

Lexical Flexibility, Natural Language, and Ontology

CHRISTOPHER A. VOGEL*

Department of Philosophy, University of Maryland, USA

The Realist that investigates questions of ontology by appeal to the quantificational structure of language assumes that the semantics for the privileged language of ontology is externalist. I argue that such a language cannot be (some variant of) a natural language, as some Realists propose. The flexibility exhibited by natural language expressions noted by Chomsky and others cannot obviously be characterized by the rigid models available to the externalist. If natural languages are hostile to externalist treatments, then the meanings of natural language expressions serve as poor guides for ontological investigation, insofar as their meanings will fail to determine the referents of their constituents. This undermines the Realist's use of natural languages to settle disputes in metaphysics.

Keywords: Ontology, realism, polysemy, metaphysical methodology, externalism, internalism.

1. *Introduction*

Metaphysical investigation has, for the better part of the past century, been conducted by way of linguistic meaning. By tracing the meanings of expressions to their worldly extensions, metaphysicians aim to determine a sentence's ontological commitments by examining the purported worldly satisfiers of its truth-conditions. As a consequence, metaphysicians embrace—either explicitly or not—a particular view about linguistic meaning in order to render their investigatory practices coherent. My task here is to undermine this assumed view about

* My thanks to Erin Eaker, Steven Gross, J. Brendan Ritchie, and Alexander Williams for fruitful discussions about this project, and for comments on earlier drafts. I am also deeply indebted to Paul Pietroski for his repeated feedback on multiple drafts of this paper, and for the profound influence of his work on my own.

linguistic meaning, and thereby undermine the metaphysical methodology that assumes this troubled view about natural languages and their semantics.

Sider (2002, 2009, 2011) embraces the metaphysical view I hope to undermine, which holds that there is both an objective structure to reality, and a unique language that mirrors this structure:

...some candidate meanings [for an expression] ‘carve nature at the joints’ more than others, and it is part of the nature of reference and meaning that candidates that carve nature [closer to] its joints are more eligible to be meant. The meaning of a word, then, is the best candidate, where strength of candidacy is based on (1) fit with meaning-determining facts about the speaker or her linguistic community, and (2) intrinsic eligibility on the part of the candidate. (Sider 2002: xxi)

Sider embraces the first of these criteria in defense of using natural languages as metaphysical guides, holding that the use of natural language expressions constrains the space of any given expression’s possible “joint-carving” meanings. Applying such criteria makes sense for ontological investigation only if the relevant language has a semantics that determines the referents for expressions in that language. On this assumption, the quantificational commitments of true sentences in such a language constitute the domain of entities that stand to satisfy the truth-conditions of sentences in that language. And insofar as that language reflects the “jointedness” or structure of reality, this domain is just the domain of things that exist.

Such a language must have a semantics that is *externalist*. Put more precisely, this metaphysical methodology is committed to the following hypothesis pertaining to the (class of) language(s) it deems useful for metaphysical investigation:

(E) For any expression e (in a given language L), the meaning of e determines e ’s truth-conditions.

For the externalist, the meaning of an expression determines its satisfaction-conditions or truth-conditions,¹ either by identifying meanings with truth-conditions, or by identifying the propositional content of a linguistic expression in a way that determines the truth-conditions of the expression. Likewise, the meaning of a word determines the mind-independent objects that constitute the word’s extension.²

¹ I use ‘truth-condition’ from here on, and in (E) as short-hand for the much clumsier ‘truth-condition or satisfaction-condition’. The purpose for marking the distinction between these types of conditions is to mark the difference between sentences and sub-sentential expressions. Since sub-sentential expressions cannot be true nor false, they cannot have truth-conditions. But I trust that adopting this convention here will not lead to much confusion.

² The exact scope of this claim, and related externalist assumptions will be teased out with more precision in what follows, particularly in §2.1 and §4.1. To advertise a bit, the meaning of indexical expressions, even according to externalists, will fail to determine their truth-conditions in the absence of a (index providing) context. This sort of concern is (I think) not terribly vexing, in that externalism can accommodate

Despite the long-standing tradition in linguistics and philosophy of treating meanings as relations words bear to worldly objects, I will argue that natural language meanings are hostile to such externalist treatment. Most prominently, natural language expressions routinely exhibit a kind of *lexical flexibility* suggested by Chomsky and others, and this flexibility is ill-captured by the rigid models available to the externalist. Accommodating this flexibility under an externalist semantics either yields implausible ontological burdens on such theories, or belies good explanations for the relevant data, lending increased credibility to an internalist³ approach to linguistic meaning.⁴

Importantly, these findings bear on the projects of metaphysicians that see natural language meanings as suitable tools for ontological investigation. If natural languages are mind-dependent, and are poorly characterized by externalist semantic machinery, then the *Realist's* metaphysical methodology that reads ontology off of the qualificational commitments of natural language meanings is without foundation. The purpose of this paper is two-fold: first to undermine the basic externalist claim that the meaning of an expression determines its truth-conditions, and second, thereby undermine the Realist use of natural language meanings—and the intuitive judgments that rely on them—as useful tools in metaphysical inquiry. The goal is to undermine (\mathcal{E}), in the following Realist argument:

(\mathcal{E}) For any expression e (in a given language L), the meaning of e determines e 's truth-conditions.

If natural language meanings determine the truth-conditions and referents of their constituent expressions, then natural languages can play an important role in ontological investigation...

...since the meanings of true natural language expressions will pick out their real-world referents to populate the worldly domain.

If the meanings of terms in a natural language fail to determine their truth-conditions (and their constituents' referents), then investigating ontology by analyzing the meanings of true natural language expressions will yield indeterminate answers to ontological questions. The thought is that, with an adequately regimented natural language, one where vague and imprecise terms have been purged from the lexicon,

such cases. So while some considerations for the role of context in determining a sentence's truth-conditions are needed, the core Fregean ideal captured in (\mathcal{E}) is at the heart of externalism.

³ While the purpose of this paper is not to directly argue for such theories, see Pietroski (2008); Pietroski (2010); Pietroski (forthcoming) and Hinzen (2006); Hinzen (2007).

⁴ Some care will be taken regarding the terminology here, since 'meaning' has been used to identify numerous different properties associated with natural language expressions, utterances, interpretations, and the content of a bit of communication. In §4.1 these differences are spelled out, but for the time being, 'linguistic meaning' here is intended to identify the meaning properties of an expression that remain constant across various contexts in which that expression is used.

philosophers can make use of meaningful expressions in that language to do metaphysics, by tracing the meanings of terms in true expressions to their referents. It's this thought, and the externalist assumption (\mathcal{E}) presupposed by it, that I argue is troubled.

2. Setup

2.1. Taxonomy

The bulk of this paper is committed to advancing arguments against externalism. But before we proceed, some distinctions may be helpful. For purposes of taxonomy we should distinguish between internalist and externalist views about both *meanings*, and *languages*. As I understand the externalist's commitments, the *meanings* of natural language expressions are relations (of a particular sort). Meanings on internalist theories are non-relational, at least where one of the *relata* is a publicly available thing.⁵

The internalist/externalist distinction pertaining to *language* regards the ontological nature of languages. For the externalist view of language, languages are mind-independent things, while for the internalist, languages are aspects of the human mind. The logical space of externalist views regarding language and meaning are exhausted by adopting either 1) an internalist or externalist view about language, and 2) an internalist or externalist view about meanings.⁶ Thus one can be an internalist about language (IL) or an externalist about language (EL), holding that language is either in the mind or not.⁷ Similarly, one can be an internalist about meaning (IM) or an externalist about meaning (EM), holding that meanings either are relations between words and objects, or not. As such, the logical space of views is displayed in the following:

⁵ Drawing the distinction along relational lines under-determines the content of internalist theories of meaning, but for our purposes here, the under-determination is immaterial. Insofar as any externalist semantics relates publicly available things to natural language expressions (and internalist theories do not), if such a semantics cannot account for the linguistic data, this undermines (\mathcal{E}).

⁶ For the sake of completeness, one could also deny the existence of either languages or meanings. The motivations for either position aside, the Realist (the target of this work) would not welcome such a deflationary view.

⁷ This is an over-simplification. One can remain agnostic about the ontology of language, while denying an externalist conception (cf. Hinzen 2007: §1.5). One could also *deny* that there are languages at all, a claim many Chomskyans seem to endorse (cf. Chomsky 1986). Such theorists are considered internalists in the literature. I ignore this distinction for the purpose of simplicity, as nothing I say here trades on this distinction.

Table 1: Externalist and Internalist View

	Internalism Meaning (IM)	Externalism Meaning (EM)
Internalism Language (IL)	IL-IM	IL-EM
Externalism Language (EL)	EL-IM	EL-EM

2.2. Assumptions/Motivations

Anti-externalist arguments target both EL and EM theories. Though these two classes of arguments can be treated distinctly, they share a series of assumptions and motivations. Most notably they are inspired by a Chomskyan approach to language generally. In part this embodies a commitment to a naturalistic methodology. On this approach, language is treated as an object of scientific investigation, as a naturally occurring phenomenon, in principle no different than biological reproduction, combustion, planetary motion, or viscosity. This focus is not merely on empirical investigation, but that such investigation can make testable predictions, provide insightful explanations, and can be integrated with other scientific disciplines—most notably psychology and biology.

That humans have the ability to communicate the content of their thoughts *via* vocalization, and that we learn to do this in a short four years, are naturally occurring phenomena that beg for an (naturalistic) explanation. Explaining these phenomena in large part requires characterizing what it is to understand a language—put flat-footedly, one must “know English” in order to “use English.” As such, the study of language should seek to answer three questions:

- (i) What constitutes knowledge of language?
 - (ii) How is knowledge of language acquired?
 - (iii) How is knowledge of language put to use?
- (Chomsky 1986: 3)

As Chomsky (1986) argues (and as we’ll see in §3), this commitment to naturalism and the guiding questions in (i)–(iii) rule out certain conceptions of language as viable candidates of study. They are excluded simply because language on these conceptions cannot be investigated through naturalistic means. Importantly, the veracity of these arguments depends on the success of the research program that insists on investigating language by naturalistic means. If the endeavors of such a research program bear no explanatory fruit, then that failure tells against treating language (and meaning) as a natural phenomenon. There is little doubt however whether the Chomskyan tradition in linguistics has failed in this regard (Baker 2002; Boeckx 2006; Piattelli-Palmarini et al. 2009).

2.3. Outline

The argument presented here is two-fold. First, I bring together various considerations, from different sources, which collectively offer sufficient reasons to doubt the truth of the claims adopted by externalists, in particular (\mathcal{E}). In the course of presenting these arguments I will both draw novel connections between the points offered by others, and consider (to then rebuff) externalist replies to them. The result then, is a series of considerations that collectively count against both the view that languages are mind-independent objects (i.e. EL theories), and the externalist thesis (\mathcal{E}) as it pertains to natural languages (i.e. EM theories).

The arguments offered here do *not* show that externalism is false. The purpose of gathering together this evidence is to show that the externalist *hypothesis* (\mathcal{E}) is contentious and troubled, and not an obvious truth to be taken as the starting point for other domains of inquiry.⁸ The second part of the argument here presses this point, since much metaphysical investigation *assumes* the truth of this externalist thesis. If the considerations offered here against the externalist proposal are indeed as troubling as I claim, then the Realist that conducts ontological inquiry by way of natural language meanings is burdened with those troubles.

I begin by addressing externalist views that treat languages as mind-external objects (i.e., EL theories), by arguing that such views forestall attempts to address central questions about the human capacity to use and acquire natural languages. In §4 I turn to externalist views about meaning (i.e., EM theories). Because of the long standing externalist tradition in linguistics and philosophy of language, I begin that section with a lengthy clarification, indicating the sort of linguistic phenomena that I do *not* think are troubling for the externalist. So clarified, I argue that the *lexical flexibility* exhibited by broad swaths of natural language expressions are poorly explained by the externalist hypothesis (\mathcal{E}). I consider possible externalist responses to the challenge presented by such flexibility before turning to the assumed role of mental content at play in EM theories. There I stress the complex relationship between linguistic expressions and the concepts they purportedly express, noting that the commonly held *labeling* theory of linguistic meaning—whereby words are merely labels for concepts—faces two profound difficulties. After presenting this rather large body of evidence against various externalist commitments, I conclude by arguing that these considerations against externalism present a fundamental challenge for Realist metaphysical methodologies, by showing that they assume the externalist theses this evidence renders contentious.

⁸ It's worth remembering here the oft-quoted claim from Lewis (1970), the progenitor of contemporary metaphysics, whereby he boldly asserts the externalist hunch, with little argument: "Semantics with no treatment of truth-conditions is not semantics."

The consequence for the Realist is that her use of natural language meanings to address metaphysical questions is unjustified.

3. *Languages as Objects*

3.1. *Arguments Against EL Theories*

There are two conceptions of the ontology of languages, an externalist (EL) view of language and an internalist (IL) view of language. The former view construes a language as a mind independent (abstract) object. Languages on this conception are abstract structures relating mind-independent objects to terms, words, or expressions—artifacts in some sense, that we use to denote (other) objects. Accordingly, one understands a language when they can identify and use the abstract structure that most sensibly coheres with the usage in their linguistic community. This conception of language can be found in Lewis (1970, 1975):

What is a language? Something that assigns meanings to certain strings of types of sounds or of marks. It could therefore be a function, a set of ordered pairs of strings and meanings. (Lewis 1975)

A commitment to naturalism speaks against thinking of languages as the abstract objects described here. In treating a language as a mind external object, as Lewis (1975) puts it, “in complete abstraction from human affairs” one wants to know how humans can come to understand or “know” languages so construed (p. 19). That is, on this Lewisian characterization, languages are abstracta: functions that take us from symbols to truth values, combined with a grammar that delineates how these symbols can be combined in acceptable ways to form interpretable expressions (or strings). If a naturalistic approach to language seeks to answer (i) what knowledge of a language amounts to, a Lewisian treatment of language renders this question intractable. As an infinitely-membered set of ordered pairs of expressions and their functional meanings, this conception of language not only gives us little direction as to how to answer (i), but seemingly gives too sparse a collection of resources to answer the question at all. Put more tangibly, all children (placed within a linguistic environment) have adult-like competence with a natural language by the age of four. A child when placed in a community of English speakers will come to “know English” by the age of four. On a Lewisian view, for a child to “know English” they must first decide which set amongst an infinite array of (infinitely-membered) sets of expression-function pairs is the English set, and then second, they must bear the right kind of epistemic relation to that set. As such, to explain what knowledge of a language (so construed) amounts to requires an account of the sort of relation that a human can bear to an abstract entity such that this is the sort of relation a four-year-old child can enter.

The worry here is not that, given certain metaphysical commitments to nominalism, any theory committing us to *abstracta* is off base. The worry Chomsky presses pertains to the conditions for explanation, and particularly whether certain conceptions of language (whatever their metaphysical commitments) forestall viable strategies to answering fundamental questions. Treating languages as abstract entities is problematic not because they are abstracta, but because abstracta *qua* objects of knowledge bear mysterious epistemic relations to human minds. As such, to explain what knowledge of a language (so construed) amounts to requires an account of the sort of relation that a human can bear to an abstract entity, and in particular a set of ordered pairs, such that this is the sort of relation a four-year-old child can enter. How one proceeds to answer these questions seems hopelessly unclear.⁹

Chomsky puts a related point about language acquisition in terms of “legibility conditions” on a natural language. If a child is to come to have “knowledge of a language,” they must come to represent that language in their mind/brain. For a given child to have “knowledge of English,” they must have come to represent both the grammar of English—the algorithms by which one can combine lexical items to form larger expressions (or sentences)—and at the very least, some internalized list of those lexical items that combine syntactically. In whatever way this information is encoded in the mind/brain, it must fit into the architecture of the mind/brain. If our four-year-old has figured out, or “knows”, which grammar (as an abstract object on an EL theory) is the English grammar, she has represented such a grammar in her mind/brain. When she wants to utter an English sentence this grammatical knowledge must be applied (to the lexical items she also “knows”) in such a way so that her articulatory systems can make the right sort of audible noise. That is, the representations she builds using her grammatical knowledge must encode information in a way that her articulatory system can make sense of—those instructions must be *legible*. Thus the structure of the abstract object that is English must abide by such legibility conditions, coming not only from the articulatory system, but from any other aspect of the mind/brain that the child’s linguistic knowledge must interact with. But once we recognize that the structure of the abstract object *English* is beholden to the legibility

⁹ A Lewisian might argue that such an explanation need not be directly forthcoming for the project of building an externalist semantics for externally construed languages to proceed. After all, humans somehow manage to learn mathematics, and the best account of the ontological nature of mathematical language is decidedly abstract, and set-theoretical. So clearly (the thought goes) humans can stand in the relevant epistemic relation to abstract objects, and in particular abstractly construed languages. Notice however, that four-year-old children do not exhibit mastery and competence with the language of mathematics, even if they do have innate mental structures that aid them in acquiring proficiency with numerosity (Carey 2009: Ch. 4). This contrast between children’s felicity with mathematical and natural languages requires an explanation, and that explanation is precisely what the Chomskyan challenge to EL views demands.

conditions imposed by the human mind/brain, this defeases the motivations for thinking of a language in this abstract way Chomsky (2000).¹⁰ Taken together, considerations of acquisition and legibility pose serious, and as far as I can see, unanswered challenges for EL theories.

None of this shows that these *abstracta* do not exist. There could well be, in addition to the mental structures hypothesized (and studied) by the Chomskyan tradition, functions filling an infinite Fregean hierarchy mapping objects of one kind (say <*e*>) to objects of another kind (say <*l*>). But these functions must earn their keep. We are told that we should believe in these *abstracta* because they are essential to successful explanations of linguistic phenomena (Lewis 1986: Ch. 1). The point here about legibility conditions is that these structures serve no explanatory purpose, and indeed present explanatory obstacles, in explaining how a child comes to “know” a language. Thus, the externalist owes us some indication of what these objects are meant to explain.

One possible explanatory virtue of adopting an externalist view about languages (and meanings) is that it yields an intuitively plausible account of successful communication. If languages (or their meanings) are external entities, then successful human communication is explained by the mutual relatedness of individuals to the language(s) they know. That is, two humans can succeed in communicating because they are related to, or come to understand, the very same *thing*—a language. But this explanatory virtue depends on the availability of a cogent specification of the individuation conditions for the external languages that individual speakers come to grasp. The prospects for these conditions, I contend, are grim, at least for the Realist.

3.2. *Grim Prospects for Linguistic Objects*

Commitments to naturalism motivate the ontological arguments against EL theories above. EL theories hold that natural languages are objects-in-the-world, whose existence is independent of minds. As such, expressions of a natural language are likewise objects-in-the-world. As we saw in the previous section, naturalist inquiry into the acquisition of natural language compels the externalist to defend the utility of EL conceptions of language. An externalist might reply, indicating that treating languages as publicly available objects yields a plausible account of linguistic communication. If languages are mind-external

¹⁰ Chomsky also argues that treating languages as abstract structures, akin to the formal languages of mathematics, renders aspects of natural languages inexplicable, in principle. Chomsky uses two examples, ‘imperfections’ in natural languages, to highlight the mismatch between the structure of natural languages and formal languages: 1) that natural languages have uninterpretable features, and 2) the displacement property (Chomsky 2000). Even if treating natural languages like formal ones leaves room for explaining such features, the point here is that there is no good justification for stipulating at the outset of investigation that the object to be investigated must meet the (optimality) conditions of a formal language (especially if even superficial differences speak against such stipulation).

things, communication between speakers can be understood by way of their respective relations to the same publicly available thing—the language they all are said to know. One strain of internalist arguments presses this ontological claim, showing that there is no scientifically respectable notion of ‘natural language expression’. More plainly, these arguments contend that mind-independent words (and thereby, mind-independent languages) do not exist, and thus cannot stand in any (reference/meaning) relation to any object in the world.

One motivation for thinking of languages as external is to account for communicative success. In order for two speakers to understand one another, they must know the same language, and this (so goes the argument) can be explained if there is some single mind-external thing they both know. Knowledge of language on this view is had when a speaker bears the right sort of relation to this external object. Thus two speakers that “know English” can make sense of each other’s speech because they bear the same relation to the same mind-external object. Thus when you use a sentence(-object) of English, I understand this sentence because I am related to that *same* English object in the *same* way you are, *via* knowing. However, this requires that languages can be differentiated ontologically, insofar as you and I must bear the same relation to the *same language*.¹¹

For purposes of differentiation, “common sense” methods will not do. As Chomsky notes, common sense treats Dutch and German as distinct languages, despite the fact that people “who live near the Dutch border can communicate quite well with those living on the German side...” (Chomsky 2000: 48). If treating languages as mind-external objects is meant to explain communicative success, then the fact that speakers of “different languages” can communicate linguistically¹² is unexplained by such theories. The common sense notion of language gets the extension wrong in the other direction as well, insofar as the Mandarin and Cantonese spoken “dialects” of “the Chinese language” are mutually

¹¹ There is a fundamental problem with this conception of language, namely that any particular human will fail to *completely* know any such “language.” Externally construed as an infinite set of expression-meaning pairs, no individual will come to have full knowledge of a language, at least if such knowledge is construed as knowing what the members of this set are. Given the limits on human cognition, no individual could know this infinitely long list of pairs. At best, we must have an incomplete knowledge of such a “language.” But the account of communicative success above assumes that two speakers stand in the same relation to the same object. If a language is an infinitely large set of sentence-meaning pairs, two speakers of that language either: (a) stand in a different (incomplete) knowing relation to that same set, or (b) stand in the knowing relation to different subsets of that set. Thus if appeals to some external entity are meant to aid in explaining how humans manage to successfully communicate, such an explanation will have to succeed despite the fact no human can fully know such a language. See Dummett (1978, 1993) for discussion.

¹² Humans, and other animals, communicate in many non-linguistic ways: body language, facial expression, gesture, etc. The important cases for this point are those instances of communication that are clearly effected with language.

unintelligible. The EL-view here fails to explain why these speakers of “Chinese” cannot communicate *via* spoken language, since (according to common sense) the relevant populations both “know Chinese.”

Since the common sense division of languages will not serve the purpose of explaining communicative success, another means of distinguishing languages (as objects) is in order. One might appeal to the *elements* of languages to distinguish them. On this proposal, two languages, (say) English and French, are distinct because of the differences between the elements that constitute them—one contains words like ‘photographer’, ‘apartment’, and ‘cat’ while the other contains words like ‘photographe’, ‘appartement’, and ‘chat’. This move requires that these elements differ along some important dimension, such that the first three belong to English and the last three belong to French. Notice that appealing to usage will not be helpful. Defining ‘French’ and ‘English’ by indicating that speakers of French use the latter and speakers of English use the former is viciously circular.

An EL theorist must make use of some other property that these words share that marks the boundary between English and French. But to what properties could an EL theorist appeal? The sonic properties of these words seem like bad candidates. ‘Cat’ and ‘Chat’ have more in common in this regard than ‘photographer’ and ‘cat’ do. Further, considering the variation seen in pronunciation across speakers of the “same word”, appeal to such properties will not distinguish words as to cohere with the communicative motivations for EL theories:

To take an example, why are ‘fotoGRAFer’ (said in Bombay) and ‘foTAH-grafer’ (said in Toronto) the same word, yet ‘foTOgrafo’ (said in Buenos Aires) is not the same word as the former two? (Stainton 2006: 918–919)

To explain communicative success and failure, the EL theorist posits the mind-external object *English* that our Bombay speaker and Toronto speaker both “know”, distinguished from other languages based on the elements of that language—words like ‘photographer’. In this example, the EL theorists wants to say that we have two words here (as opposed to one or three), one in English and one in Spanish. Appealing to the sonic properties of (utterances of) words here will clearly not help such a theorist, given variations in pronunciation.¹³ The point is familiar to phonologists, namely that there is no *sui generis* cluster of sonic properties that utterances of a word share in common. What ‘fotoGRAFer’ and ‘foTAHgrafer’ share in common is the manner in which they are represented by humans, which involves features of the system for creating phonological representations from environmental noise.¹⁴ There are then no word-objects that can be differentiated without appeal to structures of the human mind/brain—or no such objects that can be investigated by naturalistic means.

¹³ Appeals to orthography will not be any more helpful here given that illiterate individuals can communicate using spoken language quite well.

¹⁴ See Bromberger & Halle (1995) for discussion.

These problems are particularly trenchant for theorists committed to a Realist position. For the Realist the naturalist perspective is not negotiable. If there is no mind-external, naturalistically respectable notion of ‘natural language expression’, and thus no worldly objects of that sort, then there can be no (semantic) relations between words and objects. Such a Realist cannot fall back on a kind of pluralism or fictionalism about words, because their methodology prohibits such a retreat.

Thus a commitment to naturalism will preclude certain notions of language, specifically those that construe natural languages as mind-independent abstract entities. However, this naturalistic commitment does not preclude a semantics for an I(L)-language from having an externalist character. Even if language is properly construed as an aspect of the human mind, the expressions of that language, seen as mental representations, can still have contents that are cashed out in terms of mind-independent objects. Such an account of meaning is still importantly externalist. This view pushes the externalist to adopt two positions: first, an EM conception of linguistic meaning, and second an externalist account of mental content. In the remaining sections I discuss problems for both of these views. I turn next to the various arguments put forward for why (natural) I(L)-language meanings are hostile to externalist treatment.

4. Compositional Referential Semantics and Natural Language

A theory of meaning for a natural language should provide a means for pairing sentences of a natural language with their meanings. Various semantic theories accomplish this goal in varying ways, but externalist semantic theories insist that such a pairing requires relating sentences to objects in a (worldly) domain. These various EM theories adopt the thesis in (E). But a theory of meaning must also explain the distribution of competent speaker judgments about natural language constructions. This includes (but is not limited to) judgments about inferences speakers draw between expressions, and the meanings they are (not) apt to assign to expressions. A theory of meaning ought to offer a plausible explanation for this data.

The primary strain of anti-externalist arguments I will press illustrate that an externalist semantics cannot obviously accommodate the relevant data, or that in accommodating the relevant linguistic data such a semantics must take on board independently implausible assumptions. As a means of articulating the data points I contend EM theories fail to capture, I’ll begin this section by briefly indicating what these criticisms are *not* arguing.

4.1. *How Not to Argue against Externalism*

There is a line of criticisms attributable to Austin (1962), Strawson (1950), and Wittgenstein (1953, 1972) regarding externalist semantics summarized by indicating that words do not refer/denote, *users* of words do.¹⁵ One way of explaining this point is to highlight the role of indexicals (and demonstratives) like ‘I’, ‘you’, ‘here’, ‘now’, and ‘that’ in determining the meanings of expressions that contain them. The meaning of indexicals are intimately tied to the context of their use. Even if we accept that sentences like (1) have constituents with referential meanings,¹⁶ sentences like (2)–(4) do not, because they are importantly incomplete (nonsensical for Wittgenstein¹⁷) outside of their use:

- (1) Kiruna is foggy.
- (2) I think that should go there.
- (3) I am here.
- (4) I am a philosopher.

Sentences with indexicals, like (2), (3) and (4), have different meanings when uttered by different agents (at different times, different places, and using different gestures). In such cases, context plays the role of determining the referent of the indexical constituent(s) of these expressions. Call this the role played by *semantic context*. The meaning of these expressions determine their truth-conditions, in a context that supplies coordinates for the indexes proposed to be part of their interpretations, and thereby maintains the externalist idea that meanings determine the truth-conditions of expressions (Kaplan, 1977 [1989]). Such cases do *not* undermine EM theories.¹⁸

¹⁵ To quote Strawson (1950): “Referring is not something an expression does; it is something that someone can use an expression to do.”

¹⁶ Though this is far from obvious, despite what traditional semantics textbooks might suggest. Even if the conditions that satisfy the predicate ‘is foggy’ are codified, consider the ontological status of a city that moves from one location to another location two miles to the east (Rolander 2013). And even if those metaphysical difficulties can be met, it is far from clear how such a view can account for sentences like “The tallest mountain in Sweden, Kebnekaise is situated 100 km or 62 miles from Kiruna. ...Nowadays the town [of Kiruna] is not relying solely on the mine” (girontravel.se, 2013). While a city construed as a spatio-temporal object might well be situated some distance from some other object, and might even survive relocation, such an object is not of the sort that relies on anything, much less revenue. See Chomsky (2000) for examples of this sort regarding London.

¹⁷ See Wittgenstein (1972: §10).

¹⁸ Strawson and Wittgenstein would surely disagree. Broadly, the point pressed by many ordinary language philosophers is that the sort of distinction presented here between semantic and cognitive context, and indeed the semantics/pragmatics distinction generally, is misguided. My brief treatment of these concerns here is meant merely to clarify the arguments proffered in the remainder of this section. Even if the externalist can address the problems raised by Strawson and Wittgenstein, possibly in the manner indicated here, many trenchant problems remain. My thanks to an anonymous reviewer for raising this concern.

There are other ways in which the context of an expression's use can impact or "change" the meaning of the expression. Consider the following expression:

(5) Barack Obama is human.

The "literal" meaning of this expression, on an externalist understanding, indicates that some particular individual has a particular property. I can use this expression to convey a thought, the meaning of which is that some individual is a *homo sapiens*, on those occasions where I intend to convey the *linguistic meaning* of the expression. The related Austin-inspired point regarding the (externalist) meanings of expressions contends that because we use language in non-literal ways, and further that such non-literal usage is pervasive, the meanings of expressions are the conditions of their use. Thus what a sentence *means* follows from the contexts in which a speaker can felicitously use the sentence. For example, one could use (5) in various ways:

(5') Barack Obama is human. *Context*: Obama has made some mistake.

In uttering (5) in the context of Obama's mistake in (5'), I do not merely intend to communicate information about Obama's place on the phylogenetic tree. My usage communicates (or intends to communicate) the linguistic meaning acontextually attributable to the expression in (6):

(6) One should refrain from feelings of disapprobation toward Obama in this instance in light of the fallibility of humans.

Call the role context plays in such cases *cognitive context*.¹⁹ This non-literal usage of language is pervasive.²⁰

However, this feature of language does little to undermine the externalist program in semantics. That we can use sentences to convey thoughts that do not match the intuitive literal meanings of expres-

¹⁹ For a discussion of this differing role of context, and the distinction between semantic and cognitive context see Bach (1999, 2004).

²⁰ This feature of natural language seems to be at the heart of Strawson (1950), which is often misunderstood as merely indicating the context sensitivity of indexical (and demonstrative) expressions. While one way of making "different use of the same sentence" is to use an indexical expression in different contexts, this is but an instance of a more general phenomenon (Strawson 1950). The general point for Strawson is that the proposition expressed by a speaker, and thus whether what is spoken is true, depends quite heavily on the context of utterance (and the use of the sentence)—a point Russell seems to miss. Strawson focuses on the context sensitivity of indexicals, largely embedded in definite descriptions, because he argues that Russell's theory of descriptions, which sacrifices the connection between grammatical and logical form for the sake of sentential truth preservation, does so needlessly (Strawson 1950). The theory is not needed, according to Strawson, because *sentences* are not true or false, *uses* of sentences in utterances are. Russell's reply seems to miss this substantive point: "As regards 'the present King of France', [Strawson] fastens upon the egocentric word 'present' and does not seem able to grasp that, if for the word 'present' I had substituted the words 'in 1905', the whole of his argument would have collapsed" (Russell 1957).

sions *requires* that expressions have static linguistic meanings that are context-independent.²¹ And this literal meaning could be externalist, even if the communicative intention of a speaker using such a meaning heavily depends on cognitive context.

The gap between linguistic meaning and the thought inferred by an audience to a speech act (or utterance) does not present a problem for truth-conditional semantics. However, some utterances closer to the semantic-pragmatic boundary seem more troubling. Consider the following examples:

- (7) John is too smart [for this job].
- (8) John finished [writing/playing] the sonata.
- (9) John is ready [for class/to go home/...].

Each of these sentences, though acceptable to competent speakers of English, seem (in some sense) incomplete, as indicated by the bracketed content. For the externalist, the linguistic meaning of expressions of the form in (9) [NP-COP-XP] predicate properties to the individual(s) denoted by the noun phrase. ‘John is awake’, for example, predicates of the individual denoted by ‘John’ that he is not asleep. But the meaning of (9) does not predicate some general property of *readiness* to *John*.²²

Interestingly, these sentences and their completed counterparts stand in a relationship that exhibits features typically attributed to semantic properties, yet seem to be driven by pragmatic inferences. The

²¹ In uttering a sentence, I intend to communicate some thought or other. The task of my audience is to infer this thought from the sentence I used in my act of uttering. Given my overt intention to communicate a thought, my audience must identify the intended thought, in some way or other (Grice, 1957 [1989]). Contextual cues, shared biological sense modalities, a common presupposed set of knowledge, and other aspects of the uttering act all constitute the evidence available to my audience in making the correct inference about my intention. When communication is successful, they ascertain my actual intention. But paramount among the evidence considered in this inferential move is the choice of sentence used in the utterance. That I utter (5) and not (say) its negation in the context indicated in (5’) seems to matter, and it matters precisely because the sentence has a static linguistic meaning. Were the sentence void of any literal meaning, it could play no role in my audience’s inference making. That I cannot communicate the thought in (6) by using *any* expression I choose illustrates this point. There are limits on what a speaker can reasonably expect his audience to infer about his communicative intention on an occasion of utterance, and largely this is because expressions have static linguistic meanings.

²² Though this seems to be the view of Cappelen & Lepore (2005). This view however, fails to explain what needs explanation in these cases. Namely, competent speakers of English treat (9a) and (9b) as having the same meaning. The disquotational account of meaning fails to capture this data in any non-stipulative way. In fact, for Cappelen & Lepore (2005) the proposition expressed by a sentence is the disquoted sentence once we “disambiguate every ambiguous/polysemous expression in [the sentence]” (p. 145). However, the presumption that cases of ambiguity and polysemy can be resolved, *prior to* giving a semantics for the expression either denies that the data need explanation or denies that there are any data there to explain. For (9a) and (9b), that speakers treat these sentences as having the same meaning is ill-captured by a theory that insists they express distinct (disquoted) propositions.

relationship between (9a) and (9b) is different than the relationship between (5) and (6).

- (9) a. John is ready.
- b. John is ready for it.
- (5) Barack Obama is human.
- (6) One should refrain from feelings of disapprobation toward Obama in this instance in light of the fallibility of humans.

Cases like (9a) and (9b) present a problem because the role that context plays seems to be more general (or cognitive), while exhibiting (semantic) entailment patterns that pragmatic cases do not. To complete (9a) and arrive at (9b), my audience has to make use of knowledge not provided by the semantic context, specifically knowledge about *readiness*. Yet (9a) and (9b) exhibit a mutually entailing relationship that pragmatically inferred thoughts rarely have. Regardless of how (semantic) context provides the expression ‘it’ with the relevant event that John is ready for in (9b), this determination will (for the externalist) determine the truth-conditions for the expression. But any context in which the provided event makes (9b) true, also makes (9a) true.

The worry then is this: (9a) and (9b) seem to have the same truth-conditions, yet appear to have different meanings, insofar as (9a) is incomplete and (9b) is not. For the externalist, the truth-conditions of an expression are determined by its meaning. That they come apart in these cases is troubling for the externalist view.²³

The anti-externalist arguments in the remainder of this section are not the (now) traditional worry associated with Strawson, Austin, and Wittgenstein that a single expression can be used in a variety of ways to express a variety of thoughts. Nor is the worry expressed by internalists that expressions with indexical constituents require context to determine their truth-conditions. Such points do not speak against an externalist semantics (though they indicate that (\mathcal{E}) requires some clarification regarding the determination relation). The point that internalist worries stress is that, while the meanings of indexical expressions might be well-captured by appeals to (something like) the content-character distinction, they are a special case of a much more general phenomenon that is misrepresented by such treatment, and exhibited by case like (7)–(9). Natural language expressions exhibit a kind of *lexical flexibility* that is not isolated to a few problem cases to be addressed by intricate logics, but is a ubiquitous feature of natural languages—one that is importantly misrepresented by Kaplanian treatment. It is to this phenomenon that I now turn.

²³ Stanley (2000), for example, notes this trouble and offers a solution to cases of this kind by introducing unarticulated syntactic constituents. My purpose here is not to evaluate the merits of every externalist reply to such cases. I introduce them here merely to suggest the sort of worries I will (not) focus on in the remainder of this section.

4.2. *Lexical Flexibility*

An externalist semantics, or EM theory, must be capable of explaining or accommodating the distribution of meaning assignments competent speakers of a language give to expressions. In the previous section we reviewed proposed counterexamples to (E). In response to these cases (E) could be maintained, by adding characters to the meanings of expressions, as is done for indexicals. In this section, we will consider phenomena that cannot be so easily accommodated, namely those related to *lexical flexibility*.

Chomsky (1977) marks a distinction between various ways in which the meaning of an expression can be multifarious, distinguishing between expressions that exhibit *flexibility*²⁴ from those that exhibit *ambiguity*. The English word ‘trunk’ is ambiguous, with meanings used to denote both luggage, and a part of an elephant (not to mention tree parts, and humans parts). This kind of multifarious meaning is importantly different than the kind exhibited by expressions like ‘book’. Compare the following expressions:

(10) John wrote a book.

(11) This book weighs five pounds.

The use of ‘book’ in (11) is used to denote a particular, concrete (heavy) book. For (11) to be true there must be a contextually relevant physical thing with a particular heft. However, for (10) to be true there is no such requirement. As Chomsky notes, John could have the book composed in his mind, having never deployed pen and paper. This might lead us to conclude that the multifarious meaning of ‘book’ is like the *ambiguity* of ‘trunk’: we have two distinct (homophonous) lexical items, ‘book_a’ corresponding to the *abstract* usage in (10) and ‘book_c’ corresponding to the *concrete* usage in (11).

However, if this treatment of ‘book’ is apt, we should expect uses of ‘book’ to behave like uses of ‘trunk’, since ‘trunk’ *is* ambiguous is the way ‘book’ is purported to be on this explanation. That ‘book’ is less

²⁴ Often this phenomenon is termed ‘polysemy’ in the literature. I avoid the term here for two reasons, one priggish, and one substantive. The priggish reason is that the term ‘polysemy’ indicates (by its roots) that a polysemous word has multiple meanings. This is true for such terms of course, but ambiguous words are also ones with multiple meanings. Thus the contrast in the literature between polysemous terms, which are troublesome for mainstream semanticists, and ambiguous terms, which are not troublesome, is misrepresented by this use of terminology. The substantive reason for my usage here is that lexical flexibility is a property that applies rather broadly, and manifests with different semantic behavior in different contexts, two of which I discuss here. However, the explanation for the various manifestations of polysemy can be unified by internalist proposals (Pietroski 2005: §3.2). See also Pietroski (forthcoming). Lastly, my usage does not ignore the role of pragmatics in understanding the import of the examples presented below, but rather insists that the phenomena to be explained are semantic (as suggested by the discussion in §4.1).

well-behaved suggests that the semantic relationship between uses of ‘book’ is of a different sort:

- (12)a. This book, which John wrote, is five pounds.
 b. John wrote a book, this is it, and this book is five pounds.
 c. John wrote a book_a, this is it, and this book_c weighs five pounds.
- (13)a. This trunk, which Jumbo grew, is full of clothes.
 b. Jumbo grew a trunk, this is it, and this trunk is full of clothes.
 c. Jumbo grew a trunk_a, this is it, and this trunk_i is full of clothes.

The expression in (12a) is (roughly) synonymous with (12b), using ‘book’ in the two distinct ways discussed, as reflected in (12c). But (13a) is not synonymous with (13b), at least not if we interpret (13b) as making use of both lexical expressions of ‘trunk’ as in (13c). They are synonymous if we imagine Jumbo’s nose full of textiles, but this interpretation is not available for (13c). Importantly, this difference is not attributable to the syntax of these various phrases, as (12a) and (13a) appear in the same syntactic frames.

That these two expressions (‘book’ and ‘trunk’) behave in semantically disparate ways in relative clause constructions counts against a semantics that treats them as formally similar. That is, this behavior suggests that we do *not* treat the relationship between uses of ‘book’ as ambiguity (i.e., homophony). As Chomsky concludes:

Thus [in cases like ‘book’] we have a single formal element with a fixed range of meaning, and relativization is possible, despite the shift of sense. But in the case of ...[‘trunk’] (idiosyncratic ambiguity) we have two formal elements ...with the same phonetic form. (Chomsky 1977)

In cases like ‘book’, the lexical entry has a range of interpretations, exhibiting a *flexibility* that permits a kind of mixed use as in (12a). In cases like ‘trunk’ there are two lexical entries that are *homophonous*, each with distinct and unrelated meanings, rendering mixed use interpretations of expressions like (13a) unavailable. The trenchant point is that the flexibility exhibited by ‘book’ is pervasive in natural languages, and poorly captured by theories that treat them as cases of homophony (as seemingly EM theories must).²⁵

There are two kinds of flexibility explored in the remainder of this section, one based on non-linguistic knowledge, and one based on ontological type. The truth-conditions for some natural language expressions are not determined by the referents of their constituents, and the manner of their composition. To determine their truth-conditions, competent language users must deploy non-linguistic knowledge of a general sort—too general to be considered the semantic context of utterance. The second sort of flexibility permits expressions that ap-

²⁵ As a limiting case of this strategy, which treats color terms as massively homophonous at the granularity of *use*, see Rothschild & Segal (2009).

ply multiple predicates of different types to a single noun phrase. The would-be externalist meanings of such expressions require objects of a(n) (impossibly) bizarre sort.

4.2.1. *Non-linguistic Knowledge*

Consider the following expressions:

(14) Football games are played by jerks.

(15) Residential houses are robbed by jerks.

Sentences like (14) and (15) highlight the fact that we bring extra-linguistic knowledge to bear on linguistic expressions in determining the truth-conditions for sentences. The information contained in the *linguistic* properties of expressions like these (even when combined with their semantic contexts) do not determine the truth-conditions of those expressions. As such, meanings cannot be (or cannot determine) truth-conditions, and thereby cannot be externalist.

On any EM theory the difference in the truth-conditions of any two sentences (*modulo* semantic context) must be a consequence of the difference in either their constituents or the manner in which those constituents are syntactically related. So, given that (14) and (15) appear in the same syntactic frames, and that they differ only with regard to the two constituents ‘football-games’/‘playing’ and ‘residential-houses’/‘robbed’ respectively, whatever (14) indicates is true of the relationship between *football-games*, *playing*, and *jerks* should, according to (15), hold true for *residential-houses*, *robbing*, and *jerks*.²⁶

However, the expression in (14) means that *every*²⁷ football game is played by jerks, while the expression in (15) emphatically *does not* mean that *every* house is robbed by jerks. The information essential for deriving the truth-conditions for (14) involves the tight relationship between *games* and *playing*, namely that there can be no unplayed game—a relationship that does not hold between *houses* and *robbing*. But this information is not a linguistic property of the expression, and not part of the linguistic meaning of the expression. Thus in deriving the truth-conditions for (14) a competent speaker of English relies on information not present in the expression (nor even the semantic context). If meanings are (or determine) truth-conditions, then the truth-conditions of any meaningful sentence will be determined by the expression alone. But such a theory will fail to explain the differences in meaning between (14) and (15) not captured by the difference in their constituency.

²⁶ See Chomsky (1975) for similar examples.

²⁷ Admittedly, these expressions are generics, and speaker judgments in this domain are (seemingly) not concordant. However, even if the expression in (14) is not interpreted with a universal quantifier, (14) indicates (at a minimum) that *most* football games are played by unsavory individuals. In contrast, (nearly) no one will interpret (15) as indicating that most residential homes are robbed, even if most of those robberies are perpetrated by jerks. Since this difference is not syntactic, the externalist is burdened to explain why the relationship between the VPs and the NPs in these two expressions is different.

This general phenomenon is not limited to generic expressions. We bring non-linguistic knowledge to bear on linguistic expressions in other ways that do not seem to rise to the level of pragmatic inferences, but are also not a function of semantic context. The differing contribution of expressions like ‘coffee’ as a predicate in ‘coffee drink’ and ‘coffee grinder’ present a *prima facie* problem for the externalist. If the semantic contribution of an expression can be recovered by the semantic contribution of its constituents (plus the means of their composition) the meaning of ‘coffee’ in these expressions should make the same contribution across uses. But a coffee drink is one that is *composed of* coffee, while a coffee grinder is not made up of coffee at all. While these are facts that *any* theory of meaning needs to explain, the externalist will have particular difficulty dealing with this problem, insofar as the contribution of ‘coffee’ on such a theory is exhausted by its reference to a property—and in particular, a property instantiated by the indicated objects in the various expressions in which it functions as a predicate. But there seems to be no obvious single candidate for the needed property in this coffee-case. The best candidate properties for the would-be denotations of ‘coffee’ in these two expressions seem profoundly different: *being-composed-of-coffee* and *used-in-the-production-of-coffee*. The difficult task for the externalist is not only in pinpointing the relevant *single* property in such cases, but also in constructing an account of how speakers come to triangulate on such (non-obvious) properties in the many cases that exhibit these features (e.g. ‘metal shears’, ‘home loan[/inspection]’, ‘rain delay[/coat]’, ‘blue marker’, etc.). This phenomenon is pervasive in natural language, and not easily explained by the externalist.

4.2.2. *Ontology and Satisfaction*

The lexical flexibility exhibited by natural language expressions cannot be accounted for by an EM theory that demands that the meanings of expressions determine a unique referent, as a mind-independent object. Consider the following two English expressions

(16) The Hirshhorn-Museum is bankrupt.

(17) The Hirshhorn-Museum is a cylinder.

A competent speaker of English could think that these expressions are true. An EM theory accounts for this fact by indicating that such a speaker takes the following conditions to hold in the world: 1) there is an object, the one ‘Hirshhorn-Museum’ denotes, which satisfies the conditions for ‘is bankrupt’; and 2) there is an object, the one ‘Hirshhorn-Museum’ denotes, which satisfies the conditions for ‘is a cylinder’. This would require that there is some single object, denoted by ‘Hirshhorn-Museum’, that can be both bankrupt and a cylinder.

The predicate in (16) requires that this object be a financial institution, understood through an array of socio-economic notions. Whatever

these notions demand about the nature of financial institutions, be they collections of individuals or something more abstract, such objects do not seem to have a *shape*. Similarly, the kind of object that would satisfy the predicate in (17), in this case a building, is seemingly not the sort of thing that can have financial troubles. Building projects can have fiscal crises, but buildings seemingly cannot. The EM theorist then either owes us an account of the kind of object that can satisfy the predicates in both (16) and (17), or they must explain why this problem does not generate in the first place.

Taking the latter strategy, an EM theorist might appeal to the presence of ambiguity in natural languages, as in the following:

(18) The geese are by the bank.

The expression in (18) is ambiguous. Much like we saw with ‘trunk’ earlier, we can treat ‘bank’ as homophonous, indicating the distinct lexical items ‘bank_f’ and ‘bank_b’. The ambiguity of (18) is then explained by appeal to the homophony of these two distinct lexical items. The expression in (18) simply fails to determine which ‘bank’ is being used, and as such the expression can have different meanings based on which item is intended.

Likewise, one could argue that we really have two lexical entries for ‘Hirshhorn-Museum’, one that denotes the institution, and another that denotes a building. We can represent this difference between (16) and (17) as

(16’) The Hirshhorn-Museum_f is bankrupt.

(17’) The Hirshhorn-Museum_b is a cylinder.

Since ‘Hirshhorn-Museum’ identifies two distinct (though homophonous) lexical items, contextual information determines which item is used in (16) and (17) respectively, preserving the distinct meanings of the expressions, while assuaging the worry that an EM semantics requires a single (metaphysically suspect) referent for the DPs in the two expressions. So, the externalist might insist, there need not be some single object that satisfies ‘is bankrupt’ and ‘is a cylinder’, since ‘Hirshhorn-Museum_f’ and ‘Hirshhorn-Museum_b’ denote different objects.

But this reply will not do. The same speaker that endorses (16) and (17) would also endorse the following:

(19) The Hirshhorn-Museum is bankrupt and it is a cylinder.

Whatever the technical details are that govern anaphoric meanings, ‘it’ in (19) must derive its meaning and referent from ‘Hirshhorn-Museum’. Whichever lexical entry the context might supply (‘Hirshhorn-Museum_f’ or ‘Hirshhorn-Museum_b’) the meaning of the anaphoric ‘it’ is exhausted by the referent of whichever lexical item is demanded by the context. Thus, for the EM theorist to explain how competent speakers treat (16)–(19), there must be some single worldly object that is both a cylinder and bankrupt. So even if we grant the EM theorist his homophonous response, the ontological concern remains.

Pietroski (2005) also notes that while a natural language speaker could endorse expressions like (16)–(19), the sentence in (20) is strange in a way (19) is not:

(20)# The Hirshhorn-Museum is a bankrupt cylinder.

The oddity of (20) in conjunction with the acceptability of (19) (and the plethora of sentences like them) is unexplained by a semantics that treats meanings as determining truth-conditions. On any EM theory, the way the world would have to be in order for (19) to be true would also make (20) true: there is some object, the referent of ‘Hirshhorn-Museum’ that is both bankrupt and is a cylinder.²⁸ Insofar as satisfiers of mutually binding predicates of different ontological types are absent from the domain of worldly objects, this speaks against EM theories of natural language meanings that require them.

But, the Realist might bite this bullet, as some semanticists do (see Ludlow, 2003, 2011). They might just stipulate that the domain contains objects that are at once both concrete and abstract. However, this bullet biting is both unmotivated by the externalist argument, and tastes far worse than the Realist might suspect. The externalist *hypothesis* contends that, given the pedestrian objects of the world like chairs and rabbits—objects that we have good *antecedent* reasons to posit—a theory of meaning can be developed given only this domain of pre-theoretically plausible things. The externalist proposal is that natural language meanings can be rendered comprehensible without appeal to the mysterious existence of things like *Sinne*. Thus, the theory garners intuitive support because we are not forced to accept into our domain a vast hierarchy of bizarre objects, like *Sinne*.

In this light, the Realist’s bullet biting is quite strange.²⁹ As a means of avoiding the troubling consequences of viewing meanings as myste-

²⁸ While concatenating predicates does not always yield an expression with the same meaning as conjoining them, this does not seem like such a case. Consider:

- (1) This is a fake diamond.
- (2) This is fake.
- (3) This is a diamond.

While (1) implies (2), it does not entail (3)—in fact (3) must be false if (1) is true. This and other examples are problematic for straightforward applications of concatenating predicates, but the difficulty here is not an ontological one. The sort of objects that satisfy ‘is a diamond’ can also satisfy ‘is fake’—namely physical objects.

²⁹ Famously, Lewis (1986) argues that our best semantics for modal expressions posits a vast plenitude of extant possible concrete particulars that stand in spatio-temporal relations to form possible worlds. The argument on offer contends that these proliferate posits are necessary because the cost of failing to accept these concrete *possibilia* are outweighed by the theoretical benefits they confer. But, this strategy is often misconstrued. These *possibilia* are not objects we merely discover in the process of analyzing meanings. Our use of modal expressions in natural language reveals a commitment on behalf of the natural language speaker to the existence of *possibilia* (concrete or otherwise) only if one *assumes* an externalist semantics. As Ludlow rightly states,

...ontology is tied to the demands of our scientific theory of the semantics of natural language, and not the kinds of entities and objects that members of a

rious *Sinne*, the Realist accepts the existence of hybrid abstract-concreta (concrete-abstracta?), trading the mysterious for the bizarre. The further point is that in accepting the existence of ontologically bizarre hybrid objects, the Realist is *not* merely acknowledging the ontological entailments of natural language usage. Rather, she is making a prediction about the kinds of things we should expect to find in the domain, if the externalist hypothesis is correct. In the absence of any pre-theoretical evidence that there are such things, she bears the burden of providing a good reason for believing that such things exist, beyond the fact that their existence supports her hypothesis.

Notice too, that the nature of these bizarre objects is distinct from those typically associated with Chomskyan critiques of externalism. Unlike ‘flaws’ (Chomsky 1981), or ‘the average man’ (Hornstein 1984) the objects needed to satisfy expressions like (19) are not simply *abstracta* developed for the purposes of scientific theorizing (Ludlow 2011: 135–136) intricately characterized by novel logics (Kennedy & Stanley 2009). Whatever abstract-concreta are, they are far less familiar than mere abstractions, and should be less palatable to the externalist.

But maybe our dislike for the bizarre is unwarranted, based on some vestigial aspect of our human conceptual machinery. Maybe these abstract-concreta exist, happily residing in the domain and validating the Realist’s predictions. Or maybe hybrid objects are more palpable than *Sinne*.³⁰ The Realist’s perverse predictions, however, do not end there. The externalist hypothesis predicts the existence of other hybrids, including hybrid properties, and relations given the lexical flexibility of expressions like ‘bilingual’, ‘lost’, and ‘defeat’.

In a conversation about the ability to acquire a second natural language, one might use the following English expressions:

particular culture might believe in (Ludlow 2011: 142).

To engage in the latter kind of investigation is to do psychology. But, the Realist is investigating the structure of the world, not the structure of our minds. As such, to indicate that the *theory* demands certain metaphysical posits is a burden the theory has to bear, not an analytic consequence of the fact that linguistic expressions have meanings. If (\mathcal{E}) is true, then there are such things as concrete *possibilia*. In this vein, some argue that we need even more things: situations (Barwise & Perry 1983), perspectives (Schein 2002), modes of presentation (Ludlow 1995), and fictional objects (Thomasson 1999). But we only need these things if the best theory of meaning is externalist—a *fortiori* that the theory requires a domain with such things counts against the merits of the theory. Adding to this list of posits the abstract-concreta required to address lexical flexibility, the class of entities required to support the externalist hypothesis looks less and less like the pedestrian objects of everyday experience. Concomitantly, the externalist hypothesis looks less like a plausible theory that makes use of the everyday objects we are familiar with, and more like a theory with implausible commitments.

³⁰ Notice how this vein of reply assumes a false dichotomy, that one either accepts mysterious, reified meanings as things, or accepts *whatever* things an externalist semantics requires. For other alternatives again see Hinzen (2007); Hinzen (2014); Pietroski (2005); Pietroski (2008); Pietroski (2010); Pietroski (forthcoming).

(21) The child is bilingual.

(22) John is bilingual.

Likewise, one could (roughly) conjoin the meanings of these two expressions into a single sentence using either of the following acceptable expressions:

(23) The child is bilingual and so is John.

(24) The child and John are bilingual.

Of course, an externalist semantics can accommodate the meanings of these expressions, whereby their truth-conditions are satisfied just in case there is a (salient) child, a John, and both of them instantiate the same particular property. To put the matter somewhat formally, (23) and (24) are true just in case:

(24') $\iota x. \exists y. \text{CHILD}(x) \ \& \ \text{JOHN}(y) \ \& \ \text{BILINGUAL}(x) \ \& \ \text{BILINGUAL}(y)$

In the expressions (21)–(24), the expression ‘bilingual’ has a single, univocal meaning, as reflected in the single truth-conditional predicate ‘BILINGUAL’. For the externalist this identifies some *single* property, say the property had by all things that acquire/speak/know two languages. As such the sentences in (21)–(24) are well captured by an externalist semantics.

But ‘bilingual’ is lexically flexible. While walking the streets of London, Ontario, Canada, I came across an empty box outside a franchise sandwich shop. Printed on the outside of this box was the expression ‘bilingual napkins’ which presumably identified the box’s contents as napkins on which information is printed in two different languages. Sitting in this franchise with a group of friends, and noticing the features of one such napkin, a competent speaker could well say to their compatriots the following acceptable sentence:

(25) The napkin is bilingual.

Supposing this group also knew our bilingual friend John, a competent speaker could say, and the group would no doubt accept as true, the following acceptable expression:

(26) The napkin is bilingual and so is John.

However, the following truth-conditions, which for the externalist are determined by the meaning of (26) are *not* satisfied in this situation:

(26') $\iota x. \exists y. \text{NAPKIN}(x) \ \& \ \text{JOHN}(y) \ \& \ \text{BILINGUAL}(x) \ \& \ \text{BILINGUAL}(y)$

Since both bilingual predicates in (26') are derived from the single use of ‘bilingual’ in (26), they must have the same truth-conditions. As such, for the externalist they must pick out the very same property. If that property pertains to the acquisition of multiple languages, the napkin clearly fails to instantiate such a property. And if the property pertains to the kinds of orthography printed on a thing’s exterior, John does not count as having such features. For the externalist, (26) turns out to be false in the situation described—the same situation that makes both (22) and (25) true. Given that (26), on an externalist semantics, has the meaning

attained by conjoining (22) and (25) (as indicated by the relationship between (21)–(23)), externalism should predict that competent speakers accept that (26) is true. That it does not is a problem for EM theories.

As before, the externalist could claim that ‘bilingual’ is homophonous, with two lexical entries ‘bilingual_a’ and ‘bilingual_o’ pertaining to the aforementioned acquisitional and orthographic properties (respectively). But just as in the example with ‘the Hirshhorn Museum’, this reply will not do. Whichever lexical entry the context demands, the truth-conditional predicates in (26’) will have the same satisfaction conditions as each other. The ‘bilingual_a’ interpretation of these predicates leaves the first unsatisfied by *the napkin*, while the ‘bilingual_o’ interpretation leaves the second unsatisfied by *John*. Either way, externalism makes the wrong prediction that, relativized to the situation, (26) is (treated as) false (by competent speakers). The externalist’s remaining response is to accept that there is some single bizarre, hybrid (or multifarious) property that admits to having shifting satisfaction conditions within a single context.

As with the case involving the flexibility of ‘the Hirshhorn-Museum’, the Realist could bite this bullet. She can simply accept that the domain contains not only hybrid abstract-concreta, but also hybrid properties such that the very same property can be instantiated in distinct ways by disparate objects within the same context.³¹ But again, as with ‘the Hirshhorn-Museum’, even this (bizarre) concession fails to explain the distribution of competent speaker judgments. Consider

(26) The napkin is bilingual and so is John.

(27)# The napkin and John are bilingual.

The expression in (27) is strange in a way that (26) is not. This strangeness, whatever it amounts to, seems to be a fact about the meanings of the constituent expressions of the sentence, since there is no general prohibition against joint predication (as exemplified in (24)). The syntactic frames of (23) and (24) are repeated in (26) and (27), yet the latter do not bear the same semantic relations to one another as the former, insofar as they are *not* synonymous (as should be clear by the oddity of (27)).

³¹ To be clear, this is not an injunction against multiple realizability. To take the paradigm case, mental properties like *belief* can be realized in Martian brains just as well as human ones. But what is instantiated in these distinct organisms is (say) an entity that plays a particular functional role in the mind of the organism, and indeed the *same* functional role. The worry here is not that the napkin and John embody different ways of instantiating the same property, but that they instantiate *different* properties. Contrast this difference with the manner in which distinct humans are bilingual. John Kerry and Nicolas Sarkozy are both bilingual (let’s suppose), yet the bilingual property is realized in their person in distinct ways. At the very least they differ regarding the languages in which they are fluent: Kerry is fluent in English and French, while Sarkozy is fluent in French and German. The manner in which Kerry and Sarkozy differ is quite clearly not the manner in which John and the napkin differ regarding bilingualism.

The externalist, and the Realist, in order to explain the distribution of competent speaker judgments for sentences containing nouns like ‘the Hirshhorn Museum’ and predicates like ‘bilingual’ are compelled to accept some rather bizarre entities into the domain of worldly things: hybrid objects that exhibit a concrete-abstract duality, and properties that have varying conditions for instantiation across instances within a single context. Neither of these concessions seems pleasant. Worse yet, conceding in these ways still leave unexplained aspects of competent speakers’ judgments, as can be seen by the contrast between (26) and (27).

Finally, consider the following:

(28) Henry lost his key, his lawsuit, and his job.

One can easily imagine a situation in which (28) is deemed both felicitous and true by competent speakers, whereby Henry’s misplacing his car key made him late to the courthouse, which led to his termination. Yet, Henry stands in very different relations to these three objects, all of which are expressed by the single use of ‘lost’ in the sentence. Henry is no longer in possession of his key, while his lawsuit, as a complex activity he participated in, has met with a particular conclusion. And of course, Henry no longer has a job in that a contract he had with some unmentioned individual is no longer binding. On an externalist account the worldly relation that ‘lost’ denotes requires that all three objects, abstract and concrete alike, must (when paired with *Henry*) satisfy that relation, despite these differences. And as with our previous examples, an appeal to ambiguity is not available to the externalist.

Let’s not forget that examples of this sort are numerous:

(29) The chef’s kitchen ran better than an imported car.

(30) Napoleon’s defeat was worse than Kasparov’s.


Given the acceptability of sentences like these, expressions like ‘lost’, ‘run’ and ‘defeat’ seem to exhibit lexical flexibility as well. If the Realist is forced to accept the existence of hybrid *relations* as a result of such flexibility, the pedestrian nature of the objects needed to accomplish the externalist’s aims is substantively undermined—especially if this bizarre ontology remains insufficiently explanatory. The point worth underscoring here is that the examples explored in this section are not isolated aberrations in a language that is otherwise well-modeled by externalism. Lexical flexibility is rampant in natural language expressions. To account for the manner in which natural language speakers treat these expressions, the externalist is compelled to accept into her ontology metaphysically bizarre objects, properties, and relations. Rather than viewing these ontological commitments as the price to be paid for an adequate theory of meaning, such requirements might be better seen, or so I contend, as a *reductio* against the externalist hypothesis that requires them.

4.3. An Externalist Reply

The externalist might reply to these worries by leveraging the purported virtues of an externalist semantics. She might contend that even if the flexibility cases rehearsed above require that she bear an ontological burden, such burden bearing is better than the alternative. As a part of an overall externalist theory, the virtues of the theory far outweigh these burdens. One principal virtue that supporters of externalism might trumpet is that understanding meanings in this way uniquely preserves semanticists' main source of linguistic data, the truth-value judgments of competent speakers. To quote a prominent externalist:

In short, intuitions about the truth and falsity of what is said by utterances of sentences have formed the data by which theorists have tested their hypotheses about meaning. There is no other obvious source of native speaker intuitions that are related to meaning. So if we did not have robust intuitions about the truth-conditions of our utterances, it would not be clear how to test such hypotheses; there would be no firm basis on which to construct a theory of meaning. (Stanley 2007: 6)

Consider for example, the landmark insight of Davidson (1967b) in treating the logical form of action sentences as involving quantification over *events*. The sentences in (31) display a particular pattern of inference, as indicated in Figure 1, wherein the arrows represent the direction of inferences that speakers of English are apt to make.



- (31) a. Brutus stabbed Caesar slowly and skillfully.
 b. Brutus stabbed Caesar slowly.
 c. Brutus stabbed Caesar skillfully.
 d. Brutus stabbed Caesar.

Davidson's proposal aims to capture these patterns of inference. Proposing that the logical forms in (31') are indicative of the truth-conditions of the sentences in (31), this approach captures the inferential

judgments of English speakers, by modeling these inferences as logical entailment.

- (31') a. $\exists e[STAB(e, Brutus, Caesar) \ \& \ SLOW(e) \ \& \ SKILL(e)]$
 b. $\exists e[STAB(e, Brutus, Caesar) \ \& \ SLOW(e)]$
 c. $\exists e[STAB(e, Brutus, Caesar) \ \& \ SKILL(e)]$
 d. $\exists e[STAB(e, Brutus, Caesar)]$

Because the proposed logical forms for the expressions in (31) quantify over *events*, the expression in (31'a) entails the other expressions in (31') by way of conjunction reduction. As such, the “diamond-shaped” inference patterns of speakers are captured by an *externalist* theory that takes events as the worldly-satisfiers of expressions.

This reasoning, contends the externalist, only makes sense if the *explanada* of the hypothesis are the inferential judgments of speakers, as judgments about the concomitant truth of collections of sentences. Thus the purported justification for an externalist semantics is that it maintains the theoretical import of speaker judgments. To deny that sentences have externalist meanings is to deny the connection between meaning and truth that renders these judgments worthy of capture. In this manner the externalist might contend, if meanings have nothing to do with truth, then these truth-value judgments are not indicative of expression meanings and of no use for semantic investigation.

However, denying the externalist thesis (\mathcal{E}) does not also require denying that the truth-value judgments of speakers are relevant data for the purposes of semantic theorizing. Externalists are committed to a *particular* relationship between truth and meaning—namely the one codified in (\mathcal{E}). In denying this, a semanticist need not deny that meaning is related to truth. She must simply deny that meanings *determine* truth-conditions. One can hold that natural language speakers can use sentences to make utterances that are true, and still deny the externalist thesis. And this can be done without denying that there is some “systematic” manner in which meaning is related to truth (*pace* Stanley 2007: 8). Such a view merely holds that the systematic manner in which linguistic meaning relates to the external world involves the interaction of multiple non-linguistic cognitive and external systems that connect in complex ways. This complexity might be systematic, but because many of the systems involved are extra-linguistic (and not semantic) the meaning of an expression will not, in the absence of this complex interaction, determine its truth-conditions.

One can, as Stanley does, amalgamate this motley group of disparate non-linguistic systems³² under the term ‘context’. And if by ‘con-

³² The diversity of components that collectively make up the “context” of an utterance, so construed, is important to note. Shared human systems that recognize gaze following, emotional facial gestures, object detection, agency detection, and many others, not to mention the external physical systems that govern “normal” visual and auditory environments, all fall under the “context” that determinately links meanings to truth. But if one wants to know how meanings differentially interact with these various systems in order for a speaker to utter something true,

text’ one includes whatever is needed to derive truth-conditions from meanings, then trivially the meaning of an expression (plus the “context” of its use) determines its truth-conditions. But, the gap between what a sentence means and what a speaker communicates in conveying that meaning *via* a linguistic utterance admits to some marked complexity—as noted in §4.1. Respecting, and not merely masking, the complexity of this relationship between the meaning of a sentence and the truth of an utterance not only preserves the (nuanced) use of truth-value judgments as linguistic data, but it opens up *new* sources of data (Pietroski et al. 2009; Lidz et al. 2011; Vogel et al. 2014). So, far from making the semanticist’s task impossible (or without basis) denying (\mathcal{E}) expands the data-set for the theoretician, while preserving the utility (though augmenting the informativeness) of speaker judgments.

5. *Mental Content*

In the previous two sections we saw arguments that highlight the difficulty in accepting externalist accounts of language (EL views), and externalist theories of meaning (EM views). Arguments against an EL view push an externalist to adopt externalist views about *mental content*. If languages are not external objects, but rather aspects of the human mind/brain, then to rescue the EM view an externalist is committed to the view that the content of mental representations can be characterized externally, as relations between representations (or concepts) and worldly objects. As such, some arguments for internalism address both externalist theories about mental content, and the relationship between mental content and linguistic meaning. In this section I present these concerns.

5.1. *Naturalist Theories of Content*

Both sentences and thoughts seem to be about the world, and thereby exhibit intentionality. The close proximity of these disciplines gives rise to a simple solution to the problem of intentionality for language. Namely, that the problem of intentionality is solved at the level of thought, not language. An enticing view about the relationship between thought and language is that the contents and structure of our thoughts are merely mirrored in language. If the structure of natural language mirrors the structure of thought, wherein an expression in a language is merely a way of making public some particular thought composed of conceptual content, then the intentionality (and meaning) of an expression simply tracks that of the concepts used to compose the expressed thought. On such a view, natural language expressions are merely labels for thoughts, and likewise, words are merely labels

abstracting over these differences by indicating that the context somehow fills this gap is no answer at all. Worse yet, it commits one to a theory of meaning that is thereby incapable of addressing such questions.

for concepts, as a way of making them articulable. Call this the *label theory* of linguistic meaning.³³

So long as the language I speak syntactically composes in a way commensurate with the structure of my thoughts, linguistic meaning would perfectly mirror conceptual content. On this proposal, the meanings of our expressions would hook up with the world *via* conceptual content, so long as conceptual content can be characterized externally. So, to the degree that our concepts align with the “fine structure” of the world, expressions of a natural language will likewise accord with the mind external objects of the world.

But why should one insist that linguistic meaning is mediated by our conceptual system? For one, this answers the problem of intentionality at the level of language. But more importantly for the Realist, the desire here relates to naturalistic explanation. As we saw, a substantive source of contention in thinking about linguistic meaning as externalist relates to the aims of the scientific enterprise of linguistics. Facts about the acquisition and productivity of language in humans deserve explanation, and a theory of meaning ought to add to (or at least make possible) an explanation of these facts. As we saw in §3.1 these considerations strongly suggest that the object of study for a naturalistic investigation of language is in the mind. As such, to the degree one thinks that language and thought are independent, a naturalistic inquiry into these matters will address the way in which these distinct mental faculties interact.

There is no shortage of literature addressing the viability of naturalistic accounts of content.³⁴ Whether or not naturalistic accounts of intentionality are viable is beyond the scope of this work, but the point I want to emphasize here is that the force of the arguments presented so far against externalism compel the externalist to adopt two contentious views: a labeling theory of linguistic meaning, and a naturalistic account of intentionality. I'll not take the time to illustrate the contentiousness of the latter,³⁵ but the former position is worth analyzing, in part because so many philosophers seem to adopt this view without much defense.

³³ Fodor (1975) seems to hold this view. Jackendoff (2002) explicitly adopts this position, though not by this name. Oddly enough he defends an internalist proposal for linguistic meaning on the basis that no naturalistic account of external mental content is plausible.

³⁴ For a good survey see the introduction to Macdonald & Papineau (2006).

³⁵ This has been done by many, and better than I could hope to do here. See Jackendoff (2002); Loewer (1997); Boghossian (1991); Godfrey-Smith (1989); and McGinn (1982).

5.2. Labeling Theory of Meaning

The labeling theory of meaning views the relationship between words and concepts as one of labeling, whereby words are like labels for concepts. My use of ‘cat’ is just a way of tokening the CAT concept in my audience, whose content serves as the meaning of ‘cat’ in any expression that uses it, like ‘The cat is on the mat’. Linguistic meaning on this view is just conceptual content. Such a view (as indicated above) can still be meaningfully externalist, if the contents of our concepts are externalist. This labeling view is pervasive amongst philosophers. In fact, the view is often adopted as obvious, without much need to articulate that indeed adopting the view embodies a collection of commitments about the relationship between the human language faculty and the conceptual system. Burge (1975) is a paradigmatic example. In discussing the expansiveness of his famous ‘arthritis’ case, Burge writes:

On the other hand, the [arthritis] thought experiment does appear to depend on the possibility of someone’s having a propositional attitude despite an incomplete mastery of some *notion* in its content ...Suppose a subject thinks falsely that all swans are white ...that ‘swan’ means ‘white swan’ (Burge 1975: 83) (my emphasis)³⁶

Burge treats the content of a concept, or *notion*, which plays an important role in determining the content of the propositional attitude someone might hold, as no different than the meaning of a natural language word: hence the notion SWAN has the same meaning as ‘swan’. In his book-length critique of Burge’s account of wide-content Segal (2000) commits to this same theory about the relationship of words to concepts.

Zowie and Twin Zowie both say “My engagement ring is studded with diamonds.” Are the concepts expressed by their words “diamond” the same? (Segal, 2000, p. 6)

...

- Let w be the focal word
- Let c be the concept [the subject] expresses by w .

(Segal 2000: 67)

In more contemporary literature, Weber (2005) writes

The meaning of the term “gene” has changed several times in the history of twentieth-century genetics. If we distinguish between a term’s sense and its reference, it is possible that the term’s sense has changed, but not its reference...I have examined both the reference potential and the reference connected with different historical versions of the gene *concept*. (my emphasis) (Weber 2005: 228)

As a final example, consider this passage from Clark & Chalmers (1998), wherein they discuss an opponent’s (possible) response to their thought experiment involving Otto, and his purported belief about the location of a museum:

³⁶ Here Burge uses ‘notion’ as a way of talking about the content of a concept: “Talk of notions is roughly similar to talk of concepts in an informal sense”(Burge 1975: 83).

An opponent might put her foot down and insist that as she uses the term ‘belief’, or perhaps even according to standard usage, Otto simply does not qualify as believing that the museum is on 53rd street. We do not intend to debate standard usage; our broader point is that the *notion* of belief...(my emphasis) (Clark & Chalmers 1998: 14)

Here the slide from the meanings of a linguistic expression ‘belief’ (as evident in its “standard” usage) to the concept (or notion) of belief, is clear, blatant, and offered without explanation or defense.

Linguistic meaning, on this often assumed view, is simply conceptual meaning. But this view has the following consequences: first, the extension of our words must have the same extension as their underlying conceptual meanings, and second, the syntax of natural language must be mirrored in the composition of thought. Neither of these consequences seem well supported by the way natural language speakers treat the meanings of expressions.

The flexibility of natural language expressions speaks against the first consequence.

(19)The Hirshhorn-Museum is bankrupt and it is a cylinder.

If the extension of the concept HIRSHHORN-MUSEUM is to capture the meanings that natural language users apply to the term ‘Hirshhorn-Museum’ then the extension of the concept better include both the concrete building that houses artworks, and the abstract institution that employs hundreds of people. As we’ve seen, many natural language expressions bear meanings that do not track the domain of objects in this way. Thus whatever thought corresponds to (19), and thereby stands as the meaning of (19), it must either treat ‘Hirshhorn-Museum’ as labeling two distinct concepts, or have a content such that some (abstract) object (or some building) is both cylindrical and bankrupt. The former avenue belies the manifest relationship between the uses of the English term, and fails to account for the felt relatedness of these uses. The latter option has much more bizarre metaphysical commitments, since we have no other reason (other than a commitment to particular views about semantics) to postulate such an entity.

Turning to the second consequence, if language mirrored the structure of thought, then thoughts should compose much the way expressions do. More strictly, the meanings of linguistic expressions and their underlying logical forms should mirror the structure of the concepts those forms express. The deep structural syntactic frames that make up interpretable expressions in a natural language must mirror the structure of their underlying concepts. In this vein, consider the following sentence:

(32)Wilbur kicked Fred.

Paying attention to the syntax of this construction, and adopting the labeling theory commitment, we ought to conclude that the KICK concept is dyadic. The word ‘kicked’ in the complete expression in (32) takes a subject and an object, and likewise we would expect the related concept

to take two elements to form a complete thought. Thus the thought expressed must make use of a concept like:

(33) KICK(*_s_ _o*)

which when saturated with two elements, makes the complete thought

(34) KICK(*Wilbur, Fred*)

However, if this dyadic notion of KICK is supposed to underlie all meaningful uses of 'kick', as implied by the labeling theory, the following expression is an apparent counterexample:

(35) Wilbur kicked Fred with his foot.

Given the syntactic structure of the expression in (35), the underlying conceptual meaning must have a *triadic* structure, to make room for the instrument used in the kicking:

(36) KICK(*_s_ _o _i*)

For any way of differentiating concepts, surely addicity falls under the identity condition for a given concept. That is, concepts with different addicities must be *different* concepts. So, (33) and (36) cannot be the same concept. As such, the meaning of 'kicked' in (32) and (35) is different on the labeling view, insofar as 'kicked' labels concepts that are (of) different (addicities). This entailment leaves unexplained why a competent speaker would find that both (32) and (35) are felicitous descriptions of the same kicking.

A defender of the labeling view might hold that really we have only one KICK concept, with sufficient addicity to accommodate all uses of 'kick', and thereby holding their meaning constant across various uses. As such, the concept in (36) is the only KICK concept, made use of in expressions where the instrument of the kicking is unmentioned.

There are three problems with that response: first, this requires that many expressions that make use of the transitive 'kick' have implicit content of an unspoken instrument. And there seems to be no syntactic evidence that such expressions have any such implicit content. Second, such a triadic concept will not capture the meanings of expressions like

(37) Wilbur kicked Fred the ball.

(38) Wilbur kicked Fred the ball with his toe.

The underlying conceptual meaning for (38) must have a *tetradic* addicity, to make room for the indirect and direct objects in the syntactic structure of the expression. Insisting here that the single conceptual meaning for 'kick' is a *tetradic* concept is implausible. While one might entertain the plausibility that transitive uses of 'kick' leave some unspoken implicit content about what instrument was used in a given kicking, surely such uses do not leave the existence of (nonexistent??) indirect objects implicit, as would be required if (39) was the underlying conceptual meaning of (32').

(39) KICK(*WILBURs, FREDdo, NOTHINGo, FOOTi*)

(32') Wilbur kicked [nothing to] Fred [with his foot].

And plainly, (32) does not mean what (32') means.

The third problem for the super-addicity move pushes in the opposite direction. Sentences like (40) seem to require conceptual meanings that are monadic.

(40) Wilbur kicked.

Insisting here that the expression in (40) really contains implicit content that reflects the underlying tetradic conceptual structure in (39) strains good explanation.

Of course, the label theorist could respond to this data by treating all this as evidence that there are really multiple words 'kick' with multiple KICK concepts as their meanings. There are, on this reply, multiple homophones 'kick' each paired with a different concept depending on whether they have direct objects, instruments, and/or indirect objects. However, such a response treats the difference between these uses of 'kick' like the difference between 'kick' and 'punch'—they are different words, with different conceptual meanings. This of course leaves unexplained what is obvious, that the many uses of 'kick' describe quite similar actions, and are conceptually related. The events these various uses of 'kick' describe bear striking features in common—those features that make them plausible kickings in the first place. A theory about the relationship between the meaning of words and the content of concepts that leaves such basic facts unexplained is troubled.

The point then is this: the relationship between linguistic meaning and conceptual content is not nearly as simple as the labeling theory would have it, as the case of (the addicity of) 'kick' and KICK shows. And addicity is but one feature of the relationship between lexical items and concepts that admits to some *prima facie* complexity.³⁷ For the externalist, this should be troubling news, since this means that the path from the meaning of a term, through the content of the associated concept, to its worldly extension is rather complex. In short, the IL-EM

³⁷ Other aspects of the content of our concepts, apart from their structure, highlight the complex connection between words and the concepts that underwrite them. Consider the different ways in which the predicate 'is blue' applies to objects, and what this says about the complex application of the blue concept.

- (1) The house is blue.
- (2) The marker is blue.
- (3) The iris is blue.
- (4) The sky is blue.

The truth-conditions for the color predicate in (1) that would make it true of some house, would not, when applied to some marker, make (2) true, despite the fact that they appear in the same syntactic frame. So if the meaning of expressions are a result of the satisfaction conditions of their underlying concepts, the satisfaction conditions for the thoughts BLUE(HOUSE) and BLUE(MARKER) are not merely going to be attributable to the differences in the extension of HOUSE and MARKER—*mutatis mutandis* for (3). And while many uses of (4) are considered true by competent speakers, what is far from clear is which object is picked out such that it satisfies 'is blue' in any of the ways just mentioned here.

theorist is committed to two views about language and content that are contentious, one of which (given our discussion here) seems implausible. Not only are such theorists saddled with giving a sufficiently plausible *naturalistic* account of mental content, they are also saddled with the troubles articulated here for the label theory of linguistic meaning.

6. *Natural Language and Ontology*

The arguments thus far have been multifaceted, but direct. What they have shown is that externalist proposals about the nature of languages as objects, and about linguistic meaning—in particular (\mathcal{E})—are difficult theses to defend. In this closing section I'll indicate how this difficulty undermines the Realist's metaphysical methodology. Primarily, if (\mathcal{E}) is dubious, then appealing to the truth-conditions of natural language expressions, or the satisfaction conditions of their purported conceptual meanings, as justification for metaphysical conclusions is without foundation.³⁸

The problems detailed thus far for (\mathcal{E}) undermine the fruitfulness of the Realist's default metaphysical methodology.³⁹ Ontological investigation proceeds by analyzing natural language usage. Roughly, the Realist makes use of her competence with a given natural language, since such competence ensures that she understands the meanings of natural language expressions. Under the guise that such meanings are externalist, she derives from them the ontology one is committed to in accepting the truth of a given expression. Taken together, the ontological commitments of all the true sentences determine what there is. In the opening sections of this paper I sketched the Realist's methodology.

³⁸ This is especially true if the human ability to construct complex thoughts from different conceptual domains depends on the human language faculty (Carruthers 2002; Spelke 2003; Jackendoff 1990, 1996; Bloom 2000). Even if we grant that our concepts have satisfaction conditions that accord with the structure of reality, once those concepts are put to work by the language faculty in building meanings to sentences, there's no assurance that the content of the resulting construction will retain such a tight connection to the world (see Glanzberg 2011). And of course there is no guarantee that our concepts accord with reality's structure. The work of Michotte (1946 [1963]) illustrates the difficulty of such certainty with regard to our judgments of causation, where clearly non-causal scenes are judged by subjects as exhibiting causation. *A fortiori* these judgments persist even when objects interact in ways nearly identical to clearly non-causal events (Scholl & Nakayama 2002).

³⁹ Such a position is Realist because it holds that there is an objective structure to the world. However, the view is also importantly Quinean, in the sense that it adopts a methodology of looking to language, and to quantificational structure in particular, to settle ontological disputes. While these two positions seem interrelated, there is reason to think neither entails the other (Hirsch 2002). The arguments outlined here do not undermine Realism *tout court*, even if they directly undermine the Quinean Realist position adopted by many philosophers and metaphysicians. For these reasons I've highlighted the distinction here, though throughout the paper, I use 'Realist' and its neighboring forms as short hand for the more cumbersome and possibly misleading 'Quinean Realist'

I'll close here with a restatement of that widely-adopted strategy and articulate the reasons to reconsider its merits.

In §1 we saw that the Realist holds that there is a unique language (an interpretation of the existential quantifier) whose quantificational structure mirrors the structure of reality.⁴⁰ Put more explicitly, Realist

...inquiry will be guided by ...[an] assumption [that] modern logic's quantificational apparatus mirrors the structure of reality: I assume an ontology of *things*. Moreover, I assume that there is a single, objective, correct account of what things there are. (Sider 2002: xvi)

This account is captured by the meanings of expressions in (what I'll call) the Language of Ontology, or \mathcal{L}_O . With this privileged language in hand metaphysicians can proceed to answer ontological questions by investigating the meanings of expressions in that privileged language, which can be given by way of Tarskian satisfaction by sequences of domain objects. Thus the Realist holds, the objects required to account for the meanings of the true expressions in \mathcal{L}_O are the objects of reality, since this privileged language mirrors reality's (object-based) structure.

To highlight an example of this strategy in action, consider the following points made by Sider (2002) in arguing against certain conceptions of time:

The status of tense is a second issue in the philosophy of time. Tensed sentences are those which presuppose a certain position or vantage point within the whole of time, for example:

It is *now* raining.

It *was* the case that there existed dinosaurs.

I *will* one day visit Utah. (Sider 2002: 12)

In arguing against a presentist theory of time, Sider contends that the presentist cannot clearly account for the truth of sentences that (seemingly) refer to the non-present. Insofar as the presentist denies that there are any ontologically real past or future times, *any* sentence that requires the existence of past/future times must thereby be either meaningless, or simply false. Such sentences have no truth-makers given the presentist's ontology, and thus the presentist cannot account for the truth of tensed sentences.

The success of this argument clearly presumes that the meanings of these natural language expressions determine their truth-conditions, insofar as their meaningfulness depends on the existence of past/future times. The presentist denies that there are past/future times. If the sentences Sider presents are meaningful—which they surely are, given that competent speakers of English have no problem understanding

⁴⁰ Sider expresses the same Realist commitment with different, less perspicuous language elsewhere:

Clearly there are multiple (inferentially and materially adequate) interpretations of quantifiers. As I see it, the real issue is whether any of these interpretations is metaphysically distinguished, whether any of them uniquely matches the structure of the world, whether any carves nature at the joints better than the others. (Sider 2009: 392)

them—then *ex(ternalist) hypothesisi* they have truth-conditions. Those conditions are only satisfied if there is some future time where Sider is in Utah, and some past time where dinosaurs are alive and well. The presentist, contends Sider, must admit then that all tensed sentences are false or meaningless, since they have no temporal satisfiers. This consequence thereby seems bad for the presentist.

Of course, a key step in this line of argumentation asserts the truth of (E). As we have seen, this externalist hypothesis is troubled. If the meaning of tensed terms do not determine whether or not they refer to *times* (*pace* externalism), then the move from linguistic meaning to ontological commitment is without warrant. And the supposition, if not false, is (at least) difficult to defend in light of the flexibility of natural language.

However, the Realist has a ready (and plausible) response to this objection. After all, natural languages like English are awash with vexing semantic properties like vagueness, ambiguity, and (apparently) lexical flexibility. As such, there is little surprise that they are ill-suited for the purposes of ontological investigation. The language the Realist needs is one that conforms to the features of classical logic, and none of these semantic properties are tolerated by such logics. But some languages are not deficient in these ways—namely the languages invented in the process of scientific inquiry. On this reply, the privileged language \mathcal{L}_O needed for ontological investigation is the one preferred by our best sciences. After all, scientific inquiry is guided by the expressed purpose of perspicuously describing the world. This process involves making decisions about what terms to use. The results of such inquiry are languages that embody the kind of precision that natural languages like English lack. This embodiment makes these scientific languages better suited for ontological investigation, and thereby better candidates for \mathcal{L}_O .

Sider himself indicates as much:

I hold that the fundamental is determinate...First, no special-purpose vocabulary that is distinctive of indeterminacy ...carves [nature] at the joints. Second, fundamental languages obey classical logic. (Sider, 2011, p. 137)

The fundamental structure of reality does not admit to the fuzziness typical of natural language meanings. Likewise, fundamental languages, those that cleave to the structure of the world, are free of such properties. Their constituents have determinate meanings/referents. The privileged \mathcal{L}_O , the language that mirrors the fundamental structure of reality, should not admit to the kind of indeterminacy we find with natural languages.⁴¹ Instead, this privileged language should have the

⁴¹ Technically, \mathcal{L}_O must typify a class of languages, not a single language. Suppose \mathcal{L}_O contains the terms of imperial measurement. Such a language can be translated into one that uses a metric system. Neither language would cleave to realities structure any better than the other. *Mutadis mutundis* for the many logical operators. Sider's focus is on the interpretation of the quantifiers, and in particular, the existential quantifier. On this point, all the various languages in the \mathcal{L}_O class would agree, for the Realist.

semantics of a Tarskian logic, whereby domain objects constitute the meanings of its expressions.

According to Sider, the process of scientific inquiry yields languages of this ontologically determining sort:

We should believe generally what good theories say; so if a good theory makes an ontological claim, we should believe it. The ontological claim took part in a theoretical success, and therefore inherits a borrowed luster... [But] the conceptual decisions ...also took part in a theoretical success, and also inherit a borrowed luster. (Sider 2011: 12)

Sider suggests that the languages used to express our best scientific theories are better suited for questions of ontology, and that certain sciences do this better than others (Sider 2011: 6). The substantive assumption endorsed by this Realist is that scientific methodology is of sufficient epistemic heft that the languages our best sciences construct are those that match the quantificational structure of the world. And because the languages used to state our best scientific theories are designed to perspicaciously describe the world, the epistemic credentials of naturalistic inquiry assures us that such languages make use of externalist meanings. The value of this move to such languages is measured by the degree to which scientific methods deliver languages that mirror the structure of reality. I've argued elsewhere that this move might be suspect, as terms in core sciences seem to exhibit lexical flexibility in much the way natural languages do (Vogel, under review). Nonetheless, the Realist that embraces this reply must thereby abandon the use of natural languages for the purpose of ontological investigation. The very insistence that there is a distinction between natural and scientific languages (in terms of their ontological credentials) calls into question the use of natural language speaker judgments to adjudicate ontological matters.

It's worth noting how impactful this Realist retreat to scientific languages is. Much of philosophical discourse in both metaphysics proper, and in other philosophical domains where ontological questions seem to matter, relies on the use of natural language expressions, and the purported ontological commitments of speakers that endorse those expressions as true. The example above with tensed sentences is commonplace. Stotz et al. (2004) summarizes the sort of strategy philosophers deploy in pursuing their metaphysical projects:

[In analyzing concepts] [t]raditionally, philosophers have relied on their individual linguistic competence with the corresponding words. When analyzing a concept, the philosopher treats him or herself as a sociolinguistic 'sample of one' Stotz et al. (2004).

This kind of *Conceptual Analysis* tends to proceed by offering up a short *natural language* description of a situation, and then probes whether or not the situation that meets the truth-conditions of the description also serves as a satisfier of the term under dispute. To take an example of this method, consider this case from Lewis (2000) in his discussion about the metaphysics of causation:

[ROCKS]

Billy and Suzy throw rocks at bottles. Suzy throws first, or maybe throws harder. Her rock arrives first. The bottle shatters. When Billy's rock gets to where the bottle used to be, there is nothing but flying shards of glass. [...]
So Suzy's throw causes the shattering. Billy's doesn't. (Lewis 2000: 184)

The familiar method used here is to present a case which makes use of the metaphysical notion in question, and in light of the readers' comprehension of the passage, leverage their intuitive judgments about the described case with regard to that notion. Here, the case is presented to show a flaw in a simple counterfactual notion of causation, and lend support to an ancestral-counterfactual account. The judgments of natural language speakers, namely that the expression 'Billy caused the bottle to break' is false, plays an evidential role in Lewis' argument. This judgment is offered as evidence that the counterfactual conception of *causation*⁴² is troubled. The supposition is that the same situation which makes the claims in ROCKS true, makes the causal claim about Billy and his rock false. As such, the *causal structure* of the situation is not captured by a theory that belies this speaker judgment—a judgment that can only serve as evidence about *causation* if words like 'cause' refer to elements of causal structure. To assume that words so refer is just to accept the troubled externalist hypothesis (E).

Importantly Conceptual Analysis is not of a kind with a naturalistic approach to language. As Chomsky (1965) describes the related methodology for linguistic inquiry, subjects' judgments of the acceptability of a sentence need to be captured by a theory of grammar, with the hypothesis being that the acceptability of those judgments is *explained* by violations of grammatical rules. This makes sense if the *explanandum* is an aspect of the human mind/brain, since the judgments of natural language speakers are (hypothesized by the Chomskyan to be) the product of a mental faculty which includes algorithmic rules for constructing sentence—a grammar. Analogous inquiry in semantics tends to pertain to judgments about entailment, the felicity of a series of descriptions to a scene, or judgments about whether expression pairs have similar meanings. What linguists do not do is introspect into the meaning of words in a language, and certainly not with an eye toward answering ontological questions. Such a methodology plainly makes little sense if one's goal is to describe anything other than the way a particular subject views the world. Only with the added externalist

⁴² Importantly Lewis takes cases like this to illuminate the nature of causation, *not* the semantics for the English expression 'cause', as exhibited by the judgments of competent English-speakers (Lewis 1973). He addresses this question in a footnote, indicating that his proposal regards causal *facts*, *not* linguistic objects. Further, Collins et. al. echo this goal in their introductory contribution to a prominent volume on the metaphysics of causation. In fact, they indicate that the central misstep of a competing analysis defended by Davidson (1967a) is the focus on sentences instead of propositions (Collins et al. 2004: 17). They insist that the evidence brought to bear by dissecting cases, in the manner above, informs us about *propositions* and causal *facts*, not merely linguistic expressions that invoke the term 'cause'.

assumption that the meanings of expressions have real-world denotational meanings could such a methodology sensibly be applied for doing metaphysics. But it is precisely that externalist assumption that I've argued is problematic.

Nonetheless, debates about mental content (Burge 1975, 1979), the ontology of minds (Clark & Chalmers 1998), persons (Parfit 1984), causation (Lewis 1973; Collins et al. 2004), identity (Black 1952), modality (Plantinga 2003), rationality (Williams 1979), moral theory (Foot 1967), and many others require that natural language intuitions play a profound evidential role in settling ontological questions. But if natural language meanings fail to determine the truth-conditions for the expressions they serve, then the ontological commitments of speakers that endorse those expressions as true are (at best) indeterminate. Conceptual Analysis is grounded on the presumption that the truth-value judgments of competent speakers of a natural language in which the proposed thought experiment is written have ontological commitments. The Realist retreat to scientific languages acknowledges the deficiencies of natural languages for this purpose. Such a Realist, in the absence of an adequate response to the worries presented here regarding (\mathcal{E}), should abandon Conceptual Analysis as a method of ontological investigation.

The arguments presented do not rule out the Realist's methodology full stop. A Realist might well respond to these worries by abandoning their use of natural languages for languages more amenable to externalist treatment. In particular, they might adopt the suggestion made by Quine (1960: 221) and defended by Sider (2011), and look to the invented languages used to express scientific theories—languages constructed to avoid the pitfalls of natural language.⁴³ That is, the Realist might contend that while the meanings of natural language expressions do not determine their referents, terms in a scientific language do, insofar as such languages do not suffer from the vagueness and flexibility of natural language expressions.⁴⁴ For this Realist, the privileged language of ontology (\mathcal{L}_O) is a formal language, developed to express our best scientific theories. However, given both the naturalistic commitments of the Realist, and the flexibility natural languages exhibit, the arguments presented here suggest that \mathcal{L}_O cannot be a natural language, or some regimented variant of one. Consequently, the metaphysical methodologies that assume the externalist hypothesis for natural languages, like Conceptual Analysis, are without foundation, and should be abandoned as means of settling ontological disputes.

⁴³ Hence the meaning at the heart of Quine's proclamation "Language is conceived in sin and science is its redemption" (Quine 1973: 68)

⁴⁴ Though this contention too might seem dubious when one considers the various uses of biological terms like 'gene' (Weber 2005; Beurton et al. 2000; Stotz & Griffiths 2004; Stotz et al. 2004; Wilson et al. 2007).

References

- Austin, J. L. 1962. *How to Do Things with Words*. Oxford: Oxford University Press.
- Bach, K. 1999. "The semantic-pragmatic distinction: What it is and why it matters." In K. Turner (Ed.) *The Semantic-Pragmatics Interface from Different Points of View*. Oxford: Elsevier: 65–84.
- Bach, K. 2004. "Minding the gap." In *The Semantics/Pragmatics Distinction*. Stanford: CSLI Publications: 27–43.
- Baker, M. 2002. *The Atoms of Language*. New York: Basic Books.
- Barwise, J., & Perry, J. 1983. *Situations and Attitudes*. Cambridge, Mass.: Bradford Books, MIT Press.
- Beurton, P. J., Falk, R., & Rheinberger, H.J. 2000. *The Concept of the Gene in Development and Evolution: Historical and Epistemological Perspectives*. Cambridge: Cambridge University Press.
- Black, M. 1952. "The identity of indiscernibles." *Mind* 61 (242): 153–164.
- Bloom, P. 2000. *How Children Learn the Meaning of Words*. Cambridge, Mass.: MIT Press.
- Boeckx, C. 2006. *Linguistic Minimalism : Origins, Concepts, Methods, and Aims*. Oxford: Oxford University Press.
- Boghossian, P. A. 1991. "Naturalizing content." In B. M. Loewer, & G. Rey (Eds.) *Meaning in Mind: Fodor and His Critics*. Oxford: Blackwell.
- Bromberger, S., & Halle, M. 1995. "The ontology of phonology." *Studies in Logic and the Foundations of Mathematics* 134: 725–743.
- Burge, T. 1975. "On knowledge and convention." *Philosophical Review* 84 (2): 249–255.
- Burge, T. 1979. "Individualism and the mental." *Midwest Studies in Philosophy* 4 (1): 73–122.
- Cappelen, H., & Lepore, E. 2005. *Insensitive Semantics: A Defence of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell Publishing.
- Carey, S. 2009. *The Origin of Concepts*. Oxford: Oxford University Press.
- Carruthers, P. 2002. "The cognitive functions of language." *Behavior and Brain Sciences* 25: 261–316.
- Chomsky, N. 1965. *Aspects of a Theory of Syntax*. Cambridge, Mass.: MIT Press.
- Chomsky, N. 1975. *Reflections on Language*, chap. The Object of Inquiry. Pantheon Press.
- Chomsky, N. 1977. *Essays on Form and Interpretation*. Oxford: Elsevier.
- Chomsky, N. 1981. *Lectures on Government and Binding* Dordrecht: Foris.
- Chomsky, N. 1986. *Knowledge of Language*. New York: Praeger Publishing.
- Chomsky, N. 2000. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Clark, A., & Chalmers, D. 1998. "The extended mind." *Analysis* 58 (1): 7–19.
- Collins, J., Hall, N., & Paul, L. A. 2004. "Counterfactuals and causation: History, problems, and prospects." In J. Collins, N. Hall, & L. Paul (Eds.) *Causation and Counterfactuals*. Cambridge, Mass.: MIT Press: 1–57
- Davidson, D. 1967a. "Causal relations." *Journal of Philosophy* 64 (21): 691–703.

- Davidson, D. 1967b. "The logical form of action sentences." In N. Rescher (Ed.) *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.
- Dummett, M. 1978. "What do I know when I know a language." In Dummett (1993). Oxford: Oxford University Press.
- Dummett, M. 1993. *The Seas of Language*. Oxford: Oxford University Press.
- Fodor, J. A. 1975. *The Language of Thought*. Cambridge, Mass.: Harvard University Press.
- Foot, P. 1967. "The problem of abortion and the doctrine of double effect." *Oxford Review* 5: 5–15.
- girontravel.se 2013. The town of Kiruna. <http://girontravel.se/en/fun-facts/the-town-of-kiruna>
- Glanzberg, M. 2011. "Meaning, concepts, and the lexicon." *Croatian Journal of Philosophy* 11 (1): 1–29.
- Godfrey-Smith, P. 1989. "Misinformation." *Canadian Journal of Philosophy* 19 (4): 533–50.
- Grice, H. P. 1957 [1989]. "Meaning." In H. P. Grice (Ed.) *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press: 213–223.
- Hinzen, W. 2006. "Internalism and truth." *Mind & Society* 5: 139–166.
- Hinzen, W. 2007. *An Essay on Names and Truth*. Oxford: Oxford University Press.
- Hinzen, W. 2014. "Recursion and truth." In T. Roeper, & M. Speas (Eds.) *Recursion: Complexity in Cognition*, vol. 43 of *Studies in Theoretical Psycholinguistics*. Springer International Publishing: 113–137.
- Hirsch, E. 2002. "Quantifier variance and realism." *Philosophical Issues* 12 (1): 51–73.
- Hornstein, N. 1984. *Logic as Grammar*. Cambridge, Mass.: MIT Press.
- Jackendoff, R. 1990. *Semantic Structures*. Cambridge, Mass.: MIT Press.
- Jackendoff, R. 1996. "Semantics and cognition." In S. Lappin (Ed.) *The Handbook of Contemporary Semantic Theory*. Oxford: Blackwell Publishing: 539–559.
- Jackendoff, R. 2002. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford: Oxford University Press.
- Kaplan, D. 1977 [1989]. "Demonstratives". In J. Almog, J. Perry, & H. Wettstein (Eds.) *Themes from Kaplan*. Oxford: Oxford University Press: 481–563
- Kennedy, C., & Stanley, J. 2009. "On 'average'." *Mind* 118 (471): 583–646.
- Lewis, D. 1970. "General semantics." *Synthese* 22: 18–67.
- Lewis, D. 1973. "Causation". *Journal of Philosophy* 70 (17): 556–567.
- Lewis, D. 1975. "Languages and language." In K. Gunderson (Ed.) *Language, Mind, and Knowledge*. University of Minnesota Press.
- Lewis, D. 1986. *On the Plurality of Worlds*. Oxford: Blackwell Publishing.
- Lewis, D. 2000. "Causation as influence." *Journal of Philosophy* 97 (4): 182–197.
- Lidz, J., Pietroski, P., Hunter, T., & Halberda, J. 2011. "Interface transparency and the psychosemantics of most." *Natural Language Semantics* 19 (3): 227–256.
- Loewer, B. M. 1997. "A guide to naturalizing semantics." In C. Wright, & B. Hale (Eds.) *A Companion to the Philosophy of Language*. Oxford: Blackwell: 108–126.

- Ludlow, P. 1995. "Logical form and the hidden-indexical theory: A reply to Schiffer." *The Journal of Philosophy* 92 (2): 102–107.
- Ludlow, P. 2003. "Referential semantics for i-languages?" In L. M. Antony, & N. Hornstein (Eds.) *Chomsky and His Critics*. Wiley Online Library: 140–161.
- Ludlow, P. 2011. *The Philosophy of Generative Linguistics*. Oxford: Oxford University Press.
- Macdonald, G., & Papineau, D. (Eds.) 2006. *Teleosemantics*. Oxford: Oxford University Press.
- McGinn, C. 1982. "The structure of content." In A. Woodfield (Ed.) *Thought and Object*. Oxford: Oxford University Press.
- Michotte, A. 1946 [1963]. *The Perception of Causality*. New York: Basic Books.
- Parfit, D. 1984. *Reasons and Persons*. Oxford: Oxford University Press.
- Piattelli-Palmarini, M., Uriagereka, J., & Salaburu, P. 2009. *Of Minds and Language: A Dialogue with Noam Chomsky in the Basque Country*. Oxford: Oxford University Press.
- Pietroski, P. M. 2005. "Meaning before truth." In G. Preyer, & G. Peter (Eds.) *Contextualism in Philosophy*. Oxford: Oxford University Press: 255–302.
- Pietroski, P. M. 2008. "Minimalist meaning, internalist interpretation." *Biolinguistics* 4: 317–341.
- Pietroski, P. M. 2010. "Concepts, meaning and truth." *Mind & Language* 24 (3): 247–278.
- Pietroski, P. M. forthcoming. *Conjoining Meanings: Semantics Without Truth-Values*. Oxford: Oxford University Press.
- Pietroski, P. M., Lidz, J., Hunter, T., & Halberda, J. 2009. "The meaning of 'most': Semantics, numerosity and psychology." *Mind & Language* 24 (5): 554–585.
- Plantinga, A. 2003. *Essays in the Metaphysics of Modality*. Oxford: Oxford University Press.
- Quine, W. 1960. *Word and Object*. Cambridge, Mass.: MIT Press.
- Quine, W. 1973. *The Roots of Reference*. La Salle: Open Court.
- Rolander, N. 2013. "Swedish city of Kiruna plans massive relocation." *Wall Street Journal*. <http://on.wsj.com/1w3GxiL>
- Rothschild, D., & Segal, G. 2009. "Indexical predicates." *Mind & Language* 24 (4): 467–493.
- Russell, B. 1957. "Mr. Strawson on referring." *Mind* 66 (263): 385–389.
- Schein, B. 2002. "Events and the semantic content of thematic relations." In G. Preyer, & G. Peter (Eds.) *Logical Form and Language*. Oxford: Oxford University Press: 263–344.
- Scholl, B. J., & Nakayama, K. 2002. "Causal capture: Contextual effects on the perception of collision events." *Psychological Science* 13 (6): 493–498.
- Segal, G. 2000. *A Slim Book on Narrow Content*. Cambridge, Mass.: Bradford Books, MIT Press.
- Sider, T. 2002. *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford: Clarendon Press.
- Sider, T. 2009. "Ontological realism." In D. Chalmers, D. Manley, & R. Wasserman (Eds.) *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.

- Sider, T. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Spelke, E. S. 2003. "What makes us smart? Core knowledge and natural language." In *Language in Mind: Advances in the Study of Language and Thought*. Cambridge, Mass.: MIT Press: 277–311.
- Stainton, R. 2006. "Meaning and reference: Chomskian themes." In E. LePore, & B. Smith (Eds.) *The Oxford Handbook of Philosophy of Language*. Oxford: Oxford University Press: 913–940.
- Stanley, J. 2000. "Context and logical form." *Linguistics and Philosophy* 23 (4): 391–434.
- Stanley, J. 2007. *Language in Context: Selected Essays*. Oxford: Oxford University Press.
- Stotz, K., & Griffiths, P. 2004. "Genes: Philosophical analyses put to the test." *History and Philosophy of the Life Sciences* 26 (1): 5.
- Stotz, K., Griffiths, P. E., & Knight, R. 2004. "How biologists conceptualize genes: An empirical study." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 35 (4): 647–673.
- Strawson, P. F. 1950. "On referring." *Mind* 59 (235): 320–344.
- Thomasson, A. L. 1999. *Fiction and Metaphysics*. Cambridge: Cambridge University Press.
- Vogel, C. A. under review. "Realism and lexical flexibility".
- Vogel, C. A., Wellwood, A., Dudley, R., & Ritchie, J. B. 2014. "Talking about causing events." *Baltic International Yearbook of Cognition, Logic and Communication* 9.
- Weber, M. 2005. *Philosophy of Experimental Biology*. Cambridge: Cambridge University Press.
- Williams, B. 1979. "Internal and external reasons." In R. Harrison (Ed.) *Rational Action*. Cambridge: Cambridge University Press: 101–113.
- Wilson, R. A., Barker, M. J., & Brigandt, I. 2007. "When traditional essentialism fails: Biological natural kinds." *Philosophical Topics* 35 (1): 189.
- Wittgenstein, L. 1953. *Philosophical Investigations*. MacMillan.
- Wittgenstein, L. 1972. *On Certainty*. Harper Perennial.