LINGUISTIC INSTRUMENTALISM

John A. Keller and Lorraine Juliano Keller

To appear in R. Gaskin (ed.) *Linguistic Idealism*, OUP, 2025

ABSTRACT: Unless we are willing to countenance the skeptical possibility that most of our thought and talk is meaningless, it would seem that something must *guarantee* that the ontology presupposed by semantics aligns with the truth about ontology. One motivation for Linguistic Idealism is that it provides such a guarantee. We argue here that such a guarantee is not actually needed to avoid skepticism, since semantics can successfully 'do its job' of systematically pairing meanings with utterances even if it has false ontological presuppositions. More broadly, we explore and defend a form of anti-realism about linguistics according to which correct or acceptable linguistic *theories* can have false presuppositions and thus fail to be true. We argue that this is partly due to the fact that linguistic 'laws' or principles themselves—the things that govern or explain language and linguistic competence—do not need to be true in order to 'do their job'.

KEYWORDS: Realism, anti-realism, linguistics, linguistic idealism, semantics, ontology.

1. INTRODUCTION

It is widely taken to be a truism that in order for the meaning of an expression to be *x*, *x* must exist.[1] Call this *the Constraint*. Given the Constraint's neutrality about what meanings *are*—concrete objects, properties, relations, intensions, ideas, etc.—it seems difficult to deny. The Constraint merely insists that that the set of meanings is a subset of the domain of (existing) things. Platitudinous as it might sound, the Constraint has disquieting implications. If we suppose that meaningful expressions must have meanings,[2] the Constraint entails that the meaningfulness of our language requires semantics to be aligned with ontology. If the ontology presupposed by the semantics of our language turns out to be incorrect, our utterances will be rendered largely meaningless.[3] This isn't just an abstract possibility:

[1] This idea is ubiquitous. Searle (1969, 77) says 'Whatever is referred to must exist.' Williamson (2007, 20) defends the stronger claim that 'What there is determines what there is for us to mean.' And Szabó's 'First Dogma about Meanings' is that 'An expression is meaningful if and only if there is a semantic value *m* . . . and there is a relation *R* . . . such that the expression bears *R* to *m*' (2013, 35).

[2] This is not uncontroversial: see Quine 1960 and Jubien 2001. The Constraint applies most straightforwardly to 'entity-assigning' meaning theories, but it applies to other semantic theories as well: e.g., it says that the meaning of an expression can only be this pattern of use or that inferential role if those patterns and roles exist.

[3] We assume here that meaningless expressions can't be constituents of meaningful sentences, but an alternative view is that sentences containing meaningless expressions have existent but incomplete meanings: 'gappy propositions' (Braun 1993). But even if gappy propositions would save most sentences from meaninglessness, they wouldn't save them from

there are many false ontologies and only one true one, and it is unclear what would guarantee that semantics tracks the true one. But the possibility that semantics presupposes a false ontology is vertigo-inducing: the meaningfulness of our thought and talk seems to hang by a thread, if it's even hanging at all.

1.1 Satisfying the Constraint. Call the problem of finding a guarantee that semantics and ontology are aligned—and thus that the Constraint is satisfied—*the Puzzle*. There are three main approaches to solving the Puzzle:

> *Meaning Magnetism*: Ontological reality determines linguistic ontology. (Sider 2011)
> *Linguistic Idealism*: Linguistic ontology determines ontological reality. (Gaskin 2021)
> *Supernatural Design*: Linguistic ontology and ontological reality align by design.[4]

One might think we have left out an appealing option: *Easy Ontology* views that trivialize alignment by making ontology maximally permissive (Meinong 1960, Thomasson 2015). Easy Ontology views come in different flavors. The more plausible versions don't make ontology *too* permissive, but for that reason don't solve the Puzzle: linguistic ontology contains a menagerie of exotic and arguably impossible objects that only appear in *maximally* permissive accounts of ontological reality. On the other hand, maximally permissive views solve the Puzzle only by postulating implausibly exotic and extravagant ontologies. But there's a deeper problem with this approach: 'first-order' Easy Ontology views are just specific metaphysical theories, no more likely to be correct than any other. The fact that a possible metaphysical theory would satisfy the Constraint isn't any sort of *guarantee* that the Constraint is satisfied, and so doesn't solve the Puzzle. But 'second order' Easy Ontology principles that *guarantee* that the Constraint is satisfied must ultimately reduce to one of the other approaches: recherché counterexamples aside, if *X* and *Y* are non-accidentally correlated, then *X* depends on *Y*, *Y* depends on *X*, or *X* and *Y* depend on *Z*. So, it is to assessing those initial three approaches that we now turn. As we'll see, there is a real worry that these approaches generate more puzzles than they solve.

To begin with, Supernatural Design doesn't have much to recommend it: invoking *God* to solve the Puzzle would be a radical and unpopular move, and prominent supernaturalist positions contradict Supernatural Design anyway (recall the tower of Babel). Solving the Puzzle with a radical and unpopular version of a radical and unpopular view doesn't actually seem like much of a solution at all.

Meaning Magnetism is an interesting view, and there are influential arguments for it, but there remain deep concerns. On Ted Sider's development of the view, languages have a 'metaphysical semantics' that is quite different from their 'linguistic semantics'. These are not proposed as being two sides of the same coin. For example, linguistic semantics needs to be integrated with other psychological and linguistic theories (e.g., syntax), whereas metaphysical semantics doesn't (2011, 113). But that leaves the relationship between metaphysical semantics and linguistic semantics obscure, and the project of

truth-valuelessness, since gappy propositions aren't true or false. And the claim that most sentences might turn out to be truth-valueless isn't much less disquieting than the claim that they might turn out to be meaningless.

[4] Something like this view seems to have been held by the *Conimbricenses* (Ashworth 2015, 158–9; Doyle 2001, 140).

guaranteeing that they are appropriately aligned (such that the truth-conditions they assign determine the same set of worlds) seems not much different from the project of solving the Puzzle in the first place.

In addition, it is worth noting that solving the Puzzle would require a very strong form of Meaning Magnetism indeed.[5] If such a strong form of meaning magnetism is correct, it is all but guaranteed that the semantic theories we actually develop will be mistaken. According to Meaning Magnetism, the *true* ontology determines the ontology of language. But our semantic theories are unlikely to presuppose the true ontology unless we know the true ontology, and that seems rather doubtful given the history of ontological theorizing. In any case, for reasons that will become apparent, the falsity of our best semantic theories is grist to our mill, insofar as it suggests that an anti-realist attitude towards semantics is appropriate.

The last option—Linguistic Idealism—is, as Richard Gaskin himself puts it, 'quite a startling view' (this volume, xx). It would require a radical revision of our conception of the world and our place in it. As Gaskin (this volume, xx) says, 'We humans are only a tiny and insignificant bit of reality, and yet it is a condition on there being such a thing as reality in the first place—a reality for us to have evolved in—that it make sense to us. There is a sense in which, notwithstanding Copernicus, planet Earth is the centre of the universe after all.'

Even if we manage to reconcile ourselves to its radical anthropocentrism, the ontological implications of Linguistic Idealism are unsettling. When it comes to abstracta, Gaskin (2021, 169) admits that, 'whatever can consistently exist, does exist'. We would say 'forced to admit', but Gaskin seems untroubled, 'For clarity's sake, let me say [that] I see no problem with countenancing the existence of various "silly" (as they have unjustly been called) abstract objects that are often thought to put the kibosh on any policy of ontological generosity: here I am thinking of such objects as, well, sakes and kiboshes' (172). This is a radical ontological view. Things get even worse if, as seems plausible, Gaskin will be forced to admit inconsistent or impossible objects into his ontology. After all, many expressions appear to denote impossible objects if they denote anything at all. Gaskin handles problematic cases of this type by distinguishing between 'genuine names' and 'descriptive names',[6] and arguing that what's denoted by descriptive names for *impossibilia* are not impossible and hence unreal concrete objects, but perfectly real (and hence possible) *concepts of impossibilia*: concepts that actually and hence

[5] One might think that Meaning Magnetism only comes into play to settle edge cases that only philosophers care about. But if the *prima facie* mismatch between semantics and ontology is large, Meaning Magnetism will have to point reference towards meanings that are very 'unintuitive': e.g., 'trees' refers to simples arranged treewise, or to ideas of trees, etc.

[6] It's not clear that this distinction can do the work it's being asked to, unless we stipulate that empty names must be descriptive names. Gaskin (2021, 105) says, 'when we discover that [a] name is empty we may decide that, really, it is and was all along a descriptive name.' But there's no 'internal' linguistic difference between empty and non-empty names. (Linguistic analysis doesn't reveal that 'Vulcan' is empty.) So, if empty names must be descriptive names, the world is intruding on semantics in a significant way—a form of Meaning Magnetism—rather the opposite order of influence than the one Gaskin wants! Gaskin is of course alive to this worry (2021, 105–10), but his response, in short, is to accept the implication but draw a distinction: the sense in which language determines reality is *transcendental*, whereas the sense in which reality determines (semantic facts about) language is merely *empirical*. We leave an evaluation of this reply to the reader.

possibly exist, but are not possibly instantiated. Worries about such 'conceptualist' approaches to semantics aside (Kripke 1980), it is unclear that this is a fully general solution, since Russell's Paradox shows that some (purported) concepts, such as the concept of non-self-instantiation, cannot possibly exist. If such a concept existed, would it instantiate itself? Either way, contradiction looms.[7] And of course, there's no guarantee that the true ontology will end up containing concepts at all. (See Machery 2009 for an argument that it won't.)

1.2 Rejecting the Constraint. These brief remarks are not intended to be a refutation of Linguistic Idealism—far from it. *Every* attempt to solve the Puzzle has radical implications.[8] But if the only options for ensuring the satisfaction of the Constraint are radical, perhaps we might consider rejecting it—also a radical move, we admit, but perhaps the least radical of the bunch. In this paper we explore the prospects for doing just that: for rejecting the Constraint and holding that claims like 'the meaning of *x* is *y*' can be *correct* even if they are not true, and in particular even if there is no such thing as *y*.

More generally, we want to explore a form of anti-realism about linguistics according to which linguistic theories and principles can be correct without being true. Call this view *Linguistic Instrumentalism*. Van Fraassen (1980, 12) defends a form of *global* scientific anti-realism, according to which 'science aims to give us theories which are empirically adequate; and acceptance of a theory involves a belief only that it is empirically adequate.' We don't think *that's* true: we believe plenty of our scientific theories, and scientists do too. We are, in fact, realists about much of science and pretty much all of mathematics.[9] But we are only realists about *much* of science: *some* scientific theories (the ideal gas law is a notorious example) are clearly only aiming to be instrumentally useful, and accepting them involves a belief only that they are instrumentally useful. (For now, think of being instrumentally useful as a matter of making true predictions.) So, one can heartily reject *global* scientific anti-realism while accepting *restricted* or *local* forms of anti-realism about selected scientific disciplines or theories. Linguistic Instrumentalism is such a restricted form of anti-realism: '*linguistic* science aims to give us theories which are instrumentally useful; and acceptance of a *linguistic* theory involves a belief only that it is instrumentally useful.'

1.3 The Plan. This paper is exploratory. Much has been written on the subject of global scientific realism and anti-realism, and some has been written on realism and anti-realism concerning specific aspects of linguistic theory,[10] but little has been written on the subject of realism and anti-realism about

[7] See Russell 1980, and Gaskin 2021, 35–8, for a response.

[8] Anscombe (ignoring Supernatural Design) says, 'So now it looks as if either the grammar corresponded to something of the object, its real essence, which it has whether there is language or not, or the "object" were itself dependent on language. The first is like the suggestion made by Plato in the *Cratylus* [that the "logical shape" of words derives from the essence it expresses—akin to Meaning Magnetism]; the second, if it applies through and through, I call "linguistic idealism"' (1981, 113). She goes on to say, 'It is enormously difficult to steer in the narrow channel here: to avoid the falsehoods of idealism and the stupidities of empiricist realism' (1981, 115).

[9] We're also realists about morality, metaphysics, and logic. We're realists about pretty much anything it is reasonable to be a realist about. It was thus shocking to discover that we found anti-realism about linguistics plausible. But we argue below that linguistics is genuinely more amenable to anti-realism than other subjects.

[10] E.g., Rosen 1990 explicitly advocates anti-realism about possible worlds, Ludwig 2002 explicitly advocates anti-realism about T-sentences, and Ball 2018, Yalcin 2018, and Schwarz 2018 defend anti-realist-esque claims about semantics.

linguistics in general, but *only* linguistics. Our aim here is to clarify Linguistic Instrumentalism, canvas some of the 'big picture' reasons supporting it, and to knit together some of the arguments for anti-realism about isolated parts of linguistic theory into a more general argument for Linguistic Instrumentalism. We are not yet convinced of the truth of Linguistic Instrumentalism, but we *are* convinced that it merits further consideration.

2. CLARIFICATIONS

The view we are exploring is complicated and confusing. The way the terms 'realism' and 'anti-realism' are used is complicated and confusing. The different ways linguists conceive of linguistics makes talking about linguistics complicated and confusing. Clarifications are in order.

2.1 Linguistics. Linguistics is standardly defined as the scientific study of language, and disagreements about the nature of linguistics can be understood as disagreements about whether languages (*qua* the proper target of linguistic theorizing) are physical, psychological, or abstract (see, e.g., Katz 1985a, 11; Pitt 2018, 9; Ball and Rabern 2018, 33–4). *Nominalist* approaches (e.g., Bloomfield 1936) take linguistics to be the study of fundamentally *physical* entities such as sound waves and inscriptions. *Conceptualist* approaches (e.g., Chomsky 1965) take linguistics to be the study of *psychological* entities such as concepts and psychologically realized rules or grammars. Finally, *Platonist* approaches (e.g., Katz 1985b) take linguistics to be the study of *abstract* entities such as sentence-types and abstract rules or grammars.

We favor an ecumenical approach: by our lights, the study of language involves the study of physical, psychological, *and* abstract entities, and how they relate to one another (compare Santana 2016). This ecumenical perspective is clearly *descriptively* correct, in the sense that the work done in linguistics departments includes the study of the physical, psychological, and abstract aspects of language.

Many linguists favor a less ecumenical *prescriptive* approach according to which much of what goes on in linguistics departments is not *proper* linguistics. For example, Chomskyans think that linguistics 'should aim to express just what is "essential" to natural language' (Chomsky 1966, 537). We don't want to debate how to define or delineate the boundaries of linguistics, and doing so is not required for our purposes. We think that anti-realism is true of 'Chomskyan linguistics' and (most or all) other 'narrow' conceptions of linguistics as well as linguistics broadly conceived. But the specific way anti-realism is manifested will depend on one's account of linguistics. We discuss *some* such specifics below; here we simply assume that linguistics includes everything done in (most top) linguistics departments.

2.2 Correctness. We introduced Linguistic Instrumentalism as the view that linguistic principles can be correct without being true. But what do we mean by 'correct'? The concept is best grasped by example. Consider a mathematical anti-realist who thinks all mathematical claims are untrue. Most such antirealists will admit that some untrue mathematical claims are better than others. There is an important sense in which it is *better* to endorse '2 + 2 = 4' than to endorse '2 + 2 = 14'. There is *something* right about '2 + 2 = 4', and pretty much *nothing* right about '2 + 2 = 14'. Intelligent and

informed people say things like '2 + 2 = 4', but not '2 + 2 = 14'. And so forth. Even if untrue, '2 + 2 = 4' is still 'good' in an important sense.

Similarly, if you are an anti-realist about color, and think all color attributions are untrue, you will likely admit that there is still an important sense in which 'apples are red' is *better*, or more right, than 'apples are blue'. If you are a moral anti-realist and think all moral attributions are untrue, there is still an important sense in which 'murder is wrong' is better, or more right, than 'murder is obligatory'. Etc.

Everybody needs a way of sorting 'good' claims from 'bad' ones. Realists sort by (objective) truth; anti-realists sort by some other method. We are using 'correct' as a label for the category of 'good' claims. Different views yield different accounts of correctness. For realists, correctness is (objective) truth. For van Fraassen, it is empirical adequacy. For instrumentalists, it is usefulness. For cultural relativists, it is popularity. As the name suggests, Linguistic Instrumentalism is the view that correct linguistic claims are those that are *useful*. (We explain this notion of usefulness below.)

2.3 Correct Rules? But does it make sense to talk of linguistic principles being true? Such principles are often (presented as) *rules*, and rules are plausibly not truth-apt. The claim that correct rules needn't be true would then be trivial. But even if rules aren't truth-apt, they may have truth-apt *commitments* (presuppositions, entailments, etc.). For example, Reiland (2023) argues that while rules are not truth-apt, they have contents that are. More generally, all manner of non-truth-apt expressions (questions, imperatives, etc.) have truth-apt commitments. For example, regardless of whether the rule 'Do what Zeus commands' is truth-apt, it has a (false) truth-apt presupposition: that Zeus exists.

To avoid repeating the clunky phrase 'not true, or has non-true commitments', we will use 'not true' (and 'untrue') in an extended sense to cover both truth-apt things that are not true ('Snow is green'), and non-truth-apt things that have non-true but truth-apt commitments ('Do what Zeus commands').

2.4 Realism and Anti-realism. Linguistic Instrumentalism is a form of anti-realism about linguistics. On standard accounts, realism about a domain *d* requires:

> *Objectivity*: Facts in *d* are mind-independent: they do not depend in any non-trivial way on human mental activity (beliefs, desires, hopes, etc.).
>
> *Truth-aptness*: Discourse in *d* is truth-apt: it aims at truth and can be evaluated for truth and falsity (unlike, say, discourse in fictional or theatrical domains). In short, correctness = truth.[11]

As noted above, anti-realist views about a domain *d* correspond with views about what (other than objective truth) makes theories in *d* correct. For example, van Fraassen's Constructive Empiricism holds that science only aims to produce theories that are true of *observable* reality, and so correct

[11] Similar accounts of realism are common: see, e.g., J. A. Keller 2014 and Boyd 1983. Such accounts typically include a third condition, *Transparency*, according to which discourse in *d* has a transparent interpretation. Consider someone who claimed that moral statements expressed mathematical truths, and hence were objective and truth-apt. We wouldn't call such a person a moral realist, and we need something like Transparency to rule that out. But for reasons outlined in Keller 2014, 15–17, the requirement is difficult to formulate precisely.

scientific theories can fail to be true *tout court* (since they may be false of *un*observable reality). Such theories—or at least theorists' attitudes towards them—are thus not truth-apt.

2.5 Linguistics is Different. To see why Linguistic Instrumentalism might be plausible even to those who, like ourselves, are realists in other domains, consider how much less objective linguistics is than (most) other sciences. If competent speakers believe that '*x*' is grammatical and means *y* in *L*, then '*x*' *is* grammatical and *does* mean *y* in *L*. The grammaticality and meaning of expressions just aren't things competent speakers could be wrong about. Likewise with pronunciation, etc. This is as mind-dependent as things get. The contrast with sciences like physics and biology is striking.

While it is uncontroversial that relatively superficial linguistic facts such as that '*x*' is grammatical and means *y* in *L* are mind-dependent, we nonetheless think it is useful to draw attention to this point in the context of comparing linguistics with the other sciences vis-à-vis realism. But the deeper and more interesting reason to endorse Linguistic Instrumentalism derives from the *non-truth-aptness* of many fundamental linguistic theories and principles.

2.6 Laws, Principles, and Theories. We said above that we disagreed with Constructive Empiricism because we think science often *does* aim to produce theories that are true *tout court*. (Not always. Perhaps not usually. But often.) Nevertheless, we compared Linguistic Instrumentalism to a *restricted* version of Constructive Empiricism. But that was actually just a useful simplification: Linguistic Instrumentalism is something much *stranger*.

Terms like 'law', 'principle', and 'theory' are used by linguists in a systematically ambiguous way to refer to both what is being theorized about as well as our theories about it. For example, Chomsky (1965, 25) uses 'the term "theory"…with a systematic ambiguity, to refer both to the child's innate predisposition to learn a language of a certain type and to the linguist's account of this'. Using terms in a systematically ambiguous way seems liable to cause confusion, so we won't. There is a standard distinction between *laws* and law *statements*: roughly, laws govern and explain the behavior of things, and (correct) law statements state or express laws in verbal form. But talk of laws of linguistics is strained, so we'll use the term 'principles' to refer to 'linguistic laws'—the things that govern and explain linguistic behavior (including the operations of the language faculty)—and 'theories' to refer to statements of those principles. This distinction is important, since saying that linguistic *theories* can be correct without being true is very different from saying that linguistic *principles* can be correct without being true. The idea that linguistic *theories* can be correct without being true is a relatively familiar kind of *wordy* anti-realism, *à la* Constructive Empiricism, while the idea that linguistic *principles* can be correct without being true is a relatively unfamiliar kind of *worldly* anti-realism.

2.7 Worldly Anti-realism. But what, exactly, *is* worldly anti-realism? Realism was glossed in terms of objectivity and truth-aptness. Objectivity straightforwardly applies to both principles and theories: 'fact' is another one of those systematically ambiguous words that has both a wordy and a worldly meaning. But the above account of 'truth-aptness' only applies to theories. Still, the core idea behind

truth-aptness—that correctness is truth—can be generalized to cover principles. Worldly anti-realism only makes sense in cases where what is being theorized about—the *target* of our theorizing—is something that can be untrue: something truth-apt, or with truth-apt presuppositions. Principles fit this description. Accordingly, the generalization of truth-aptness to linguistic principles is the idea that linguistic principles—the principles *actually* governing and explaining linguistic phenomena—must be true. One might, however, wonder how the principles actually governing and explaining linguistic phenomena could *fail* to be true. After all, principles are linguistic *laws*, and it's generally taken to be *analytic* that laws are true.

The idea that laws must be true has been famously challenged by Nancy Cartwright. On her view (Cartwright 1983), *the laws of nature themselves lie*, so that correct, fully accurate, statements of those laws may lie as well. Here's a simple illustration. Assume that *that* $F = ma$ (force equals mass times acceleration) is an *actual* law of physics: a principle governing and explaining the behavior of physical objects. On a Cartwrightian view,[12] this could be the case even if it isn't *true* that force equals mass times acceleration. So, *that* $F = ma$ could be a false law. The law *statement* '$F = ma$' would also be false, since it's false that $F = ma$. But '$F = ma$' would be a fully accurate—not just empirically adequate—statement of an actual law. Theorists might be aiming to *literally* state a law, and *succeeding*, but nonetheless say something (correct but) untrue. Hence it is not an objection to this view—unlike many forms of anti-realism—that it doesn't *seem* to theorists that they are engaging in pretense or some other non-truth-apt form of discourse. Someone who asserts '$F = ma$' succeeds in literally and accurately describing a law, just as someone who asserts 'Bilbo found a magic ring' succeeds in literally and accurately describing *The Hobbit*. What they assert is nonetheless untrue. The problem, in both cases, is that what they are literally and accurately describing is not itself true.

This view rejects truth-aptness (for both theories and principles) and is thus anti-realist, but the fundamental source of the anti-realism isn't the theory (law statement), or the intentions of the theorizer, but the world (the law or principle itself). Indeed, the deviation from truth-aptness in the case of *theories* can easily be remedied: even if '$F = ma$' is false, 'it's a law that $F = ma$' is true. (Likewise with 'In *The Hobbit*, Bilbo found a magic ring'.) Call law statements like '$F = ma$' *unadorned* and law statements like 'it's a law that $F = ma$' *adorned*. On this view, whether correct law *statements* (theories) are true or false will depend on how they are formulated: on whether they are adorned or unadorned. Hence (in part) the focus on the claim that the laws themselves are false.

We don't agree with Cartwright's view of science any more than we agree with van Fraassen's: it's hard for us to wrap our heads around the idea that the actual laws of physics, for example—the principles actually governing and explaining physical phenomena—are false. But while we reject Cartwright's view as an account of science *in general*, we think it is surprisingly plausible as an account of linguistics: the linguistic principles that actually explain the linguistic facts and govern linguistic behavior plausibly have false presuppositions, and hence do not accurately describe the phenomena they explain. They are, in short, correct but untrue. As Cartwright might put it, *the laws of linguistics lie*.

[12] Cartwright says things along these lines, but her view is complicated. The aspect of her view that we're focusing on is similar to the one emphasized in Russell (forthcoming).

2.8 Linguistic Correctness. But if the fundamental theories and principles of linguistics don't aim at truth, what *do* they aim at? Instrumentalism replaces questions about whether a theory or principle is *true* with questions about whether it is *useful*. But to know whether linguistic theories or principles are useful, we have to know what they are used *for*. What is the use or purpose of language (and linguistic *principles*), and what is the use or purpose of linguistic *theories*? The details are matters of dispute among linguists, but there is *some* agreement about the broad outlines. Human beings have thoughts and desires they would like to communicate.[13] Language facilitates this. Call the contents or meanings of complete thoughts and sentences 'meanings$_S$'. If I am thinking a thought with content *p* , I can communicate that to you by producing sounds, signs, or marks ('utterances') that you interpret to mean *p*. For this to work in a general and efficient way, there must be mutually recognized systematic *principles* for associating utterances with meanings$_S$ (in context). Linguistic *theories* are attempts to state those principles. But, and this is the crucial point, those principles do not need to be *true* in order to fulfill their function: in order to enable us to efficiently and systematically associate meanings$_S$ with utterances.

This brings us to an important clarification. We are happy to say that utterances have meanings$_S$, and that it is the job of linguistics to systematically pair utterances with meanings$_S$. But we think the linguistic machinery used for systematically pairing utterances with meanings can do its job even if it presupposes a false ontology and is thus untrue. Hence, we are skeptical about sub-sentential meanings, syntactic structures, and all the rest of it. We are realists about utterances and meanings$_S$, but not about what comes in between.

If we think of the language faculty as *computing* the meanings$_S$ of utterances in context, what matters is only that the computational principles they use *work*—whether they deliver the right results—not whether they're true. And computational principles *can* work while being untrue. Using the heights of shadows to compute the heights of objects obviously don't require thinking that shadows actually exist any more than computing the number of deadly sins by putting them in 1–1 correspondence with the seven dwarves requires thinking that there are dwarves.

Alternatively, we might think of linguistic principles as being used by hearers to *predict* what speakers are thinking given the utterances they produce, and by speakers to *predict* what utterances will enable them to be understood. It is uncontroversial that false principles can reliably make true predictions, and so if linguistic principles are in the business of prediction, that would make anti-realism about them more plausible than anti-realism in other more descriptive domains. Regardless of whether linguistic principles are used for computation or prediction, the fact that they have a practical *use*—unlike, say, the laws of physics—makes anti-realism about them more plausible than anti-realism about physics. Other sorts of scientific *theories* clearly have uses (such as prediction and explanation), and one

[13] While this is widely accepted, Chomsky (2017, 87–8) (in)famously dismisses the claim that communication is a central function of language, holding instead that language is fundamentally an instrument of thought.

of the main motivations for *wordy* anti-realism about those theories derives from that fact. But here we are talking about laws (principles), not law statements (theories), and it's unclear what it would mean to say that the laws of physics have a practical use.

Another alternative would be to say that linguistic principles are *instructions* for pairing meanings$_S$ with utterances (cf. Pietroski 2018). But instructions with false presuppositions can do their job, whether they're instructions for computing a solution, getting from *A* to *B*, writing a college essay, or assembling a desk. This is, we think, obvious upon reflection: outside of the philosophy room, non-literal speech is *de rigueur* (with certain exceptions like computer programming). I can successfully explain how to fly from Rome to Reggio Calabria by saying, 'Fly south along the coast until you get to the tip of the boot'; instructions for how to build a desk might make reference to such ontologically suspect entities as 'left sides', 'right sides', 'holes', etc. If instructions had to be literally true in order to be effective, nothing would get done! Or at least, nothing would get done if instructions weren't written by lawyers or philosophers. That's clearly false.

The claim that untrue linguistic principles can 'do their job' is the strangest part of Linguistic Instrumentalism, so it's worth dwelling on the point. Consider an analogy. In blindfold chess, the participants play the game 'in their heads' (without a chessboard or pieces). But it remains *correct* to talk of pieces having a location on the board, players moving their pieces on the board, and capturing, gaining, and losing those pieces by saying things like 'B8 to A6'.[14] That's a sort of wordy anti-realism. But the rules governing, or instructions for playing, blindfold chess are just the rules or instructions for chess: they mention pieces, the board, etc., even though none of those things are actually there. And that's perfectly fine! Playing blindfold chess is, in fact, a way of playing chess. The standard chess rules *correctly* apply to blindfold chess, even though they don't *literally* apply, even though the state descriptions to which they appeal (e.g., white's king being in such-and-such position) don't actually *obtain*. The point of the rules is to enable play, and the rules do that regardless of whether they literally describe the play they enable. That's a sort of worldly anti-realism.

The chess example illustrates an important point. Anti-realism is not a blank check. Van Fraassen says that acceptance of a theory *t* involves a belief only that *t* is empirically adequate. If so, the belief that *t* is empirically adequate is nonetheless truth-apt, and so is correct only if it is true. On van Fraassen's view, *t* is only correct if '*t* is empirically adequate' is true. On instrumentalist views, *t* is only correct if '*t* is useful' is true. *Every* claim about boards and pieces in a game of blindfold chess is untrue, and *many* are also incorrect. So correctness isn't trivial to come by: *something* must be true for, e.g., 'White has three pieces on the board' to be correct in a game of blindfold chess. It just isn't that White has three pieces on the board.

3. STRUCTURES

[14] In standard notation, move the piece on B8—Black's knight at the start of the game—to A6.

According to the conception of linguistics outlined in §2.8, what makes the core linguistic subdisciplines (semantics, syntax, morphology, and phonology) *core* subdisciplines (unlike, say, neurolinguistics or sociolinguistics) is the fact that they are essential for the systematic association of utterances with meanings$_S$. (Recall that we are using 'utterances' to refer to physical sequences of sounds, movements, and marks.) Syntax is essential because it determines which utterances are 'in' the language and thus eligible to be assigned meanings$_S$ in the first place (ungrammatical 'utterances' like 'Dog cat chased' don't have meanings$_S$), and because the syntactic structures it assigns to utterances are required for semantic processing. Morphology and phonology are needed—in conjunction with syntax—for identifying the meaningful constituents of utterances. That is essential because the *identities* and *meanings*$_S$ of utterances are functions of their syntax and the identities and meanings of their meaningful parts.[15]

We allow that *some* linguistic subdisciplines, such as corpus linguistics, aim at truth. Linguistic Instrumentalism is a view about the *core* linguistic subdisciplines: semantics, syntax, morphology, and phonology. We canvassed some of the 'big picture' reasons supporting Linguistic Instrumentalism in §2; here we focus on fine-grained considerations supporting Linguistic Instrumentalism in phonology, morphology, and syntax. (We reserve discussion of semantics for §4.) Call the things to which we attribute phonological, morphological, and syntactic properties/structures *sentences*. In broad and sketchy outline, the case for instrumentalism about phonology, morphology, and syntax is given by *The Argument*:

1. The actual morphological, phonetic, and syntactic *principles* governing the language faculty attribute morphological, phonetic, and syntactic *properties* or *structures* to sentences. So do (typical/unadorned formulations of) linguistic *theories* that state or describe those principles.
2. Sentences do not have the morphological, phonetic, or syntactic properties or structures attributed to them by the abovementioned principles (and theories).

So, 3. Morphological, phonetic, and syntactic principles and theories are untrue (because not truth-apt).

The Argument seems valid: if the actual *principles* governing language and underlying linguistic competence are (correct but) untrue, that would establish *worldly* linguistic anti-realism. It would follow that unadorned *statements* of those principles are (correct but) untrue. That would establish *wordy* linguistic anti-realism. But is The Argument sound? To support its premises, we must show that morphological, phonetic, and syntactic principles attribute properties to sentences that they do not have. Our initial arguments assume that sentences are utterances—physical sounds, movements, or marks—but we argue in §3.5 that a parallel argument goes through on Conceptualist and Platonist assumptions.

[15] While morphology, phonology, and syntax are distinct sub-disciplines, they cannot be pursued in total isolation: how a sentence should be divided up into meaningful parts depends on morphological, phonetic, and syntactic considerations.

3.1 Phonology. We begin with phonology.[16] Phonologists and phoneticians have long recognized that the speech stream does not present hearers with a set of discrete, articulated sounds, let alone words; instead, hearers perceive features in the acoustic stream that are not really there. One example is the phenomenon of categorical perception. Experiments have confirmed that humans, chinchillas, monkeys, and rats perceive nonexistent sharp differences in streams of sound that differ along a continuum: hearers perceive discrete differences between phones (sounds), even in cases where a computer manipulates a phone so that it continuously changes from, e.g., /ba/ to /da/ (Adger 2019, ch. 6). Hearers also perceive non-existent spaces between words, even in cases of languages with no writing systems and languages like Mandarin whose writing systems do not use discrete units for phonemes (sets of phones that make morphemic or lexical distinctions within a language). Another example is the well-documented phenomenon of phonetic illusion, which Edward Sapir studied in speakers of various languages, including Paiute, Sarcee, and English. Sapir's English-speaking phonetics students transcribed non-existent glottal stops, influenced by internalized rules about the pronunciation of American English (Zsiga 2020, 13). A further example of phonetic misrepresentation is provided by experiments showing that Japanese speakers, influenced by rules forbidding consonant clusters, report hearing non-existent vowels in the midst of consonant clusters: e.g. they hear [ebzo] as [ebuzo] (Dupoux et al. 1999, 14). Finally, it is well known that transcriptions of phones into the International Phonetic Alphabet vary considerably.

The upshot is that the way hearers represent the speech stream differs significantly from how it actually is: the acoustic signals produced by speakers do not have the phonetic forms attributed to them by hearers. It's *possible* to posit something outside of the minds of language-users that has the phonetic features attributed by the language faculty, but it's not plausible. As Chomsky says,

> Communication could then be described in terms of such (partially) shared entities, which are easy enough to construct: take '*a' to be the singleton set {a}, or {3, a}; or if one wants a more realistic feel, some construct based on motions of molecules. With sufficient heroism, one could defend such a view, though no one does, because it's clear that we are just spinning wheels. (2000, 129)

The evidence thus strongly suggests that *nothing* outside of the minds of language-users has the phonetic structure they impute to the sound stream.[17] Phonetic structure is being *projected onto* the acoustic stream by hearers. While the acoustic stream is experienced and processed *as if* it were segmented into discrete phones, the reality is that it is continuous and unsegmented. As phonologist Elizabeth Zsiga (2020, 29) writes, 'if segments didn't exist, we would have to invent them.' Indeed.

[16] Phonology, the study of sound patterns in a language, is standardly taken to deal with speech sounds as discrete, psychological entities. Phonetics, the study of how speech sounds are produced, is concerned with speech sounds as continuous, physical entities. The distinction between phonology and phonetics has been (and, to some extent, still is) fraught, with some historically prominent linguists (e.g. Saussure, Trubetzkoy, and Hjelmslev) arguing that phonetics is not part of linguistics proper (Zsiga 2020, ch. 2).

[17] For simplicity, we've focused on spoken language in this passage, but the same points apply to signed and tactile languages, *mutatis mutandis*.

The anti-realism here is fundamentally worldly and only derivatively wordy. It is the language faculty itself, in its interface with the sensorimotor system, that is misrepresenting the acoustic stream. But despite this misrepresentation, or perhaps even because of it, the language faculty is able to do its job effectively. The *actual* principles governing the language faculty are *eo ipso* correct, regardless of whether they mispresent the sound stream. Of course, whether this worldly anti-realism percolates up into our linguistic theories will depend on whether the 'law statements' in those theories are adorned or unadorned. Either way, the worldly anti-realism remains.

3.2 Morphology. Morphology is the study of the forms of words.[18] Lexemes, the fundamental units of the lexicon, are abstract representations that include all inflected forms of a word: for example, there is one lexeme associated with the forms 'sing', 'singing', and 'sung'. While it's true that language-users *represent* distinctions between phonemes, morphemes, and lexemes, in the sense that such distinctions affect how they produce, respond to, and understand spoken, written, or signed speech, it's doubtful that anything outside of their minds has the features language-users are representing—i.e., the syntactic, semantic, and functional properties by which (e.g.) morphemes are distinguished. Language-users are responding to sounds, signs, and inscriptions (we assume), but physical tokens do not have the relevant semantic and syntactic properties. For example, consider a scrap of paper with 'IS' inscribed on it. Is the inscription the numeral '15'? Or the English word 'IS'? Or perhaps the paper is upside down and it's the Spanish word 'SI'. Or—if the inscription bled through and we're looking at the wrong side—could it be the approximately-equal sign '≃'? The properties of the inscription itself do not settle the matter.

And of course, spoken morphemes are ultimately composed of phonemes, which are sets of phones. We saw above that phonetic phenomena themselves are illusions of a sort—they are projected onto the acoustic stream by hearers. The acoustic stream itself is not actually divided into phones. But phones are the ultimate building blocks of (spoken) morphemes. The result is that the structures they compose (morphemes, etc.) 'would seem to be enormous and intricate buildings resting on—nothing. Or, if there turned out to be something, this would be entirely accidental' (Rey 2006, 250). The principles governing the language faculty attribute properties to parts of utterances, but those parts don't actually have those properties and may not actually exist. But, as before, the principles governing the language faculty are *eo ipso* correct, and likewise with unadorned statements of them. They are just not true.

3.2.1 Wordy Morphological Anti-realism. The fact that linguists generally assume that expressions have their meanings and other linguistic properties *necessarily* provides strong support for anti-realism about current linguistic practice: this is an idealized assumption that is clearly false of expressions in actual human languages. We don't deny that this idealization is theoretically useful, but it *is* an idealization: linguistics isn't (primarily) in the business of uncovering necessary truths. The meanings of natural language expressions—and of the sentences in which they appear—are at least largely contingent and

[18] Morphemes are the smallest linguistic constituents that have a meaning or function: e.g., the suffix '-m'—a morpheme that marks the accusative case in 'whom'—has a function but not a meaning.

empirical matters. This way of individuating expressions and indeed languages for the purpose of theorizing is strong evidence that that theorizing is not being undertaken in a wholly realist vein. While this is mere wordy anti-realism that doesn't seem to have a worldly correlate, we think it is worth pointing out since, as Yalcin (2014, 35) says, it is 'standard in syntax and semantics'.

It is worth stressing that there is a *lot* of merely wordy anti-realism in linguistics. According to Williamson (2018, 21), 'As we learn ever more of the extraordinary complexity underlying even the most ordinary conversations, philosophers of language and linguists may have to rely increasingly on a model-building methodology.'[19] Models proceed by way of *idealization*, and 'To say that it is good often to proceed by way of idealization is to argue that sometimes, in thinking about the world, truth isn't what you need . . .. [A]n idealization is a useful untruth' (Appiah 2017, xii), Linguistic theorizing by way of model-building is not truth-apt.

Such merely wordy anti-realism is not our focus because there is a fair amount of merely wordy anti-realism—including idealization and model-building[20]—in most scientific disciplines, and adjudicating whether there is proportionately more in linguistics would be difficult and not particularly interesting. Even more importantly, merely wordy anti-realism doesn't dissolve the Puzzle. The Puzzle isn't about the alignment of our current or even ultimate semantic *theories* with the true ontology, but about the alignment of the semantic *principles* actually governing natural languages with the true ontology.

3.3 Syntax. Finally, consider syntax, and in particular the sort of tree structure commonly used to represent the syntax of sentences like

(W) John seems to Bill to want to help himself.[21]

[19] This perspective is explicitly defended in Yalcin 2018 and only a bit less explicitly in Ball 2018.
[20] We know it is controversial to claim that modelling is an anti-realist enterprise: if you disagree feel free to substitute 'non-truth-apt' for 'anti-realist'.
[21] We borrow this example from Rey (2006, 245).

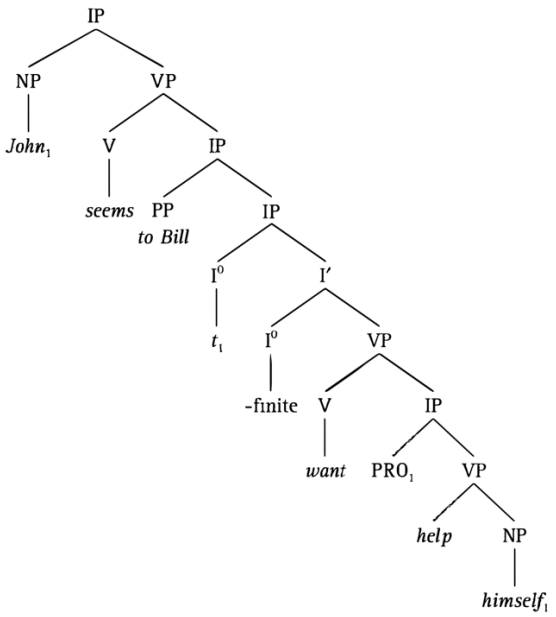

Call this structure '(S_W)'. One argument for syntactic instrumentalism is that (S_W) contains 'empty' categories—trace (t_1) and PRO_1—that, as Rey (2006, 245) puts it,

> indicate a node in the tree that for one reason or another has no [constituent] attached to it, for example, because [a constituent] has been 'moved' to a different node in some 'transformation' of the sentence, and/or needs to be co-indexed with elements at other nodes: thus t_1 is co-indexed with *John*, PRO, and with *himself*, capturing the fact that the subject of *want* and *help* are both *John*, and not *Bill*.

We have worries about the non-metaphorical meaning of 'moved' and 'transformed', but our fundamental concern is that utterances don't seem to instantiate anything like (S_W). Intuitively, (W) is supposed to instantiate the structure (S_W), but no actual physical utterances that would be used to express (W)—neither the string of marks that appears to the right of '(W)' above nor the sequence of sounds one would produce if one read it—instantiate (S_W). Indeed, given that (S_W) contains empty categories, it's unclear what it would mean for something to instantiate it. Only things with three sides can instantiate a triangular structure. Does something instantiating (S_W) need to have 14 elements? But since trace (t_1) and PRO_1 are empty, 14 seems like too many. But how could something with less than 14 elements instantiate (S_W)? That sounds like saying that something with two sides could instantiate a triangular structure.

But set that worry aside. The fact that utterances don't have the syntactic structures posited by our syntactic theories is entailed by the existence of structurally ambiguous utterances (strings that can be associated with more than one syntactic form). E.g.,

(K) I killed the king with a knife

can be associated with

(S_{K1}) $[_{TP}[_{NP}$I $[_{VP}$killed $[_{NP}$the king$][_{PP}$with $[_{NP}$a knife$]]]]]$ (I killed the king by using a knife)

or with

(S_{K2}) [$_{TP}$[$_{NP}$I [$_{VP}$killed [$_{NP}$[the king][$_{PP}$with[$_{NP}$a knife]]]]]] (I killed the king who was wielding a knife).

This sort of structural ambiguity is pervasive. Utterances thus do not have the syntactic structures attributed to them by the linguistic principles governing our language faculties and explaining linguistic competence. But, to repeat, those principles are *eo ipso* correct. They are thus correct but untrue. Likewise with unadorned statements of them.

3.4 Non-physical Sentences. Our argument so far has assumed that sentences are utterances: physical entities. But what if sentences are mental or abstract entities, as Conceptualists and Platonists hold? Let's grant for the sake of argument that such entities exist, in addition to physical utterances. Call physical sentences 'sentences$_P$', abstract sentences 'sentences$_A$', and mental sentences 'sentences$_M$'. Might sentences$_A$ or sentences$_M$ *be* or at least literally *have* syntactic structures like (S_W), as well as phonetic and morphological structure? Could such a view vindicate phonological, morphological, or syntactic realism?

Consider first the view that sentences are sentences$_A$. One obvious worry about that view is a version of Benacerraf's problem: there are too many candidates among abstract entities for playing the role of sentences and no principled way of choosing from among them (Benacerraf 1965).[22]

But there is a more fundamental problem with using sentences$_A$ to vindicate realism. Think about the poverty of the stimulus argument (Chomsky 1965, 30–1). A key premise of that argument is that the stimuli language learners encounter are compatible with a wide a variety of possible grammars (sets of linguistic principles). That's true if stimuli are sentences$_P$ (utterances) like (W), but much *less* true if they are sentences$_A$, or indeed *anything* that literally has a tree structure.[23] Call the tree *diagram* above '(D)'. If our stimuli included tree-diagrams like (D)—if, say, that's how English was written—then our (written) stimuli *wouldn't* be compatible with (such) a wide variety of grammars. (D) contains much more grammatical information than (W). So, anyone who thinks the poverty of the stimulus argument is onto something cannot think that our stimuli are things like (D). But in any case, we obviously *don't* communicate using tree diagrams like (D), or indeed anything that is or has a rich tree structure. Evidence for the obvious is provided by the lamentable but indisputable fact that there are, in reality, structurally ambiguous utterances like (K). One can insist that utterances like (K)—our structurally ambiguous stimuli—are not genuine sentences. The cost of this move is that it would make genuine sentences irrelevant to communication. As Rey argues, abstract objects like numbers earn their keep by being indispensable to physics, but the syntactic structures of sentences (like sentences themselves) "*have no role to play independently of our representations of them*" (2006, 250).

[22] See King 2007, 47; and L. J. Keller 2019.

[23] Of course, our stimuli would still be impoverished by not including (many) 'negative' (ungrammatical) examples. See Rey (2020, §1.3).

Rey is a Conceptualist, and seems to be suggesting that sentences$_P$ are just as irrelevant as sentences$_A$: that it's sentences$_M$ that matter for linguistics. But the same considerations regarding the poverty of the stimulus apply to sentences$_M$. *If* sentences$_M$ have rich syntactic tree structures, they cannot be among our stimuli (given that our stimuli are impoverished). Of course, sentences$_M$ plausibly have a significant role to play in language comprehension and production. They are certainly not linguistically irrelevant. But the idea that sentences$_M$—or sentences$_A$ for that matter—are among our stimuli is deeply implausible. Just as children cannot hear (or see) abstracta, nor can they hear (or see) the constituents of their interlocutors' minds. In any case, forget The Argument, and consider *The Argument**:

1*. The actual morphological, phonetic, and syntactic *principles* governing the language faculty attribute morphological, phonetic, and syntactic *properties* or *structures* to sentences$_P$. So do (standard formulations of) linguistic *theories* that state or describe those principles.

2*. Sentences$_P$ do not have the morphological, phonetic, or syntactic properties or structures attributed to them by the abovementioned principles (and theories).

So, 3. Morphological, phonetic, and syntactic principles and theories are untrue (because not truth-apt).

The soundness of The Argument*—assuming it is valid—only requires that phonetic, morphological, and syntactic structures are attributed to sentences$_P$, and that sentences$_P$ don't have them. Perhaps such structures are *also* attributed to sentences$_M$ and sentences$_A$, and perhaps they *do* have them. That would do nothing to undermine The Argument*.

It does not matter for The Argument* what sentences are. Let them be sentences$_A$, or sentences$_M$, or tree diagrams, or actual trees for that matter. The fact remains that linguistic principles guide us, or our language faculties, to attribute phonetic, morphological, and syntactic structures to sentences$_P$ like (W) and (K). In so doing, they are attributing structure that is not actually there. But of course, the linguistic principles governing our language faculties—the principles explaining linguistic competence—are *eo ipso* correct. They are thus correct but untrue. Likewise with unadorned statements of those principles. Hence, premises (1*) and (2*) of The Argument* are true, thus establishing Linguistic Instrumentalism with respect to morphology, phonetics, and syntax.

The claim that we represent utterances as having properties they do not have, and treat them as if they did, is a form of *projectionism*: we (or our language faculties) are *projecting* such properties onto utterances. Such 'projection' talk is not uncommon in linguistics and philosophy of linguistics.[24] Since projectionism is a form of anti-realism, this buttresses our claim that morphology, phonetics, and syntax involve a form of worldly anti-realism.

[24] E.g., in the course of defending the claim that sentences and words are mental entities, McGilvray (1998, n. 19) says, 'What about the offended complaint: "But we hear words issuing from another person's mouth, and surely other people's mouths are outside the head!" The same, however, is true of colours, and they too are inside the head. We need a projectivist account here.'

4. SEMANTICS

Let's reorient ourselves. What, exactly, do we need an instrumentalist semantics to do? According to the Constraint, the meaning of an expression can only be x if x exists. The Constraint makes meaningfulness beholden to ontology in a way that has seemingly skeptical implications. The Puzzle was to find a way to ward off skepticism by *guaranteeing* that the Constraint was satisfied. But since every *solution* to the Puzzle raised puzzles of its own, we've been exploring the idea that instead of solving the Puzzle, we should *dissolve* it: reject the Constraint and hold that the meaning of an expression *can* be x even if x fails to exist. More broadly, we have been exploring the prospects of Linguistic Instrumentalism, a broad form of anti-realism about linguistics.

Linguistic Instrumentalism allows that there are utterances and meanings$_S$, but is instrumentalist about everything in between.[25] Linguistic principles systematically pair utterances with meanings$_S$, such that given an utterance we can compute its meaning$_S$, and given an expressible meaning$_S$, we can compute an utterance that expresses it. Very broadly, Linguistic Instrumentalism is the view that the principles pairing utterances with meanings$_S$ can be successful—can reliably enable the systematic pairing of utterances with meanings$_S$—even if they are untrue.[26] This *helps* dissolve The Puzzle, since many proposed principles for pairing utterances and meanings$_S$ presuppose implausible ontologies. If the ontological presuppositions of linguistic principles had to be true for them to 'work', there would be a real worry that they *didn't* work, and hence that most of our thought and talk was meaningless. Linguistic Instrumentalism alleviates that worry by holding that linguistic principles can do their job of pairing utterances with meanings$_S$ even if they have false ontological presuppositions. It *insulates* linguistics from ontology.

4.1 The Decoupling Objection. One might worry that insulating linguistics from ontology is ill-advised: that Linguistic Instrumentalism dissolves the Puzzle only by decoupling language and reality, rendering utterances devoid of truth-conditions and thus compatible with any state of the world whatsoever. Call this *The Decoupling Objection*.

It is widely held to be a datum that competent speakers of a language L are able to recognize the truth-conditions of sentences in L, and that a desideratum on semantic theories is to explain that fact. For example, Stanley (2001, 41) says that it is a constraint on semantics that 'we are able smoothly to grasp the truth conditions of novel…sentences on the basis of our familiarity with their parts', but the claim is ubiquitous.[27]

---

[25] It is neutral on the *nature* of meanings$_S$: perhaps they're propositions, or thoughts, or truth-conditions, or whatever.

[26] Compare Grice's 'ontological Marxism', where he defends commitment to 'queer' entities so long as they 'help with the housework' and 'are not detected in illicit logical behaviour' (1975, 30–1).

[27] See, e.g., Heim and Kratzer 1998, 1; Jacobson 2014, 21; Yalcin 2014, 20; King 2018, 784, etc. Radical Contextualists might object, holding that we need pragmatics to derive (full) truth-conditions (Cappelen and Lepore 2010). We do not mean to dismiss Radical Contextualism, but we are assuming for the sake of argument that semantics assigns truth-conditions. If it doesn't, that just makes The Decoupling Objection easier to answer.

Aristotle would insist that there must be at least a *grain* of truth in this common consensus, and so there is. Competent speakers must recognize the truth-conditions of sentences—but only in a shallow and superficial sense. Linguistic competence enables speakers to recognize *some* facts about truth-conditions, and it is a constraint on semantics that it should explain how. Language and reality are not *completely* decoupled. But linguistic competence only requires speakers to recognize *shallow* facts about truth-conditions, and emphatically *doesn't* require speakers to recognize *philosophically interesting* facts about truth-conditions, if, indeed, it is true that ordinary assertions *have* philosophically interesting truth-conditions. So, language and the fundamental nature of reality are *largely* decoupled. This is what allows The Puzzle to be dissolved. Such significant but incomplete decoupling is both plausible and desirable. The reason it is *desirable* is that it dissolves The Puzzle. But explanation is required to see why it is *plausible*. The challenge, as Stanley (2001, 44) puts it, it to 'give a successful account of how we could assign ontologically innocent truth conditions to ontologically promiscuous discourse'.

4.2 Instrumentalist Semantics. The two leading traditions in semantics are inspired by the pioneering work of Donald Davidson (1967a) and Richard Montague (1970). In what follows we present a simplified 'toy' semantic theory in each of these traditions and investigate the role of instrumentalism therein.

4.2.1 Davidsonian Semantics. Davidsonian approaches to semantics minimize the assignment of entities as meanings. Instead, there are meaning *axioms* for sub-sentential expressions which systematically entail meaning *theorems*—T-sentences; statements of truth-conditions—for whole (declarative) sentences. Sentences are thus not paired with propositions, but assigned (typically homophonic) truth-conditions without any sort of propositional intermediary. To get the flavor of this approach, consider the following meaning axioms and the truth-condition they generate:

(DR) The name 'Pegasus' refers to Pegasus.

(DP) For any name *N*, '*N* flies' is true iff the referent of *N* flies.

(DTC) 'Pegasus flies' is true iff Pegasus flies.

Because Davidsonian approaches minimize the assignment of entities as meanings, there is relatively little opportunity for ontology to be misaligned with semantics. But only *relatively* little: (DR) (classically) entails that Pegasus exists.[28] It seems that ontology must cooperate with Davidsonian semantics by supplying referents for names and other singular terms. While this may be a problem—not all sentences containing empty names seem meaningless—it is a relatively *limited* problem. If we could be assured that most names weren't empty, it might not be worth worrying about. But we *can't* be assured that most names aren't empty. Compositional nihilism entails that pretty much all names are empty, organicism entails that the names of inorganic things are empty, etc. (van Inwagen 1990). Davidsonian semantics thus remains hostage to ontological fortune.

[28] Further developments of Davidson's approach plausibly generate further commitments: domains of quantification need to be populated, perhaps with events in addition to individuals (Davidson 1967b), etc.

4.2.2 Davidsonian Instrumentalism. Or at least apparently: while some Davidson*ians* adopt non-classical 'free' logics to skirt the problem of empty names (Sainsbury 2005), Davidson himself seemed to think that such maneuvers were unnecessary:

> words, meanings of words, reference, and satisfaction are posits we need to implement a theory of truth. They serve this purpose without needing independent confirmation or empirical basis…[Satisfaction and reference are] notions we must treat as theoretical constructs whose function is exhausted in stating the truth conditions for sentences. (Davidson 1977, 254–5)

This sounds amenable to instrumentalism:[29] the linguistic principles ('meaning axioms') for singular terms are *useful* for systematically generating truth-conditions for sentences, but they can do that while being untrue. Once sentential truth-conditions (meanings_S) have been generated, the ladder of reference can be kicked away. This sort of view, if tenable, is exactly what we are looking for: a way of generating meanings_S without generating ontological commitments in the process.

But what about the ontological commitments of meanings_S themselves? Common sense requires that semantics place *some* constraint on the truth-conditions of sentences, but dissolving the Puzzle requires that constraint to be *shallow*: largely or wholly neutral with respect to questions of abstruse metaphysics. Disquotational truth-conditions satisfy this joint desideratum nicely. Consider

(S) 'Snow is white' is true iff snow is white.

(S) rules out a lot: it allows us to make firm judgements about the truth-value of 'snow is white' in many cases. But (S) is largely or wholly metaphysically *neutral*. The truth-conditions it gives seem compatible with any account of the metaphysics of color or the ontology of snow. As Davidson claims, 'The theory reveals nothing new about the conditions under which an individual sentence is true; it does not make those conditions any clearer than the sentence itself does' (1967a, 311).

What we know when we grasp Davidsonian truth-conditions is shallow and philosophically uninteresting. If you took high-school German, you know that 'Smith weiß, dass Jones einen Ford besitzt' is true iff Smith knows that Jones owns a Ford. That's a useful piece of information, but it doesn't tell you anything *philosophically interesting* about the conditions under which 'Smith weiß, dass Jones einen Ford besitzt' is true. After all, if you knew what those conditions were in a philosophically interesting sense, you'd be a famous epistemologist (cf. J. A. Keller 2015, 117).

But what about 'Pegasus flies'? Just how robust a grasp of the truth-conditions of 'Pegasus flies' does competence require? Not very: after all, competent speakers don't *have* a robust grasp of those truth-conditions. We know the relevant 'facts on the ground', as it were, regarding Pegasus. But if we had to bet our lives on whether 'Pegasus flies' were true, we'd be terrified. Some competent speakers think 'Pegasus flies' is false, since 'Pegasus' doesn't refer. Some think it's true, since 'Pegasus' refers to a flying creature of fiction. As competent speakers, we grasp (DTC). But that 'reveals nothing new', as

[29] Talk of Davidson's (flirtation if not consummation with) instrumentalism is ubiquitous: see, e.g., Williams 2013.

Davidson puts it, about the conditions under which 'Pegasus flies' is true. Our grasp of those conditions is shallow and metaphysically neutral.[30]

Note that this isn't just a point about *Davidsonian* truth-conditions. It's a point about our actual grasp of truth-conditions. As a matter of empirical fact, that grasp is very shallow, just as required for dissolving The Puzzle. The distinction between shallow-linguistic and deep-metaphysical truth-conditions is intuitive, although rarely drawn.[31] As Sarah-Jane Leslie says, metaphysical truth-conditions:

> should not be mistaken for semantically derived truth conditions…[If] a dispositionalist theory of color is correct,…'Bob is red'…is true if and only if Bob is experienced as red by standard observers in standard conditions. This is a specification of the circumstances in the world that must obtain for 'Bob is red' to be true. Such a specification does not tell us anything about the semantically derived, compositionally determined truth conditions for 'Bob is red', however…It is in no way part of semantic competence to recognize that the truth of 'Bob is red' entails that there exist standard observers…This is not plausibly a semantic entailment, but merely a metaphysical one. The semantic truth conditions for 'Bob is red' may well be no more than Red(Bob). This respects the compositional structure of the sentence…For this reason, and others, it is very often desirable to simply disquote individual expressions when giving semantic truth conditions. Any further analysis of individual expressions very often belongs to metaphysics rather than to semantics…(2008, 43–4)

4.2.3 Logical Form. A close cousin to The Decoupling Objection is the concern that semantics should ground competent speakers' knowledge of the *logical forms* of expressions.[32] Logical form in this sense is at least conceptually-distinct from LF; it refers to the 'formal' properties of utterances or meanings$_S$ in virtue of which they stand in (formal) logical relations (entailment, consistency, etc.).

If competent speakers knew the logical forms of expressions, it would be reasonable to think that semantics must explain that. But they don't, so it isn't. Competent speakers must recognize trivial analytic entailments, like 'Gabriel is short' entails 'Gabriel is not tall'. But they don't *have to* recognize the logical forms of expressions, because they *don't* recognize the logical forms of expressions. 'Translating' ordinary sentences into the regimented idiom of logic is non-trivial even for trained philosophers, much less ordinary speakers. The logical forms of, e.g., 'It's raining', 'Pegasus does not exist', and 'Izzy believes that Pegasus flies' remain highly controversial.[33]

---

[30] The main texts suggests that (possibly untrue) meaning axioms are used to derive true meaning theorems (statements of truth-conditions). But (DTC) itself may not be metaphysically neutral, since it arguably entails the existence of Pegasus, or at least something that flies iff 'Pegasus flies' is true. Perhaps, then (DTC) is false. It is nonetheless correct: grasping it enables us to determine the truth value of 'Pegasus flies' in many circumstances. (E.g., 'Pegasus flies' is false in circumstances where Pegasus is an existent but non-flying horse.) Note that we can clearly *use* non-truth-apt conditions for sorting: e.g., I can reliably (albeit imperfectly) determine what things satisfy the condition 'bigger than Bilbo'.

[31] Similar *distinctions* are defended in King 2002, Williams 2010, and Sider 2011, and similar *views* are defended in Johnson 2007, Szabó 2013, Glanzberg 2014, and Yalcin 2018.

[32] We take this to be a plausible explication of the common claim that semantics grounds speaker's knowledge of 'entailment facts' (Yalcin 2014, 20, King 2018, 784).

[33] See J. A. Keller 2015, L. J. Keller 2019, King 2002, and Szabó 2012 for discussion.

Given this, we focus in what follows on the plausibility of instrumentalism about semantic *principles* without worrying about the truth-conditions or entailments of the meanings$_S$ those principles generate.

4.2.4 Montague Semantics. To get a sense of how Montagovian theories work, consider the following toy semantics:

(MR) [[Pegasus]] = Pegasus

(MP) [[flies]] = the *f* such that for any *x*, $f(x) = \begin{cases} 1, & x \text{ flies} \\ 0, & \text{otherwise} \end{cases}$

(MS) [[Pegasus flies]] = $[[flies]]([[Pegasus]]) = \begin{cases} 1, & Pegasus\ flies \\ 0, & otherwise \end{cases}$

Understanding this requires a little unpacking. '[[*x*]]' denotes the semantic value of *x*. Here, the idea is that [[Pegasus]] is an entity and [[flies]] is a function from entities to truth values. Truth values—truth and falsity—are typically identified with 1 and 0. The function [[flies]] thus maps entities that fly to 1 and entities that don't to 0. (MR) and (MP) thereby generate (MS), the meaning$_S$ of 'Pegasus flies'. Our toy semantics entails that [[Pegasus flies]] is 1 or 0, depending on whether Pegasus flies. But that individuates meanings$_S$ much too coarsely: there are only two of them! Non-toy versions of Montague semantics identify meanings$_S$ with *intensions*: functions from possible words to truth-values, or more typically their characteristic sets. Thus, [[Pegasus flies]] would be the set of worlds at which Pegasus flies:

(IMS) [[Pegasus flies]] = {*w*: [[flies]]^w([[Pegasus]]^w)=1} = {*w*: Pegasus flies @ *w*}.

Montague semantics essentially *identifies* meanings$_S$ with truth-conditions.[34] Because semantic values are things, there appears to be abundant opportunity for ontology and semantics to misalign on this view. Our toy semantics appealed to (purported) ordinary entities like Pegasus, functions, and truth values; (IMS) adds possible worlds to the mix, and typical non-toy theories add times, events, degrees, kinds, situations, and vectors (see Rett 2020).

4.2.5 Montagovian Instrumentalism. How plausible is instrumentalism about Montagovian semantics? Can we use Montagovian principles to pair utterances (or, really, the output of phonological and morphological processing) with meanings$_S$, and then kick away the ladder once we're done? There is evidence that Montagovian *theorizing* is not generally undertaken in a realist vein. The identification of truth values with 1 and 0 is a tell: whatever truth and falsity are, they aren't 1 and 0! Sometimes the anti-realism is explicit, as when Pauline Jacobson says 'There is little, if any, reason to choose between the unary rule and the silent lexical item approach; they ultimately may just be different metaphors for the same thing' (2014, 123–4).

More generally, theorizing in this tradition exhibits a sort of reckless indifference to ontology. This seems like strong evidence that the resulting theories are not intended to be truth-apt. For example,

[34] Note that this truth-condition—'Pegasus flies' is true @ w iff Pegasus flies @ w—is just as shallow as the Davidsonian truth-condition discussed above.

Derek Ball counsels against ontological scruples in semantic theorizing in his response to a version of Benacerraf's dilemma for the semantics of gradable adjectives:

> Semantics is said to be the study of meaning. If we took a formal semantic theory to be making a claim about what meanings are (identical to)—for example, that the meaning of *tall* is a function from individuals to sets—then the proliferation of candidates would be a problem and formal semantics would be in trouble. But that should not be our attitude. (2018, 389)

Similarly, Zoltan Szabó writes:

> Linguists want to say what the meanings of various expressions are without having to say too much about them. The situation resembles that of the mathematician: she wants to say enough about the number 2 to prove that its square root is irrational, but not so much as to take sides in the metaphysical debate among Platonists, intuitionists and formalists. The aim is to present a set of minimal commitments. (2013, 39)

Semantic theorizing in the Montagovian tradition thus seems not only compatible with instrumentalism, but (often) to *be* instrumentalist.[35] Of course, this is mere *wordy* anti-realism. How plausible is it that there is a worldly analogue? At least somewhat. We argued in §2.8 that linguistic principles *can* do their job of associating utterances with meanings$_S$ even if they are untrue.[36] Given that, it's plausible that that's how they *do* do their job. Language is, we assume, the product of evolution. Evolution cares about *efficiency*, not truth—and there is almost always a trade-off between them. This at least suggests that we should not expect evolved linguistic principles to have only true presuppositions.

4.2.6 Mathematical Meanings. We have said that we are realists about mathematics. Some hold that formal semantics in the Montagovian tradition *is* basically just math (Pickel 2019). Does Linguistic Instrumentalism thus conflict with mathematical realism? We don't think so. Linguistic Instrumentalism allows that *some* linguistic theories are truth-apt. So if formal semantics turned out to be truth-apt, that wouldn't necessarily be an objection. But we don't think formal semantics *is* truth-apt. Formal semantics might be *basically* math, but it isn't *just* math: it's mathematical operations on atomic meanings (the meanings of morphemes). And so if atomic meanings don't exist, formal semantics will have false ontological presuppositions.

Let's see how all of this fits together. Hearers are presented with 'utterances'—relatively undifferentiated and 'fuzzy' sequences of sounds, movements, or marks—something like an auditory version of

(U) **pegasusflies.**

On standard views, phonological and morphological principles assign (project onto) (U) structure that it doesn't have, generating something like the following (mis)representation of (U)

(R1) Pegasus flies.

---

[35] See Ball 2018, Schwarz 2018, and Yalcin 2018 for fairly explicit arguments to this effect.

[36] Note that Davidsonian and Montagovian semantic *theories* both work—both associate utterances with meanings$_S$—regardless of their truth.

Syntactic principles then assign (project onto) (U) further structure, and *lexical* semantics assigns (plausibly non-existent) lexical semantic values, thus generating a 'richer' (mis)representation of (U):

(R2) $[[\text{flies}]]^{w}([[\text{Pegasus}]]^{w})$.

Finally, *compositional* semantics computes the function in (R2), yielding something like

(IMS) $\{w\text{: Pegasus flies @ } w\}$.

We are realists about utterances like (U) and truth-conditions like (IMS)—and about meaningss more generally—but instrumentalists about the attributions/projections involved in the intermediate steps between them: about the correctness of representations like (R1) and (R2). That is, we think representations like (R1) and (R2)—and the linguistic principles that generate them—are (often) correct but untrue. The mathematical step to (IMS) doesn't involve attributions or projections, and so we are happy to be realists about that. We are neutral on whether there *is* such a mathematical step—this is just one common picture—but if there is we are happy to be realists about it.

We have also argued that speakers' grasp on truth-conditions like (IMS) is shallow and largely metaphysically neutral. (Likewise with their grasp of logical forms.) Thus, Linguistic Instrumentalism minimizes potential conflicts between the ontology and the truth-conditions of anodyne utterances—and the ontological commitments of the linguistic principles governing them.

Our view is not entirely novel. Rey (2020) defends a similar picture—he is skeptical about the existence of words, syntactic structure, sub-sentential semantic values, etc.—but his *interpretation* of the picture is quite different. Because Rey is a linguistic *realist*, he endorses a conglomerate of murky and implausible ontological and ideological theses about 'intentional inexistence' in order to make sense of his picture. It seems preferable to simply acknowledge that linguistic principles can be correct—*useful* for generating (shallow) truth-conditions—without being true.

5. CONCLUSION

We end with a note of possible *rapprochement*. Wittgenstein (1929, 164–5) famously said that 'we can draw no conclusions—except very vague ones—from [grammatical form] as to the actual logical form of the phenomena described'. Gaskin *contrasts* this realist sentiment with Linguistic Idealism, according to which,

> there is no question of our 'having to worry', in general, that an object 'perhaps does not exist': if my language is meaningful—and the criteria for that are internal to the language-game itself, a matter of whether the signs I purport to employ do indeed have use—then whatever objects my words can be considered to refer to *eo ipso* exist. To say that my words refer to objects is just another way of saying that these words are indeed meaningful. (this volume, xx)

We, on the other hand, take these sentiments to be perfectly compatible, and hope the congruence between these passages and what we've argued here is as striking to the reader as it is to us. The idea that linguistic principles should be evaluated 'internally' is a core component of Linguistic Instrumentalism, but also of Linguistic Idealism. We are drawn to anti-realism about linguistics because we wish to maintain a robust realism about both metaphysics and our ordinary thought and talk. Perhaps Gaskin would say that our realism is merely *empirical*, perfectly compatibly with some

sort of *transcendental* idealism. We brushed aside the distinction between empirical and transcendental realism above (in n. 7) because, frankly, we can make neither head nor tail of it. We take the inability to make sense of that distinction to be a mark of a true realist. But perhaps Gaskin would argue that it is instead the mark of a true transcendental idealist. *Whereof one cannot speak thereof one must be silent.*[37]

BIBLIOGRAPHY

Adger, D. 2019: *Language Unlimited* (Oxford: OUP).

Anscombe, G. E. M. 1981: *From Parmenides to Wittgenstein* (Cambridge: CUP).

Appiah, A. 2017: *As If: Idealization and Ideals* (Cambridge: Harvard University Press).

Ashworth, E. J. 2015: 'Medieval Theories of Signification to John Locke', in M. Cameron and R. Stainton (eds.), *Linguistic Content: New Essays in the Philosophy of Language*, (Oxford: OUP), 156–75.

Ball, D. 2018: 'Semantics as Measurement', in Ball and Rabern 2018, 381–410.

Ball, D. and Rabern, B. 2018: *The Science of Meaning: Essays on the Metatheory of Natural Language Semantics* (Oxford: OUP).

Benacerraf, P. 1965: 'What Numbers Could Not Be,' *The Philosophical Review* 74, 47–73.

Bloomfield, L. 1936: 'Language or Ideas', *Language* 12, 89–95.

Boyd, R. 1983: 'On the Current Status of Scientific Realism', *Erkenntnis* 5/19, 45–90.

Braun, D. 1993: 'Empty Names', *Nous* 27, 449–69.

Cappelen, H. and Lepore, E. 2005: *Insensitive Semantics* (Oxford: Blackwell).

Cartwright, N. 1983: *How the Laws of Physics Lie*, (Oxford: OUP).

Chomsky, N. 1965: *Aspects of the Theory of Syntax* (Cambridge: MIT Press).

——1966: 'Explanatory Models in Linguistics', in *Studies in Logic and the Foundations of Mathematics* 44, 528–50.

——2000: *New Horizons in the Study of Language* (Cambridge: CUP).

——2017: *On Language* (New York: The New Press).

Davidson, D. 1967a: 'Truth and Meaning, *Synthese* 17, 304–23.

——1967b: 'The Logical Form of Action Sentences', in N. Rescher (ed.), *The Logic of Decision and Action*, (Pittsburgh, PA: University of Pittsburgh), 81–95.

——1977: 'Reality Without Reference', *Dialectica* 3, 247–58.

Doyle, J. P. 2001: *The Conimbricenses: Some Questions on Signs, Translated and Introduction with Notes* (Milwaukee: Marquette University Press).

Dupoux, E., Kakehi, Y., Hirose, Y., Pallier, C., and Mehler, J. 1999: 'Epenthetic vowels in Japanese: a perceptual illusion?' *Journal of Experimental Psychology: Human Perception and Performance* 25, 1568–78.

Gaskin, R. 2019: 'From the Unity of the Proposition to Linguistic Idealism', *Synthese* 196, 1325–42.

——2021: *Language and World: A Defense of Linguistic Idealism* (New York: Routledge).

Glanzberg, M. J. 2014: 'Explanation and Partiality in Semantic Theory', in A. Burgess & B. Sherman (eds.) *Metasemantics: New Essays on the Foundations of Meaning* (Oxford: OUP), 259–92.

Grice, H. P. 1975: 'Method in Philosophical Psychology: From the Banal to the Bizarre', *Proceedings and Addresses of the American Philosophical Association* 48, 23–53.

Heim, I, and Kratzer, A. 1998: *Semantics in Generative Grammar* (Malden, MA: Blackwell).

Jacobson, P. I. 2014: *Compositional Semantics* (New York: OUP).

[37] Wittgenstein (1922, 23). Thanks to Derek Ball, Audre Brokes, Nathaniel Bulthuis, Joseph Corabi, Michael Kates, Bryan Pickel, Jeff Speaks, and especially Richard Gaskin for helpful comments and discussion. They should not be held accountable for any of the many deep flaws we are sure this paper contains.

Johnson, K. 2007: 'The Legacy of Methodological Dualism', *Mind and Language* 22, 366–401.
Jubien, M. 2001: 'Propositions and the Objects of Thought', *Philosophical Studies* 104, 47–62.
Katz, J. J. 1985a: *The Philosophy of Linguistics* (Oxford: OUP).
——1985b: 'An Outline of Platonist Grammar', in Katz 1985a, 17–48.
Keller, J. A. 2014: 'Theological Anti-Realism', *Journal of Analytic Theology* 2, 13–42.
——2015: 'Paraphrase, Semantics, and Ontology', in K. Bennett and D. Zimmerman (eds.) *Oxford Studies in Metaphysics* 9, (Oxford: OUP) 89–128.
Keller, L. J. 2019: 'What Propositional Structure Could Not Be', *Synthese* 196, 1529–53.
King, J. C. 2002: 'Two Sorts of Claim about 'Logical Form'', G. Preyer and G. Peter (eds.) *Logical Form and Language* (Oxford: OUP), 118–31.
——2007: *The Nature and Structure of Content* (Oxford: OUP).
——2018: 'W(h)ither Semantics!(?)', *Noûs* 52, 772–95.
Kripke, S. 1980: *Naming and Necessity*, (Cambridge, MA: Harvard University Press).
Leslie, S. J. 2008: 'Generics: Cognition and Acquisition', *Philosophical Review*,117, 1–47.
Ludwig, K. 2002: 'What is the Role of a Truth Theory in a Meaning Theory?', in J. K. Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth: Investigations in Philosophical Semantics* (New York: Seven Bridges Press), 142–63.
Machery, E. 2009: *Doing Without Concepts* (New York: OUP).
McGilvray, J. 1998: 'Meanings are Syntactically Individuated and Found in the Head', *Mind & Language* 13, 225–80.
Meinong, A. 1960: 'The Theory of Objects', in R. M. Chisolm (ed.) *Realism and the Background of Phenomenology* (Glencoe, IL: Free Press), 76–117.
Montague, R. 1970: 'English as a formal language', in B. Visentini, et al. (eds.), *Linguaggi nella Società e nella Tecnica* (Milan: Edizioni di Communita), 189–224.
Pickel, B. 2019: 'Structured Propositions in a Generative Grammar', *Mind* 128, 329–66.
Pietroski, P. 2018: *Conjoining Meanings: Semantics without Truth Values* (Oxford: OUP).
Pitt, D. 2018: 'What Kind of Science is Linguistics?' in C. Behme and M. Neef (eds.), *Essays on Linguistic Realism* (Amsterdam: John Benjamins Publishing Company), 7–20.
Quine, W. V. O. 1960: *Word and Object* (Cambridge: MIT Press).
Reiland, I. 2023: 'Regulative Rules: A Distinctive Normative Kind', *Philosophy and Phenomenological Research*, 2 July 2023, 1–20 (**https://doi.org/10.1111/phpr.13008**).
Rett, J. 2022: 'A Typology of Semantic Entities', D. Altshuler (ed.) *Linguistics Meets Philosophy* (New York: CUP), 277–301.
Rey, G. 2006: 'The Intentional Inexistence of Language—but not Cars', in R. Stainton (ed.) *Contemporary Debates in Cognitive Science*, (Oxford: Wiley-Blackwell), 237–55.
——2020: *Representation of Language: Philosophical Issues in a Chomskyan Linguistics* (Oxford: OUP).
Rosen, G. 1990: 'Modal Fictionalism', *Mind* 99, 327–54.
Russell, B. 1980: 'Correspondence with Frege', in G. Frege, *Philosophical and Mathematical Correspondence* (Chicago: University of Chicago Press), 130–1.
Russell, G. K. forthcoming: 'How the Laws of Logic Lie', *Episteme*.
Sainsbury, R. M. 2005: *Reference Without Referents* (Oxford: OUP).
Santana, C. 2016: 'What is Language?', *Ergo* 3, 501–23
Sapir, E. 1921: *Language* (New York: Harcourt).
Schwarz, W. 2018: 'Semantic Possibility', in Ball and Rabern 2018, 361–80.
Searle, J. 1969: *Speech Acts* (New York: CUP).
Sider, T. 2011: *Writing the Book of the World* (New York: OUP).
Stanley, J. 2001: 'Hermeneutic Fictionalism', *Midwest Studies in Philosophy* 25, 36–71.

Szabó, Z. G. 2012: 'Against Logical Form', in I. G. Preyer (ed.) *Donald Davidson on Truth, Meaning, and the Mental*, (Oxford: OUP), 105–26.
——2013: *Problems of Compositionality* (Hoboken: Routledge).
Thomasson, A. L. 2015: *Ontology Made Easy* (New York: OUP).
van Fraassen, B. 1980: *The Scientific Image* (Oxford: OUP).
van Inwagen, P. 1990: *Material Beings* (Ithaca: Cornell University Press).
Williamson, T. 2007: *The Philosophy of Philosophy* (Oxford: OUP).
——2018: 'Model-Building as a Philosophical Method', *Phenomenology and Mind* 15, 16–22.
Wittgenstein, L. 1922. *Tractatus Logico-Philosophicus*, C. K. Ogden (trans.), (London: Kegan Paul).
——1929: 'Some Remarks on Logical Form', *Proceedings of the Aristotelian Society,* Supplementary Volume 9, 162–71.
Williams, J. R. G. 2010: 'Fundamental and Derivative Truths', *Mind* 119, 103–41.
——2013: 'Reference' in E. Lepore and K. Ludwig (eds.) *A Companion to Donald Davidson* (Oxford: Wiley-Blackwell), 264–86.
Yalcin, S. 2014: 'Semantics and Metasemantics in the Context of Generative Grammar', in A. Burgess and B. Sherman (eds.) *Metasemantics*, (Oxford: OUP), 17–54.
——2018: 'Semantics as a Model-Based Science', in Ball and Rabern 2018, 334–60.
Zsiga, E. 2020. *The Phonetics/Phonology Interface* (Edinburgh: Edinburgh University Press).