense of this, consider the following: if an adequate semantic theory must give up all of philosophy of language's mainstay concepts (e.g., intentionality, normativity, etc.), and if our semantic theory must be intensional, internalist, and individualist, then what possible reason could we have for not endorsing Chomsky's approach to meaning? After all, by Chomsky's own admission, his approach is intensional, internalist, and individualist (1992, p. 211), and so he has already satisfied the minimal set of requirements which we arrived at via our criticism of Millian externalism. Additionally, his approach is undoubtedly free (for obvious reasons) of all the troublesome elements which the Chomskian critique identified in Millian

externalism. In fact, when we consider these things, it seems we have been totally unjustified in our endeavour to find a suitable version of CRS, since an adequate version (namely, Chomsky's) has been staring us in the face all along. On the surface of it, this may not sound like a colossal problem: the pursuit of CRS was just a clumsy oversight, and we can now find satisfaction in following Chomsky. However, there is a massive implication here. To elaborate, it follows from all this that if we are committed to a scientific theory of meaning, then there is no reason whatsoever to not opt for the methodology of linguistics.

This revelation may come across as deeply unsatisfactory to the philosopher of language, as well it should. In a word, it would mean the dissolution of philosophy of language's traditional subject matter, and would effectively dissolve the distinction between linguistics and philosophy of language. And if that were the case – if philosophy of language could no longer have any pretense to having a specialized subject matter – then there would seem to be no grounds upon which to say that there even is a philosophy of language at all. There would only be philosophers doing linguistics.

So then, is this the end of the matter? Did our pursuit of a suitable version of CRS amount to a needless labour? I do not think so. I believe there is a way out of our present difficulties, though conservative philosophers of language may like it even less than the idea of subordinating philosophy of language to linguistics. We shall see. Whatever the case, readers may already have a sense of what this alternative is, since I have dropped hints in past pages.

21. In the concluding section of the previous chapter, I noted that the Wittgensteinian critique is not susceptible to the conclusion of the Chomskian critique; it was of no concern to us that our Wittgenstein-inspired argument made explicit appeal to public linguistic items (and other intrinsically normative objects/concepts) in its criticism of Millian externalism. This was the case since the Wittgensteinian critique did not purport to issue from a scientific perspective, but rather belonged to the framework of common sense and ordinary discourse. Here, one finds a hint as to how philosophy of language can remain distinctive from linguistics. The solution is simply this: philosophy of language must renounce the goal of describing language as it really is, and henceforth may not attempt to approximate to a scientific theory of meaning.

Naturally, it is one thing to state this solution, and quite another to articulate what it entails. As one might expect, such a solution enables us to preserve philosophy of language's traditional subject matter by allowing us to evade the Chomskian critique. But less expectedly, this solution has the effect that one must endorse a form of quietism.

Now as the reader might already suspect, I have been setting the stage for a version of CRS because I believe there is a way in which CRS can accommodate all the insights we generated from our criticisms of Millian externalism, while at the same time being consistent with the notion that philosophy of language ought to be, in a way, quietist. Of course, now is not the time to discuss how CRS might achieve these goals. Much is needed in the way of explaining what quietism consists in, and why we must be committed to it if we are intent on maintaining the distinctiveness of philosophy of language from linguistics. As such, let us

shelve the issue of which version of CRS we ought to endorse, and let us turn our attention to the Wittgensteinian picture of philosophy – since no picture better illustrates what quietism is all about.

Half a century ago, when philosophers perceived themselves as engaging with the ‘essential problems’ of reality (as opposed to the ‘inessential’ ones dealt with by science), many were horrified to hear of Wittgenstein’s anti-theoretical, therapeutic approach to philosophy (Baker & Hacker, 2009d, p. 274). He was not the recipient of much love when he wrote, for example, that “it can never be [philosophy’s] job to reduce anything to anything, or to explain anything,” since “philosophy really is ‘purely descriptive’” (1958, p. 18). And similarly, he was not applauded when he decreed that “all the facts that concern us [already] lie open before us” (1958, p. 6).²⁸ Rather, these seemingly-dogmatic proclamations earned him a curious mix of awe and contempt.

Yet, there was in fact extensive reasoning behind his high-handed decrees. Wittgenstein was the first to popularize the idea that traditional philosophical problems are not like empirical problems; they cannot be investigated and subsequently solved. Rather, he argued that a philosophical problem is created when we are aware (albeit, not usually self-consciously aware) of a kind of disorder in our concepts. The story goes something like this: since our concepts are revealed through our language-use, we come to be tacitly aware of these

²⁸ These sentiments are similarly expressed in the *Philosophical Investigations* (1953) with respect to philosophy of language: “Philosophy must not interfere in any way with the actual use of language, so it can in the end only describe it. For it cannot justify it either. It leaves everything as it is” (§124). Then, with respect to philosophy generally, he says: “Philosophy just puts everything before us, and neither explains nor deduces anything. – Since everything lies open to view, there is nothing to explain. For whatever may be hidden is of no interest to us” (§128).

conceptual disorders whenever a form of expression compels us to think within the terms of an antinomy that we cannot recognize as being an antinomy. This occurs because the trouble-causing form of expression has different uses which correspond to different, mutually-contradicting modes of thought (which, taken individually, are each plausible since they correspond to correct uses), and it is these modes of thought which generate the antinomy. Moreover, we find ourselves unaware that we are in the grip of an antinomy, because we are only self-consciously aware of being in the grip of a philosophical question. When this happens, we are compelled to think we are encountering a real problem, one which may be solved, when in fact we are wrestling with our inability to think about the question in a way which transcends the particular modes of thought that generated the antinomy in the first place. When we do this – that is, try to transcend the particular modes of thought that led to the antinomy – what we are in effect doing is attempting to transcend the language-games which contain the criteria for the correct usage of the form of expression which is the source of all our trouble. The hope here is that by finding the pre-eminently ‘correct’ or ‘true’ usage for the form of expression which led to the conceptual discomfort, we might thereby acquire a means of arbitrating among the modes of thought which led to the antinomy – which would presumably solve the perceived problem. But this actually just makes the problem all the more deeply felt, since there are no pre-eminently ‘correct’ or ‘true’ usages for *any* form of expression. There are just correct and incorrect usages within *particular* language-games. Consequently, the attempt to overcome the conceptual discomfort by transcending language-games is one which fails by necessity, since there are no criteria for usage to be

found outside of language-games. The question's hold over us therefore strengthens. Beyond language-games, we find ourselves in a frictionless, criterial void wherein nothing constrains the use of our language and concepts.

Considerations such as these led Wittgenstein to argue that philosophy should not consist in theories (where a theory is the sort of thing one sees in the natural sciences),²⁹ since philosophers rarely find themselves investigating actual phenomena in the course of addressing philosophy's traditional problems. Rather, the means by which philosophers should address philosophical problems should consist of investigations into the grammars of words – which, by extension, are investigations into modes of thought. The basic idea here is that a philosophical problem is not a *real* problem, but rather a pseudo-problem which is comparable to a psychological neurosis. So, in much the same way that a therapist does not 'solve' a patient's neurosis, but rather helps the patient to identify the underlying causes of the ailment in order to free the patient from its grip, so too does the philosopher aim to diagnose and cure. When one encounters a philosophical problem, the job of the philosopher is to show how the problem was generated by the modes of thought which correspond to the uses of the troublesome form of expression which is the problem's base cause. By showing the afflicted the underlying causes of the problem, the philosopher helps her to shed the illusion of its being a genuine problem. And as I said above, the

²⁹ It is important to remember this when we consider Wittgenstein's quietism. What Wittgenstein calls a 'theory' is actually a scientific theory or one that approximates to a scientific theory. Consequently, there is no tension between Wittgenstein's quietism and my proposal for a non-scientific theory of meaning.

philosopher goes about this task by investigating the grammar of the form of expression which led the afflicted into difficulty.

This therapeutic conception of the role of philosophers obviously brings with it an equally unique conception of philosophy as a whole. For Wittgenstein, it is not the purpose of philosophy to make tangible contributions to our knowledge through the accumulation of new facts. That is, philosophy should not imitate the natural sciences. Rather, the purpose of philosophy is to make contributions to our *understanding* through the dissolution of pseudo-problems.³⁰

Now having outlined Wittgenstein's picture of philosophy, the reader might be in a position to see how something similar to this picture can apply not necessarily to philosophy as a whole, but to philosophy of language specifically. Oddly enough, it is not any sort of Wittgensteinian reasoning that allows us to see why this picture might be true of philosophy of language, but rather the Chomskian reasoning. To begin with, let us remember that if we don't merely want to do linguistics, then the Chomskian critique blocks us from advancing a scientific theory of meaning. That is to say, we must stay within the framework of common sense and ordinary discourse. But if this is the case, then we cannot posit underlying realities which stand in causal relations to our intrinsically normative objects of study. The reason for this has already been stated several

³⁰ Sellars articulated a similar purpose of philosophy. In 'Philosophy and the Scientific Image of Man' (1962), he says, "The aim of philosophy, abstractly formulated, is to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term" (p. 369). He then went on to characterize success in philosophy as consisting of "know[ing] one's way around" things, where 'knowing' in this context refers to a kind of 'knowing *how*,' as opposed to 'knowing that' (p. 369).

times: when one moves to a finer-grained level of observation, one realizes that philosophy of language's objects of study do not correspond to objective differences in kind. In a manner of speaking, moving beyond a manifest image of language results in the total disappearance of anything that might be called a 'sentence' or 'natural language.' It would further follow from this that we could not attempt to reduce our objects of study to sets of more fundamental entities, which is the hallmark of scientific analysis. So already, Wittgenstein's claim that 'philosophers can't reduce anything to anything' would seem to be necessarily true of philosophy of language, given the pressures of the Chomskian critique.³¹ And furthermore, at this point we could not deny his claim (with respect to philosophy of language) that philosophical investigations do not parallel scientific investigations, given the normative dimension of philosophy of language's objects of study. Whereas science is exclusively concerned with what *is* the case, philosophy must take equal interest in what *ought* to be the case. And in the course of investigating what ought to be case, one cannot say they are making tangible contributions to knowledge (i.e., they are producing new facts). After all, it genuinely is the case that all the relevant facts are already known; a competent language user already knows the appropriate usages of the expressions which lead her into conceptual difficulty. She merely needs to be attentive to them for the purposes of evading philosophical puzzles.

³¹ These observations also speak to his point that philosophy does not 'explain' anything. Here, it is important to note, that Wittgenstein had in mind explanations of the sort one encounters in science – that is, reductive analyses. (There will therefore be no tension between quietism and my future claim that a common sense philosophy does – given my principle in the previous chapter – explain language.)

To be clear then, I am *not* proposing that philosophers study language simply on a different level of abstraction from that of linguistics. It is not the case that philosophy of language can occupy a place relative to linguistics comparable to that of biology to chemistry – assuming philosophy of language is, in any way, concerned with the framework of common sense and ordinary discourse, along with its intrinsically normative concepts and objects. Such an interpretation of what I have proposed would miss the point entirely. It is the job of scientific theories, like those provided by linguistics, to pursue integration with the core natural sciences. But here, no such thing is possible for a philosophical theory which refuses to relinquish such concepts as normativity and intentionality. Additionally, any such attempt to salvage philosophy of language along these lines would invariably repeat the mistakes of Millian externalism, and would therefore find itself standing on the sharp end of the Chomskian critique.

All of this might be a lot to swallow. Nonetheless, I would argue that this conception of philosophy of language is infinitely preferable to simply opting to make philosophy of language indistinguishable from linguistics. Concepts such as intentionality and normativity are integral to our ability to make sense of innumerable events in our day-to-day lives, and so philosophy is entirely justified in its felt need to accommodate them. The quietist alternative allows for this. Moreover, as I have already said, this alternative allows for the preservation of a distinct subject matter for philosophy of language. (In fact, if I am correct that philosophers must choose between linguistics and a non-scientific philosophy of language, the fact is there is no philosophy of language without the quietist alternative.) So, the main selling point here seems to be the following: the non-

scientific approach, despite entailing quietism, allows philosophy of language to retain something of a resemblance to its former self. However, I might also add that the benefits of the quietist perspective are not limited to this. Contrary to what might be said by young philosophy students who half-heartedly embrace Wittgenstein as a means of one-upmanship over their peers, Wittgenstein *did* have a profound respect for the importance of philosophy. He did not regard it as an idle amusement (Baker & Hacker, 2009d, p. 300). In keeping with that respect, I would urge that it is not necessary that philosophy should make discoveries on par with those of the natural sciences in order to be useful; it is enough that it liberates us from the confusion and discomfort wrought upon us by those innately specified concepts which belong to our biological endowment.³²

22. We are finally in a position to give a general sketch of a semantic theory which might accommodate all the insights we gathered over the course of this paper. Naturally, it will help us to take inventory of those insights, and list the various requirements that an adequate semantic theory must fulfill.

First, if we are striving for a semantic theory which genuinely explains the phenomenon of meaning, we must opt for a theory of understanding. That is, our semantic theory must provide an account of what it is someone *knows* when she understands a language. Second, our approach to semantic content must be intensional, internalist, and individualist. This means the following: our theory

³² This comment is in obvious contradiction to Wittgenstein's belief that philosophical problems seem to be the bane of philosophers alone. Frankly, he was self-evidently wrong in this matter. He would have thought otherwise if ever he had the pleasure of frequenting either an internet forum or a Youtube comments section.

must concern itself with the contents of concepts; it must determine meaning by way of identifying a speaker's relevant mental states; and lastly, it must regard speakers as having first-person authority over the meanings of their utterances. Third, if we refuse to relinquish the categories and concepts which belong to ordinary discourse and common sense, then we cannot endorse a theory of meaning which approximates to a scientific theory. And fourth, if we do not choose to advance a scientific theory of meaning (i.e., if we do not adopt the methodology of linguistics), then our theory of meaning must confine itself to the framework of ordinary discourse and common sense, as well as accept a mild form of philosophical quietism.

Now I have already proposed that there might be a version of CRS which can rise to the above expectations. In anticipation of sketching this version, I provided a general outline of what CRS consists in, and provided a few remarks on the some issues which divide its proponents. Here, it should be obvious why I had little to say about non-normative approaches which employ methodologies which are explicitly scientific and naturalist. These approaches simply don't enter into consideration here, since we have long conceded that a philosophical theory of language should not aspire to be a scientific theory. And yet, our desired version of CRS is not technically a normative approach either, since the identification of an expression's meaning with the inferential transitions one *ought* to make from that expression would run contrary to our claim that a semantic theory ought to be internalist and individualist. After all, if the meaning of a word just is the set of rules and norms which govern its usage, then it is difficult to make sense of the idea that meanings are in the head.

Having said all this, our desired version of CRS is something like this: the meaning of an expression is to be identified with the set of concepts which the speaker brings to bear in the expression's use. That's it, that's all.

Of course, this evidently simplistic formulation of linguistic meaning deserves massive amounts of qualification. For one, it is absolutely essential to recognize that by invoking concepts as an explanatory device, we are not reducing meanings down to sets of physical entities, observable at a finer-grained level of observation, any more than Wittgenstein was when he identified meanings with 'functions within a rule-bound context.' Rather, concepts (in our theory) are real in the same sense in which public linguistic items and other normative objects are real. To elaborate, they are real and in the mind, but the way in which concepts are individuated will not correspond to objective differences in kind; our notion of a concept is wholly invisible to the scientific perspective. So, while the use of concepts will obviously always correspond to some neurophysiological state, there are no sharply-demarkated physical entities which might be strictly identified as concepts. We cannot therefore be accused of some manner of physical reductionism which is characteristic of a scientific theory.

Unfortunately, this alone does not put us in the clear. What I have written above is not at all inconsistent with ways in which concepts have been explicated in scientific discourses, and thus we are still vulnerable to the charge that we are positing an underlying reality to explain meaning. In a word, we are still vulnerable to the charge of practicing science.

The way out of this difficulty then is to stress the point that our notion of a concept does not carry with it the technical connotations or theoretical baggage

that it might carry elsewhere. Our use of the term ‘concept’ is exactly as it occurs within the scope of ordinary discourse; as I shall discuss shortly, we have arrived at it via Wittgenstein’s method of examining the grammar of the word ‘concept.’ Moreover, it *must* be this way if our theory is to be reconcilable to quietism. If we fail at any point to explain meaning by way of concepts belonging to the framework of ordinary discourse and common sense, then we have, in a manner of speaking, already given up the game. It therefore follows that we cannot do as the methodological naturalists are wont to do – that is, modify and alter common sense concepts for the purposes of increasing their explanatory power. Here, we find the reason for why we cannot be accused of engaging in a sort of scientific enterprise: the common sense notion of a concept is completely ill-suited to perform any sort of explanatory function within a scientific theory, whether that theory is to be found in cognitive science or elsewhere. As a result, our theory remains within the framework of common sense and ordinary discourse, and excludes similar versions of CRS in which concepts are technical innovations.

Now the fact remains that I promised a theory, and here I may appear to have abstained from being theoretical in every possible way. In fact, to those familiar with Wittgenstein, I seem to have very aggressively adopted his anti-theoretical stance. This is true, and yet I have not been duplicitous in my promises. As was noted earlier, Wittgenstein understood ‘theory’ as denoting something which is properly found in the natural sciences. Consequently, to adopt his brand of quietism is to assume his anti-theoretical stance on philosophy, but only if ‘theory’ is understood in his particular sense of the word.

If then I seem ‘anti-theoretical,’ it is only because the reader is inclined to think of theories in much the same way as Wittgenstein did.

For my part, I have been inclined to treat something as a theory of x so long as it consists in descriptions which explain x . This has been the case ever since I presented a principle concerning what it is for a description of x to be an explanation of x . Given this, what I have presented in this chapter is indeed a theory – albeit, hardly a complete one. While the common sense notion of a concept, for example, is ill-suited to performing an explanatory function within a scientific theory, it works just fine within the framework of common sense and ordinary discourse. Similarly, my account of meaning – simplistic as it may appear – has all the explanatory power it requires for the purpose to which it is suited. But before I can illustrate as much, let us bring some much-needed light to our notion of a concept, since this is the first step towards getting a sense of all that is entailed by the claim that meaning consists in concepts.

Earlier, I mentioned that we arrived at our use of the word ‘concept’ by examining the word’s grammar. Given the lines of discussion explored in the previous chapter, it should already be evident why investigating the patterns of use surrounding an expression is an effective way of getting at the expression’s meaning. When we look at the grammar of ‘concept,’ examining its use within ordinary discourse among ordinary language users, what we find is that the word refers to an indeterminate ‘something’ which is only intelligible insofar as it corresponds to ways in which speakers are inclined to behave. Thus, if one has a concept of ‘moral good’ which differs from my own, she will be inclined to act in ways that I would not when she proclaims she is about to perform a righteous

feat. But not only that, she will identify instances of moral good which I would not, and she will be disposed to make different inferential transitions from sentences which feature the concept of moral good. All of this is to say that in ordinary language, while we openly recognize concepts as being mental, they are only intelligible as sets of dispositions to particular forms of behavior.

It should be obvious that by 'behavior,' I mean both linguistic and non-linguistic acts. It is no accident that I chose my examples as I did; what we recognize as being concepts from the perspective of common sense just are dispositions to make certain language departure transitions, language entry transitions, and intra-linguistic moves. As a result, to say the meaning of an expression is the set of such-and-such concepts a speaker deploys when she uses that expression, is just to say that the meaning of an expression is identifiable via the speaker's readiness to behave in particular ways.

The natural objection here is to point out the obvious flaws of any account of concept-use which reeks as strongly of behaviorism as the one I've presented above. But again, I must urge the reader to remember that we are not looking for deep, theoretical insights. We are not attempting to describe meaning 'as it really is.' After all, we are limited to the concepts afforded to us by the framework of common sense and ordinary discourse, and must explicate accordingly. It is true that our account of concepts will fail to explain certain kinds of phenomena, but here we are only concerned with explanations of *meaning*. And where meaning is concerned in day-to-day life, the common sense behaviorist construal of it

seems to do the trick just fine. Without the behaviorist approach, we cannot reconcile philosophy of language to quietism and thereby salvage it.³³

³³ It is important to note the sense in which my proposal is behaviorist, since my position need not provoke the sort of anxieties which are characteristic of contemporary views of behaviorism. My proposal is behaviorist to the extent that Wittgenstein was behaviorist. In other words, my proposal does not imply a denial of mental phenomena (i.e., psychological behaviorism), nor does it suggest that mental concepts can be exhaustively explained in terms of behavioral concepts (i.e., logical behaviorism). Rather, the idea here is that while mental phenomena exist and are not conceptually reducible to observable behavior, they nonetheless depend upon public, observable criteria to render them intelligible. That is, we are capable of talking about mental phenomena (and therefore capable of identifying them) only because their identities are fixed by corresponding them to observable behaviors. Consequently, the Wittgensteinian position is behaviorist only with respect to the *intelligibility* of mental phenomena (and other private entities, processes, events, etc.). For Wittgenstein's defence against the label 'behaviorist,' see §§304-307 (1953).

Afterword

23. If the general conclusion of this paper could be summarized with an image, it would be this: traditional philosophical theories of meaning have attempted to occupy points along a spectrum where the framework of science is on one end, and the framework of common sense is on the other. Millian externalism is one such traditional theory. But over the course of our investigation, it has come to be made evident that there can be no spectrum; there just are the two frameworks, and nothing resides between. In their attempt to approximate to scientific theories while at the same time invoking the categories of common sense, theories such as Millian externalism have failed to achieve adequacy within either framework.

The Wittgensteinian picture I have presented remains the only option for a philosophy of language which doesn't merely imitate the methods and doctrines of linguistics. Moreover, I have gone to great lengths to show why it can be made palatable: it allows philosophy to retain something of itself. If we are inclined to think the Wittgensteinian picture is as intolerable as the prospect of dissolving philosophy of language in favor of linguistics, this is only because we are uncomfortable with the prospect of philosophy being unable to maintain the pretense that it is in pursuit of real discoveries.

References

- Bach, K. (1987.) *Thought and reference*. Oxford: Clarendon Press, Oxford University Press.
- Baker, G. P. and Hacker, P. M. S. (2009a.) The Augustinian conception of language. In *Wittgenstein: Understanding and meaning 2nd edition, part I* (pp. 1-28). Malden: Wiley-Blackwell.
- Baker, G. P. and Hacker, P. M. S. (2009b.) The language-game method. In *Wittgenstein: Understanding and meaning 2nd edition, part I* (pp. 46-64). Malden: Wiley-Blackwell.
- Baker, G. P. and Hacker, P. M. S. (2009c.) Ostensive definition and its ramifications. In *Wittgenstein: Understanding and meaning 2nd edition, part I* (pp. 81-106). Malden: Wiley-Blackwell.
- Baker, G. P. and Hacker, P. M. S. (2009d.) Philosophy. In *Wittgenstein: Understanding and meaning 2nd edition, part I* (pp. 271-306). Malden: Wiley-Blackwell.
- Block, N. (1987.) Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10(1), 615-678.
- Brandom, R. (1994.) *Making it explicit: Reasoning, representing, and discursive commitment*. Cambridge: Harvard University Press.
- Bianchi, A. (2012.) Two ways of being a (direct) referentialist. In J. Almog & P. Leonardi (Eds.), *Having in mind* (pp. 79-92). Oxford: Oxford University Press.
- Boër, S. E. and Lycan, W. G. (1974.) Knowing who. *Philosophical Studies*, 28(5), 299-344.
- Burge, T. (1979.) Individualism and the mental. *Midwest Studies in Philosophy*, 4(1), 73-121.

- Chomsky, N. (1992.) Explaining language use. *Philosophical Topics*, 20(1), 205-231.
- Chomsky, N. (2000.) Language from an internalist perspective. In *New horizons in the study of language and mind* (pp. 134-163). Cambridge: Cambridge University Press.
- Davidson, D. (1992.) The second person. *Midwest Studies in Philosophy*, 17(1), 255-267.
- Devitt, M. (1981.) *Designation*. New York: Columbia University Press.
- Donnellan, K. (1966.) Reference and definite descriptions. *The Philosophical Review*, 75(3), 281-304.
- Dummett, M. (1996.) What is a theory of meaning? In *The seas of language* (pp. 1-33). Oxford: Oxford University Press.
- Frege, G. (1892.) Sense and reference. *The Philosophical Review*, 57(3), 209-230.
- Greenberg, M. and Harman, G. (2006.) Conceptual role semantics. In E. Lepore & B. C. Smith (Eds.), *The Oxford handbook of philosophy of language* (pp. 295-322). Oxford: Oxford University Press.
- Kaplan, D. (1989.) Afterthoughts. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan*, pp. 565-614.
- Kripke, S. (1972.) *Naming and necessity*. Malden: Blackwell Publishing.
- McDowell, J. (1994.) *Mind and world*. Cambridge: Harvard University Press.
- Putnam, H. (1973.) Meaning and reference. *The Journal of Philosophy*, 70(19), 699-711.
- Putnam, H. (1975.) The meaning of 'meaning'. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science, Vol. 7: Language, Mind, and Knowledge* (pp. 131-193). Minneapolis: University of Minnesota Press.
- Rozemond, M. (1993.) Evans on de re thought. *Philosophia*, 22(3-4), 275-298.
- Salmon, N. (1989.) Illogical belief. *Philosophical Perspective*, 3, 243-285.
- Salmon, N. (2004.) The good, the bad, and the ugly. In *Content, cognition, and communication* (pp. 309-338). Oxford: Clarendon Press.

- Savigny, E. von (1988.) *The social foundations of meaning*. New York: Springer-Verlag.
- Sellars, W. (1953.) Inference and meaning. In K Scharp & R. B. Brandom (Eds.), *In the space of reasons: Selected essays by Wilfrid Sellars*, pp. 1-27.
- Sellars, W. (1954.) Some reflections on language games. In K Scharp & R. B. Brandom (Eds.), *In the space of reasons: Selected essays by Wilfrid Sellars*, pp. 28-56.
- Sellars, W. (1962.) Philosophy and the scientific image of man. In K Scharp & R. B. Brandom (Eds.), *In the space of reasons: Selected essays by Wilfrid Sellars*, pp. 369-408.
- Skorupski, J. (1997.) Meaning, use, verification. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language*, pp. 29-59.
- Soames, S. (1989.) Semantics and semantic competence. *Philosophical Perspectives*, 3, 575-596.
- Soames, S. (2002.) *Beyond rigidity*. Oxford: Oxford University Press.
- Stainton, R. (2006.) Meaning and reference: Some Chomskian themes. In E. Lepore & B. C. Smith (Eds.), *The Oxford handbook of philosophy of language* (pp. 913-940). Oxford: Oxford University Press.
- Sullivan, A. (2010.) Millian externalism. In R. Jesion (Ed.), *New essays on singular thought* (pp. 246-270). New York: Oxford University Press.
- Thau, M. (2002.) *Consciousness and cognition*. Oxford: Oxford University Press.
- Travis, C. (2000.) *Unshadowed thought: Representation in thought and language*. Cambridge: Harvard University Press.
- Wettstein, H. (1986.) Has semantics rested on a mistake? *The Journal of Philosophy*, 83(4), 185-209.
- Wettstein, H. (2004.) *The magic prism: An essay in the philosophy of language*. Oxford: Oxford University Press.
- Wittgenstein, L. (1953.) *Philosophical investigations*. Malden: Wiley-Blackwell.
- Wittgenstein, L. (1958.) *The blue and brown books*. Malden: Blackwell Publishing.