Chapter 3
Priorities and Diversities in Language and Thought

Elisabeth Camp

Abstract Philosophers have long debated the relative priority of thought and language, both at the deepest level, in asking what makes us distinctively human, and more superficially, in explaining why we find it so natural to communicate with words. The "linguistic turn" in analytic philosophy accorded pride of place to language in the order of investigation, but only because it treated language as a window onto thought, which it took to be fundamental in the order of explanation. The Chomskian linguistic program tips the balance further toward language, by construing the language faculty as an independent, distinctively human biological mechanism. Devitt (Ignorance of language. Clarendon Press, Oxford, 2006) attempts to swing the pendulum back toward the other extreme, by proposing that thought itself is fundamentally sentential, and that there is little or nothing for language to do beyond reflecting the structure and content of thought. I argue that both thought and language involve a greater diversity of function and form than either the Chomskian model or Devitt's antithesis acknowledge. Both thought and language are better seen as complex, mutually supporting suites of interacting abilities.

Keywords Systematicity · Language of thought hypothesis · Non-sentential logic · Modularity · Maps · Diagrams · Discourse structure · Illocutionary-force-indicating devices · Williams Syndrome

This article grew out of an Author Meets Critics session at the 2007 Pacific APA; thanks to audiences there, and especially to Michael Devitt. I also presented some of this material at the 2007 Workshop in Philosophy of Linguistics in Dubrovnik, addressing work by Peter Ludlow; thanks to audiences there, and especially to Peter Ludlow, Brian Epstein, and Gurpreet Rattan. Additional thanks to Josh Armstrong and Paul Pietroski, and to Carolina Flores for editorial work and advice.

E. Camp (✉)
Department of Philosophy, Rutgers University, New Brunswick, NJ, USA
e-mail: emc233@philosophy.rutgers.edu

A. Bianchi (ed.), *Language and Reality from a Naturalistic Perspective*, Philosophical Studies Series 143, https://doi.org/10.1007/978-3-030-47641-0_3

Which comes first, thought or language? Some sort of thought-first model has considerable intuitive pull. Indeed, although analytic philosophy has fixated on language since its inception, this interest has generally been driven by the assumption that language is important primarily or only because it affords our most direct, transparent window onto the structure and content of thought (Dummett 1994). In particular, while Frege and Russell often inveighed against the inconstancies and confusions of ordinary discourse, they proposed formal logics as an improved means for accomplishing the aim they took actual languages to achieve only imperfectly: of transparently reflecting the structure and content of thought. And much subsequent philosophical analysis has been driven by the assumption that ordinary language can itself be revealed to be in perfect logical-conceptual order, given a sufficiently ingenious mapping from surface to logical forms (Wittgenstein 1921/2001) and a sufficiently sophisticated understanding of the relation between words' meanings and the uses to which speakers put them (Grice 1975).

Thus, despite its intense focus on language as a topic of investigation, much of twentieth-century analytic philosophy embraced a more fundamental focus on thought at the level of explanation. The Chomskian linguistic program extended the "linguistic turn" to this more fundamental level at which philosophers still prioritized thought and concepts, instead treating language as an explanatory end in its own right. But it did this in part by understanding 'language' in a very particular way, which it takes to be more theoretically tractable and scientifically pertinent than either the ordinary or philosophical construals. Analytic philosophers and laypeople typically treat language as a public phenomenon: a system of conventions for using certain sounds to stand for certain ideas, which each of us accede to by "tacit consent" (Locke 1689) in order to make ourselves understood. By contrast, Chomskians accord primacy to the sets of sentences generated by an individual speaker's linguistic knowledge, and inquire into the essential characteristics of the mechanism that generates them. More specifically, they investigate the 'language faculty', construed as an innate biological mechanism for generating syntactically complex representational forms which map conceptual meanings to sounds within an individual speaker's idiolect. Any connection to public conventions for use in communication is for them secondary at best, and perhaps altogether illusory.

It is easy to feel that something has been lost in the shift away from language conceived as a public system for expressing thoughts and toward language conceived as an individual system for generating formal structures. In *Ignorance of Language*, Michael Devitt aims not just to restore the balance, but to shift it entirely to the side of thought. While much of the book focuses on various negative arguments against what he takes to be current linguistic methodology, his positive argument for the priority of language is fairly direct. He begins with the claim that "language expresses thought" (2006: 128), which he takes to be "relatively uncontroversial" in itself, but to entail the exuberantly controversial tenet that "[t]hought has a certain priority to language ontologically, explanatorily, temporally, and in theoretical interest" (276). He then argues that thought itself is sententially structured; and concludes that there is thus "little or nothing" for the language faculty to do beyond matching sounds to complex mental representations, in a way that can be accomplished by "fairly brute-causal associationist processes." All the wondrous complexity of contemporary linguistics really belongs to thought instead: "Humans are predisposed to learn languages that conform to the rules specified by UG because those rules are, largely if not entirely, innate structure rules of thought" (276).

While Devitt's conclusion is strong and surprising, the main premisses are widely accepted in some form. My own theoretical proclivities also lie on the thought-first side of the seesaw; and I share Devitt's conception of language as a public, social construction. However, I see no reason to reject the existence of a substantive, biologically-based language faculty. More importantly, I think humans employ a range of formats for thought, both naturally and by enculturation. And I think language does much more than express thought. Thus, I will argue that neither thought nor language can be assigned clear explanatory priority over the other. In particular, instead of either a single "language faculty" or a single set of "rules of thought," it is more plausible to posit complex suites of distinct, interacting abilities that add up to make certain ways of talking and thinking very natural for us. I consider the function and format of thought and language in §3.1 and §3.2, respectively. In §3.3, I argue that the constraints imposed by Universal Grammar are more plausibly explained as originating from language, as Chomskians maintain, rather than thought, as Devitt proposes.

3.1 The Language of Thought and Diversities in Cognitive Format

The first step in Devitt's broadside for the priority of thought is establishing that thought itself has a sentential structure. To support this conclusion, he invokes the familiar Fodorian inference to the best explanation from systematicity and productivity in observed behaviors to the existence of a representational system with recurrent, systematically recombinable parts. Devitt forthrightly admits that the Language of Thought Hypothesis (LOTH) is "controversial" (142). However, it plays a central role in his overall argument for the priority thesis – and indeed, the specific role it plays lends it an additional degree of controversy. That is, for Devitt, as for Fodor and many others, the most direct and compelling source of evidence for LOTH is the systematicity and productivity of human speech. As we'll discuss in §3.2, the claim that human speech is indeed so highly systematic is contested by various philosophers of language and linguists. But even accepting that we do observe such systematicity, the direction of causal influence remains an open question: perhaps thought is systematic because and to the extent that language is. While this is a thesis on which most proponents of LOTH can remain neutral, in order for Devitt to establish that thought is language-like in a non-question-begging-way, he needs independent evidence for the systematicity of thought that doesn't rely on language. Moreover, to establish his ultimate conclusion that language exhibits the particular features it does *because* thought possesses those features, he needs independent evidence not just that thought is highly systematic, but that it has a specifically sentential structure. In this section, I argue that if we bracket off linguistic evidence about the format of thought, then the case that thought has distinctively sentential format becomes much weaker.

The Language of Thought Hypothesis is amenable to at least two construals (Camp 2007). On the stronger construal, thought is claimed to possess a distinctively linguistic structure; on the weaker one, it is merely like language in being a compositional representational system. For Fodor's central aim of defending computationalism against connectionism (Fodor and Pylyshyn 1988), the weaker construal suffices. Fodor himself consistently extends his arguments to nonhuman animals (e.g. 1987), and offhandedly assumes that pictorial representational models are compositional (2007). Like Fodor, Devitt recognizes that an appeal to representational complexity doesn't entail the strong claim about specifically sentential structure, because maps and other non-linguistic representations have a syntax that is "very different" from language (146). And like Fodor, Devitt appeals to the cognitive states of non-human animals as evidence about the nature of thought – in his case to establish thought's temporal priority over, and contemporary independence from, language (131).

So, the argument from systematicity alone does not justify an inference to sentential structure, especially in the current argumentative context; and Devitt acknowledges this. However, he argues that the need to explain the *processes* of thought does. "Formal logic," he says, "gives us a very good idea of how thinking might proceed" (146–147); by contrast, we "have very little idea how thinking could proceed if thoughts were not language-like" (147). Devitt says very little about what he means by 'formal logic', but he appears to have something like a traditional predicate calculus in mind. Bermudez (2003: 111) makes the same claim more explicitly: "We understand inference in formal terms – in terms of rules that operate on representations in virtue of their structure. But we have no theory at all of formal inferential transitions between thoughts that do not have linguistic vehicles" (see also e.g. Rey 1995: 207).

Although this assumption is common – and understandable, given the intimate historical connection between analytic theorizing about inference and the development of predicative logic – it's not true that formal sentential logic provides our only model for "how thinking could proceed" in general. Recent successes with connectionist models of "deep learning" have challenged the computational orthodoxy (e.g. Schmidhuber 2015); while hierarchical Bayesian models have introduced probabilistic inference to computational methodology in ways that function very differently from traditional logics (Tenenbaum et al. 2011). So it is not obvious that systematic cognitive abilities must be implemented by a system of representational vehicles which are comprised of recurrent symbolic parts governed by fixed, formally-specified rules. However, even assuming that they must be implemented by such a system, a diversity of representational formats can satisfy this criterion.

First, maps – both those, like seating charts, that exploit a finite base of elements and principles of composition, and also those, like road atlases, that exploit potentially continuously varying shapes, colors, and textures – can be constructed and interpreted by means of formal principles (Pratt 1993; Casati and Varzi 1999; MacEachren 2004). These principles depart substantively from those of language (Rescorla 2009; Camp 2007, 2018a). And they can be exploited to define rules for updating and integrating distinct maps within a larger cartographic system (or from inter-translatable systems), so long as they represent regions that are themselves related in spatially appropriate ways (e.g. that are at least partially contiguous). Given a definition of validity that is not specifically linguistic, we can assess such transformations for validity (Sloman 1978). Finally, there is substantive psychological and neurophysiological evidence that both people and other animals do process spatial information, including abstract information about spatial relationships, in a distinctively spatial way (Morgan et al. 2011; Franconeri et al. 2012; Marchette et al. 2017).

Devitt's second reason for rejecting the hypothesis that thought might be structured like a map rather than a set of sentences is that maps are expressively limited in comparison to language (146). In comparison to the invocation of constraints on explaining processes of thought, this argument is more compelling. While the expressive limitations of maps are often exaggerated – in particular, ordinary maps can be enriched to represent negation, tense, disjunction, and conditionals in various ways – it is true and important that maps cannot represent information that is not spatial. Most notably, they cannot represent abstract quantificational information (Camp 2007, 2018a).

At the same time, though, there also exist diagrammatic systems, which are likewise formally defined and differ substantively from language (and from one another), and which have a much richer expressive range than maps (Shin 1994; Allwein and Barwise 1996). Moreover, some of these diagrammatic systems have robust, rigorous practical applications in science and mathematics (Tufte 1983; Giardino and Greenberg 2014). Indeed, De Toffoli (2017) argues that diagrams are useful in mathematical practice precisely because the *process* of using them – by manipulating constituent algebraic elements – constitutes a valid form of inferential 'calculation'. Likewise, diagrams can be distinctively useful in tracking information about abstract relations in the real world. In particular, directed graphs or 'Bayes nets' offer a rigorously defined diagrammatic format for representing and manipulating causal information, one that is arguably more effective than sentential logics at least for certain purposes (Pearl 2000; Elwert 2013), and that has been argued to implement causal knowledge in children (Gopnik et al. 2004) and possibly non-human animals (Camp and Shupe 2017).

Thus, given this diversity of formally definable, practically relevant representational systems, there can be no in-principle argument that thought *per se* must be sentential. At the same time, a more modest version of the appeal to expressive power can be used to establish that at least some human thought does have a distinctively sentential structure (Camp 2015). Language is distinguished from other representational formats by its abstractness, in at least three respects. First, it employs a highly *arbitrary* semantic principle mapping basic elements to values. Second, it employs a highly *neutral* or general combinatorial principle (e.g. predication, functional application, or Merge), which itself has only minimal representational significance. And third, its principles of construction and interpretation are defined entirely in terms of operations on the values of the basic elements, rather than on the vehicular elements themselves.

Most diagrammatic systems are like language, and unlike most maps, in employing a highly arbitrary semantic principle, largely freeing them of significant constraints on the types of values their constituents can denote. In contrast to maps, some diagrammatic systems also employ highly neutral combinatorial principles: for instance, Venn diagrams use spatial relations to represent set-theoretic relations among denoted entities. The relatively high abstractness and generality of those set-theoretic relations permits such diagrams to represent relations among a correspondingly wide range of entities. (By contrast, other diagrammatic systems employ principles with more robust significance, which impose commensurate expressive restrictions: for instance, because phylogenetic tree diagrams assign branching tree structures the significance of branching ancestry, they invariably represent the entities denoted by the nodes in a branching tree as related by ancestry and descent; Camp 2009a.)

However, even the most general diagrammatic systems fail to be fully abstract along the third dimension, of vehicular implementation. That is, simply in virtue of being diagrams, their construction and interpretation rules exploit the spatial (or topological) structure of their representational vehicles. And this inevitably generates some expressive restrictions: for instance, even sophisticated Venn diagrams can only represent relations among sets that can be implemented with closed continuous figures in a single plane (Lemon and Pratt 1997).

Thus, Devitt is right that language is distinctively expressively powerful, in virtue of its distinctively abstract semantic and combinatorial properties. Still, the class of complex relations that exceed the scope of diagrammatic representation is rather rarified, and so might not seem to constitute much of an argument from expressive generality for a sentential structure of thought. In response, an advocate for LOTH might point to the fact that ordinary human thought is highly intensional in order to suggest a more pervasive and relevant potential expressive restriction on non-sentential systems. Diagrammatic systems can represent at least some kinds of modality – for instance, Pearl (2009) argues that directed graphs are uniquely equipped to capture counterfactual causal inference. But most diagrammatic systems are extensional; and the best-developed and most general intensional logics are all extensions of the predicate calculus.

Nonetheless, even if we grant that intensional relations, as well as certain extensional relations among sets, can only be expressed in language, this still falls well short of establishing that “the innate structure rules of thought” have a sentential syntax and semantics, in the way Devitt needs. First, like the basic argument for LOTH, inferring that the logic of intensionality must be predicative relies on an appeal to a lack of available alternatives that is vulnerable to subsequent counterexemplification. Second and more generally, even if a formal predicate calculus does constitute our most rigorous general model for “how thinking could proceed” when we analyze “thought” in terms of the prescriptive “laws” of thought, it is frustratingly obvious that much, even most actual human thinking fails to conform to this model (Evans and Over 1996). And indeed, the various species of intensionality have proven to be especially recalcitrant to systematic formal analysis. Given this, models of thought that appeal to schemas and other partly abstract, partly iconic modes of representation may hold more promise for capturing the distinctive contours of actual ordinary human cognition, including especially intensionality (Johnson-Laird 2005).

Finally and most importantly, establishing that some of the contents that people sometimes think about can only, or most easily, be represented and manipulated sententially doesn't establish that all thought takes that form. Devitt, like many advocates of LOTH, implicitly assumes that thought is governed by a single set of innate structure rules; but empirical evidence suggests that humans regularly and spontaneously employ multiple representational formats. Here, one might argue for the centrality of sententially-structured thought on the grounds that its expressive generality uniquely equips it to integrate thoughts encoded in distinct formats (Carruthers 2003). But this too is a substantive argument by exclusion, which proponents of modularity can resist in various ways (Rice 2011). More importantly, it would still not establish language as the exclusive format for thought, only as the privileged vehicle for integration when it occurs. And advocates of cognitive modularity often point to the pervasive failure of full substantive integration in human cognition in support of a multiplicity of representational forms and structures (Fiddick et al. 2000).

Thus, we have multiple reasons to think that human cognition can, and does, take multiple forms. And while I think we do have good reasons to accept that a significant amount of human thought is indeed sententially structured (Camp 2015), it is very much an open possibility that this reflects the influence of language as a biologically-endowed and overlearned communicative medium, rather than the other way around.

3.2 Language as Expressing Thought: Diversities in Linguistic Function

The central lesson of §3.1 was that we have good reasons to reject Devitt's claim that human thought in general takes a sentential form, akin to a predicate calculus. Suppose, though, that we do accept that assumption. Shifting from thought to language, the next big move in Devitt's argument for the priority of thought is the claim that language takes the form it does *because* it expresses thought, and in particular because the structure of language reflects the structure of thought. (As Dummett (1989: 197) puts it, "a fully explicit verbal expression is the only vehicle whose structure must reflect the structure of the thought.") In this section, I argue that while expressing thought is indeed one central thing that language does, it also has other important functions.

Devitt doesn't offer much detail about what it means for language to express thoughts. 'Thoughts', for him, are "mental states with meanings" (142): "propositional attitudes, mental states like beliefs, desires, hopes, and wondering whethers" (125). 'Expressing' is a matter of "convey[ing] a 'message'" by "uttering a sentence of the language to express a thought with the meaning that the sentence has in that language" (127), where that 'meaning' is determined by public conventions for use (132). So his overall picture is that language expresses thought by combining words whose conventional meanings match the concepts that constitute the propositional attitude expressed, in a structure that mirrors the structure of that propositional attitude.

An initial, somewhat ancillary worry focuses on the role Devitt assigns to conventional meaning here. He needs to do this to establish his overall negative conclusion, that "the primary concern in linguistics should not be with idiolects but with linguistic expressions that share meanings in idiolects" (12). However, the move from the claim that language expresses thought to the conclusion that linguistic meaning is conventional is too quick. Even many theorists who embrace a conception of language as a communicative device and who accept that linguistics should study "shared meanings" reject the conventionality of meaning. In particular, where Devitt simply assumes that the conventional meaning of an uttered sentence "often" matches the thought that the speaker intends to express with it (132), 'radical contextualists' like Recanati (2004) argue that many if not all utterances involve significant context-local influences that are not triggered by elements within the sentences uttered; and they often conclude, with Davidson (1986), that any appeal to convention is an irrelevant chimera. I agree with contextualists that most utterances involve context-local influences on communicated meaning. But I also agree with Devitt that conventional meaning plays an important role in the theoretical explanation of linguistic communication (Camp 2016). However, establishing this latter conclusion requires closer attention to the dynamics of ordinary discourse than Devitt provides; and I am suspicious of the claim that language as such, shorn of pragmatic modulation and amplification, typically expresses complete thoughts that speakers would be willing to endorse, let alone care to communicate.

Let's put general worries about the existence and role of linguistic convention aside, though, and focus just on what conventions for direct and literal use might actually be like. Crucially, linguistic terms and constructions implement a variety of conventional functions, not all of which can be smoothly assimilated under the rubric of 'expressing thought'. One key source of complexity centers around illocutionary force, which lies at the intersection of syntax, semantics, and pragmatics. Standard linguistic theories now reject the traditional 'marker' model, on which different sentence types conventionally mark distinct forces applied to a common propositional core. Instead, declarative sentences are standardly taken to denote propositions, while questions denote partitions of possible worlds and imperatives denote goals or properties that are indexed to the addressee (see e.g. Roberts 1996/2012, 2018). None of these denoted objects are themselves "thoughts," in Devitt's intuitive sense; rather, utterances of sentences of these three syntactic types conventionally function to undertake the speech acts of assertion, interrogation, and direction, and those speech acts have the conventional effect of altering the discourse in a certain way, for instance by adding the denoted object to the common ground (Stalnaker 1978).

There are obviously often intimate causal and normative connections between those speech acts and speakers' psychological attitudes, especially beliefs and intentions. But the claim that those acts function in their entirety only to express those attitudes is implausible. To take the most straightforward case, of assertion, the view that asserting is the expression of belief (e.g. Bach and Harnish 1979) is at a minimum incomplete, because it fails to distinguish assertion from other modes of linguistic belief-expression such as presupposition and implicature. More seriously, it also fails to allow for assertions that do not even purport to be grounded in belief, such as bald-faced lies (Sorensen 2007), 'selfless' assertions (Lackey 2007), and suppositions and other contributions for sake of the current conversation (Stalnaker 1978). Thus, rather than functioning exclusively or ultimately to express psychological attitudes, it is more plausible that utterances of sentences of the relevant syntactic types have the operative conventional function of *doing* something in discourse, either by altering the structure and contents of the common ground (Stalnaker 1978; Roberts 1996/2012; Murray and Starr 2020), and/or by undertaking a public commitment to produce other, suitably related speech acts in appropriate circumstances (Brandom 1983; MacFarlane 2011), where this action may be linked to, but is not identical with, the expression of belief.

Further, sentential type is not the only morpho-syntactic element with the conventional function of indicating or modulating illocutionary force. Other "illocutionary-force-indicating devices" include performative verbs like 'I apologize'; appositive clauses like 'as I claim' (Searle and Vanderveken 1985; Green 2007); and adverbials like 'frankly' or 'admittedly' (Bach 1999). Related terms and constructions, such as 'while', 'but', 'therefore', 'actually', and 'all in all', function to regulate the structure of discourse, by indicating and modulating relations among utterances of distinct sentences so that they form a coherent whole (Asher and Lascarides 2003; Kehler 2004). (Indeed, Clark and Fox Tree 2002 argue that apparent disfluencies in spontaneous speech, like 'um' and 'uh', are conventional English words which function to implicate that the speaker is initiating a major or minor delay in speaking.) Discourse particles like 'Man,' (McCready 2008) and 'like' (Siegel 2002) function to intensify or hedge the semantic contents of their focal terms. Evidentials, like 'I heard', 'I saw' or 'as they told me', function to indicate the evidential status of the illocutionary act's core at-issue content (Murray 2014). And a wide range of performative terms function to regulate social dynamics, to display emotional attitudes, and to mark social affiliation: thus, expressives like 'damn' express the speaker's emotional state (Potts 2007); honorifics like French 'vous' implicate a social or attitudinal relation between speaker and addressee (McCready 2010); and slurs like 'kike' undertake a commitment to the appropriateness of a derogating attitude toward the target group (Camp 2013).

The point of mentioning all of these classes of terms and constructions is not that they are disconnected from psychology; on the contrary, they are some of the most nuanced linguistic tools we have for coordinating minds and behaviors. In a suitably capacious sense of the words ‘express’ and ‘thought’, on which ‘expression’ is the outward showing of an inner state and ‘thought’ includes any kind of inner state (e.g. Green 2007), we can presumably identify psychological states correlated with each of these term-types – although an analysis of their meanings as simply the expression of those states will often encounter challenges like those facing a purely expressive account of assertion, and the relation between term and state will often not be aptly characterized on the familiar model of Gricean non-natural meaning in the form of reflexive intention-recognition.

Rather, the key point is that the conventional functions of many of these terms and constructions do not arise simply from the need to exteriorize inner thoughts, but rather in significant part from the need to manage distinctively social dynamics. As such, they are not functions we would expect to be manifested in a Mentalese prior to and independent of linguistic communication, which are then merely exteriorized by speech. Specifically, many of these terms and constructions function to provide higher-order comments on primary, first-order speech acts (Neale 1999), in a way that renders correlative mental states at least partly dependent on that lower-order speech act. More generally, the presence of such terms and constructions reflects the status of natural language as a deeply social construction. *Pace* Chomsky, language is not just a biological mechanism for constructing complex meaningful strings; but neither is it just the public avatar of a commonly instantiated but ultimately essentially individual Language of Thought. To make sense of, or even notice, these linguistic phenomena, we need to approach language on its own terms: as a shared tool for achieving various species of coordination beyond just representation.

In addition to the type of meaning they have, these terms and constructions are also theoretically interesting because of the way they interact with the rest of the linguistic machinery. Specifically, their conventional contributions are typically rhetorically peripheral rather than ‘at-issue’ (Horn 2014), so that they are not the natural target of direct anaphoric agreement and denial, like ‘That’s true’ or ‘I disagree’. Many of them resist syntactic embedding under more complex constructions like negation and conditionalization. And when they do embed, they are typically interpreted as ‘projecting out’ of those constructions, so that the speaker is interpreted as undertaking a straightforward, unmodified commitment to their associated contribution even as the ‘core’ content is negated, conditionalized, etc.

Thus, much as in the case of cognition, when we examine how language actually works, we find significant functional diversity, which is reflected in significant semantic and syntactic diversity. A parallel response here, as in the case of cognition, is to ‘go modular’, by segregating peripheral contributions from the ‘core’ compositional machinery (Potts 2005), leaving the latter free to be analyzed in ways that more closely approximate the traditional model of a predicate calculus. Such a segregationist model is formally attractive. But it cannot accommodate the fact that the resistance of such terms and constructions to embedding is in many cases merely a default status, which can be overridden by syntactic and pragmatic factors in particular cases. In particular, imperatives and interrogatives, expressives, and slurs can also sometimes receive embedded interpretations, when their contribution is rendered at-issue relative to the larger discourse structure (Siegel 2006; Simons et al. 2010; Camp 2018b). Given this, an empirically adequate linguistic theory needs not only to acknowledge and explain these 'peripheral' constructions in isolation, but also to analyze the familiar core logical machinery, including negation, disjunction and conditionalization, in a way that reflects the diversity of uses to which that machinery can be put in natural language, which, as we've seen, includes operating on non-representational semantic values.

'Dynamic' approaches to linguistic meaning, which analyze the meanings of words in terms of their compositional contributions to the 'context change potentials' or 'update instructions' associated with sentences in which they occur (Heim 1983; Groenendijk and Stokhof 1991; Veltman 1996), appear to be especially well-equipped to provide the requisite flexibility in a theoretically motivated way. They may even support a resuscitated version of the thesis that language expresses thought (Charlow 2015). But they are also likely to have radical consequences, both for the analysis of logical machinery in natural language and also potentially for our theoretical understanding of the cognitive states expressed. These are consequences that need to be articulated and assessed in detail. But they are consequencees that many more traditional philosophers, including especially Devitt, are likely to want to resist.

3.3 Universal Grammar and the Psychology of Language Processing

3.3.1 UG-Violating Strings

The final major step in Devitt's argument for the elimination of the language faculty, after establishing that thought has sentential form and that language expresses thought, is the claim that linguistic competence, and hence the language faculty, is no more than "the ability that matches token sounds and thoughts for meaning" (129). Once we shift the entirety of the explanatory burden onto cognition, the argument goes, and accept that language merely transduces the contents and structure of thoughts, there is "little or nothing" left for the language faculty to do; the little work that does remain can be performed by "fairly brute-causal associationist processes."

In §3.1, I rejected the claim that thought itself is universally sentential; and in §3.2 I argued against a monolithic model of language as expressing "beliefs, desires, hopes, and wondering whethers." So we already have significant reasons to doubt that language universally functions to implement, in publically observable form, a structure that is antecedently instantiated by a univocally structured mental state like belief. In this section, I assess the claim that the processes that govern distinctively linguistic processing are merely associationist, with all or most observed constraints on grammatical structure arising at the level of the thoughts expressed.

If UG really constituted the rules of thought itself, this would seem to entail that ordinary people are unable to generate or classify, let alone comprehend, UG-violating strings. However, it appears that we do regularly make sense of UG-violating strings, at a minimum in the course of correcting other speakers' disfluencies. Devitt acknowledges that we can indeed make sense of UG-violating strings, but suggests that we do so, "not by carrying its syntax into our thought but by translating it into a thought with a syntax that is like a sentence in our language" (151) – where this process of 'translation' is presumably also achieved by "fairly brute-causal association."

However, empirical evidence does not support such an 'associationist translation' view of the interpretation of UG-violating strings. Typical humans can learn to construct and classify strings using both UG-conforming and UG-violating rules. Specifically, although they may have difficulty extrapolating UG-violating rules from unstructured data (Smith et al. 1993), they can learn to deploy rules that violate UG in virtue of utilizing "rigid" linear distance between words, when those rules are stated explicitly (Musso et al. 2003). At the same time, though, UG-conforming and -violating rules are not on a cognitive par. In particular, they are implemented in distinct neural regions; specifically, the regions within Broca's area that are also activated during ordinary natural language processing are only activated when deploying artificial syntactic rules that conform to UG (Embick et al. 2000; Moro et al. 2001). Indeed, Musso et al. (2003) found that Broca's area progressively *disengaged* as subjects learned the UG-violating grammar, without any other distinctive pattern of brain activity being manifested.

The first, most straightforward implication of these findings for the current discussion is that the overall cognitive abilities of normal subjects can underwrite at least some UG-violating 'thought', in the sense of rule-governed classification, without any translation into or activation of UG-conforming structures. Second, however, the crucial mechanism that does process UG-conforming strings is not part of 'thought' in Devitt's favored sense, of "mental states with meanings." In particular, while some of the relevant experiments contrasted real and artificial rules for languages like Italian and Japanese, others contrasted rules for classifying meaningless symbols as 'agreeing' with respect to patterns involving color and size (Tettamanti et al. 2009). Thus, the distinctive quality that activates the neural areas especially associated with UG seems to be a purely abstract, structural one. More specifically, the crucial feature is whether the rule for 'agreement' among elements is "non-rigid": concerning structural relations among features that are neutral with respect to position, in contrast to "rigid" rules about linear distance between elements. Further, these same neural areas also appear to subserve cognitive processing for other domains that involve the same sort of hierarchical structure, such as music (Patel 2003) and planning complex actions (Koechlin and Jubault 2006).

Devitt might take this last fact, that these neural regions are activated for domains other than language, to support his claim that the relevant structures are processed at a level of 'thought' rather than language, and so that there is "little or nothing to the language faculty" after all. However, the sense in which this holds is at best terminological. Chomsky and colleagues take the "faculty of language" in the relevant, "narrow" sense ('FLN') to be a mechanism that generates complex internal representations by recursion, and then "pairs sound and meaning" by interfacing with the "sensory-motor" and "conceptual-intentional" systems (Hauser et al. 2002: 1571). The fact that the core mechanism is also utilized for other cognitive purposes does not undermine the existence of biologically innate, distinctively linguistic package of a hierarchical recursive syntax plus phonology and semantics. Further, if that core recursive mechanism is what generates the complex matching structures that are utilized by both the articulatory and conceptual systems, then this mechanism is what implements the mapping from sounds to meanings that Devitt himself calls the 'language faculty' – but in a way that is the very opposite of a "brute-causal associative process."

3.3.2 UG-Conforming Complexities

Classificatory tussling aside, and even ignoring all the varieties of pragmatic and 'peripheral' conventional aspects of meaning cited in §3.2, there remains much more to natural language than the pure recursive operation of predication or Merge. 'Universal Grammar' encompasses all the initial constraints and operations required to derive the full complexities of adult linguistic competence. Devitt is committed to the claim that these linguistic complexities, like the more obviously systematic operations of predication and functional application, are to be explained as manifestations of the "innate structure rules of thought," rather than as arising from language itself. We saw in §3.1 that there are good reasons to doubt that human thought innately has any one universal format. But even if we focus just on the sorts of thoughts that are most plausibly canonically expressed in language – "beliefs, desires, hopes, and wonderings whether" – it is still implausible that those aspects of UG that aren't directly derivable from a Merge-like core operation of hierarchical recursion are indeed 'largely' derived from innate rules of thought alone. Rather, they are more plausibly generated by distinctive features of the structure of natural language instead.

This is shown first, by various types of constraints on well-formedness that are not plausibly motivated by anything about the thoughts expressed, but that also don't vary in a conventional way across languages. Ludlow (2009) invokes the case of filler-gap constructions to make this point, using the following minimal pair:

(1) Who(m) did John hear that Fred said that Bill hit?
(2) # Who(m) did John hear the story that Bill hit?

(1) is a perfectly well-formed question; and (2) is lexically and structurally highly similar to (1). But while the question that a speaker of (2) might be *trying* to ask is perfectly comprehensible – who was the subject of the Bill-hitting story that John heard – the string itself is irredeemably ill-formed. Explaining the difference between (1) and (2), and the vast range of analogous minimal pairs, has motivated

linguists to posit highly complex unpronounced syntactic structures and transformation constraints that are specific to natural language; in the case of (1) and (2), these are so-called ‘island constraints’ on movement across phrases.

Where island constraints are almost purely syntactic, other classic cases demonstrate syntactic constraints that are responsive to the meanings of their constituent elements, but again not in a way that is derivable from constraints on thought itself. Thus, NPIs or negative polarity items are expressions like ‘any’ and ‘lift a finger’ that can only appear in certain environments. In the following minimal pairs,

(3a) I’ve got some money.	(4a) # I don’t have some money.
(3b) # I’ve got any money.	(4b) I don’t have any money.

the co-numbered pairs of sentences seem to express equivalent thoughts, but only one is well-formed. Negation, as in (4a) and (4b), is the most obvious NPI-licensing environment, and many licensers are like negation in being downward entailing (Ladusaw 1979). But some contexts which are not straightforwardly downward entailing, such as antecedents of conditionals, questions, and expressions of surprise, can also license NPIs; and some licensing appears to depend on more fully pragmatic factors (Krifka 1995; von Fintel 1999; Israel 2011).

Finally, some constraints on well-formedness appear not to be motivated by any systematic property at all, syntactic, semantic or pragmatic. Johnson (2004) cites the contrast between ‘put’ and ‘stow’ as illustration:

(5a) John put his gear down.	(6a) ∗John put his gear.
(5b) ∗John stowed his gear down.	(6b) John stowed his gear.

Here again, the co-numbered pairs seem to express equivalent thoughts, but only one is well-formed.

Devitt could resist the conclusion that these last two classes of cases reveal constraints on well-formedness that are not derivable from thought by insisting that in fact, all of the minimal pairs in (3) through (6) actually express distinct thoughts, and so that their different linguistic statuses are inherited from thought after all. However, going this route requires embracing a notion of ‘thought’ that is so fine-grained as to verge on equivalence to the inner voicing of sentences. This would stipulatively rule out the intuitively plausible and linguistically significant possibility that distinct sentence-types can express the same thought. More importantly, it would threaten to undermine any independent grip on what a ‘thought’ is, or else to suggest that Devitt’s ‘thoughts’ are ultimately individuated by the public-language sentences that express them rather than the other way around.

A second route to demonstrating the irreducibility of language to thought appeals not to constraints on well-formedness, but to grammatical rules that concern meaning in a way that is difficult to generate from Mentalese alone. Perhaps the clearest cases of this involve what Pietroski and Crain (2012) call “unambiguities”: constraints on mappings from forms to meaning that can’t be derived from either the bare sentential structure or from the thoughts themselves. Thus, the sentence

(7) John is eager to please

appears to have two implicit slots for pronoun assignment, and there appear to be at least two distinct, perfectly coherent thoughts that would result from filling in those slots:

(7a) JOHN$_1$ IS EAGER THAT HE$_1$ PLEASE US
(7b) JOHN$_1$ IS EAGER THAT WE PLEASE HIM$_1$.

But only (7a) is available as an interpretation of (7) (Pietroski and Crain 2012).

Finally, in addition to it being unclear how to generate or even define all the constraints exhibited by natural languages from the hypothesis that they reflect the "innate structure rules of thought," there is also empirical evidence that UG can govern linguistic production in the absence of commensurately complex thought. Advocates of linguistic modularity (e.g. Pinker 1999) often cite people with Williams Syndrome in this context, because they display differentially robust linguistic abilities – specifically, implicit grasp of syntactic principles like c-command, scope, and binding – against a strongly impaired general cognitive background. Although this interpretation has been resisted by 'neuroconstructivists' (e.g. Thomas and Karmiloff-Smith 2005), it is increasingly well-established that adults with WS do process syntax and morphology using the same mechanisms as typical subjects (Brock 2007). Differences in their performance on grammatical tasks from normal subjects are more plausibly attributed to limitations in handling complexity – that is, from limitations arising from extra-grammatical cognitive resources like working memory, in a way that parallels limitations exhibited by typically-developing children matched to WS adults for overall cognitive function (Musolino and Landau 2012).

Devitt acknowledges that there are people who utter complex grammatical sentences despite being highly cognitively impaired. But he argues that in order for them to count against his priority claim, "we would need to establish *both* (a), that the savants [e.g. people with WS] cannot think thoughts with certain meanings, *and* (b), that sentences out of their mouths really have those meanings" (165). He argues against (a) by suggesting that such people might just be bad at reasoning with the thoughts they do have, and against (b) by saying that if (a) were true, then we would thereby be forced to conclude that the sounds coming out of their mouths were "mere noise."

Against this latter claim, note first that even if the words uttered by people like those with Williams Syndrome were "mere noise," this would still demonstrate the existence of a psychological mechanism for generating specifically grammatical complexity, independent from commensurately complex thought – a mechanism, that is, very close to the sort of language faculty under dispute. But second, it is not plausible that people with WS are in fact just making "noise". For instance, Musolino and Landau (2012) probed grammatical knowledge in people with WS by asking subjects to match lexically matched but syntactically distinct sentences, such as

(8a) The cat who meows will not be given a fish or milk t6.1
(8b) The cat who does not meow will be given a fish or milk t6.2

to animated vignettes – a task that they were indeed able to perform, and that requires assigning truth-conditions. More generally, people with WS are theoretically notable because their utterances are typically not just syntactically well-formed but also semantically coherent, both internally and in relation to one another, and at least basically pragmatically appropriate. For instance, they are often good at spontaneously generating coherent and engaging narratives from pictures; thus, Rossen et al. (1996: 367) cite the following spontaneous description by a 16-year-old Williams Syndrome subject, Crystal, about her future aspirations:

> You are looking at a professional bookwriter. My books will be filled with drama, action, and excitement. And everyone will want to read them … I am going to write books, page after page, stack after stack. I'm going to start on Monday.

Meanwhile, this same patient "fails all Piagetian seriation and conservation tasks (milestones normally attained by in the age range of 7 to 9 years); has reading, writing, and math skills comparable to those of a first- or second-grader, and requires a babysitter for supervision." It's hard to make sense of just what is going on in the mind of a person like this. But a flat-footed appeal to general "stupidity" and "failure in practical reasoning" (Devitt 2006: 165) won't be satisfying unless it can explain how such subjects do produce such complex, sophisticated, and specifically verbal behavior.

3.4 Priorities, Sufficiencies, and Speculations

Establishing the strong conclusion that there is little or nothing to the language faculty requires commensurately strong assumptions: first, that thought has a sentential format, specifically one that conforms to UG; second, that language expresses thought, directly and in virtue of conventions for use; and third, that the transparent expression of thought leaves no further work for psychological mechanisms distinctively associated with language beyond associating surface forms with complex, independently meaningful mental representations. These assumptions can seem highly plausible, even inevitable, when formulated within the context of a model of both thought and language as monolithic instances of a "rational calculus." This picture has held philosophers captive since before the founding of analytic philosophy. But when we look at how both thought and language work, we find that the actual contours of human cognition and natural languages are more complex, and less systematic, than the strong argument requires. The broad terms 'thought' and 'language' encompass a range of importantly diverse functions, each implemented by a suite of intimately interacting but at least partially dissociable abilities and mechanisms.

At the same time, those strong assumptions appear plausible, and have dominated analytic philosophy, for a reason. It is not merely wishful thinking that people are rational animals: humans really do, at least sometimes, engage in logical reasoning – though we also often blithely ignore or spectacularly fail to follow the laws of logic. Likewise, even if compositionality is better seen as a regulative methodological principle than a truistic observation about natural languages (Szabo 2012), it has still proven to be an enormously productive principle, with apparently recalcitrant constructions, such as epistemic modals, receiving compelling analysis given more sophisticated formal tools.

More specifically, we have seen that both thought and language do manifest a common core that does approximate to the predicate calculus, not just in the general sense of being highly systematic but in the narrower one of employing a hierarchical recursive combinatorial principle. Moreover, this core appears to be implemented by a common neural mechanism, which plausibly plays a central role in making both distinctively human thought and talk possible.

As I noted in §3.1, one crucial, much-discussed feature of distinctively human cognition is expressive generality: the ability to think about a wide range of contents of indefinite complexity. Hierarchical recursive syntactic structure plays a key role in underwriting expressive generality. But it does not suffice on its own. In addition, a representational system must also employ a semantic principle that is arbitrary enough to represent a wide range of types of values. More relevantly, its combinatorial principle must also be highly neutral, or else the system as a whole will only have the capacity to represent the sorts of values that can be meaningfully related by that principle.

At least some primates appear to possess hierarchically-structured but domain-limited cognitive abilities. In particular, there is good behavioral evidence that baboons represent hierarchical, recursively structured relations of social dominance (Cheney and Seyfarth 2007); but they don't seem to think similarly complex thoughts about other domains. This suggests that they may employ something like a branching tree structure with a dedicated significance, much like a phylogenetic tree (Camp 2009a). Conversely, representational systems can also achieve a high degree of expressive generality without employing a branching tree structure, as in Venn diagrams. Given these dissociations, hierarchical recursive syntax should not be viewed as the necessary and sufficient essence of a distinctively powerful, and therefore distinctively human, capacity for thought.

If we want to speculate about the evolutionary origins of human thought and language, it is at least *prima facie* plausible that evolutionary pressures for a neutrally interpreted recursive tree structure stemmed as much from a need to communicate as from the need to represent hierarchically complex contents. Perhaps our pre-human ancestors possessed a dedicated module for social cognition plus a syntactically simple signaling system, much as baboons do. Syntactic complexity is only advantageous once the range of potential signals exceeds a certain threshold (Nowak et al. 2000); and in principle, any sort of combinatorial system could satisfy the need to generate an indefinitely large number of signals from a restricted base. But the communicative media plausibly most reliably accessible to those pre-human ancestors – sounds and gestures – are saliently distinguished by having a unidimensional, specifically temporal structure. An operation that merges multiple branches into single nodes, which are themselves hierarchically ordered, permits the representation of complex contents in a linear order. Thus, the exaptation of the basic syntactic structure of the social dominance module, by means of abstracting away from the semantic significance of branching trees within that module, would permit communication of a wide range of contents in a single, readily available, and flexibly implementable medium.

But even granting this highly speculative step, the bare potential to represent an indefinitely wide range of complex contents would be practically irrelevant without a robust actual ability to form, connect, and transform a wide range of representations from within any given context – that is, without a significant degree of stimulus-independence (Camp 2009b). Such active cognitive flexibility is a crucial ingredient in instrumental reasoning. But it too can be implemented in a variety of ways, including by imagistic simulation, and it is not restricted to humans (Camp and Shupe 2017). And here again, an argument can be made that the distinctively communicative use of language facilitated (and continues to facilitate) the development of imaginative flexibility, by giving thinkers a means to simulate what someone else, or they themselves, would say about a given problem or possibility (Carruthers 1998; McGeer and Pettit 2002).

Again, this is highly speculative. But a synthetic view along these general lines requires far fewer strong assumptions and big leaps than either pure Chomskian structuralism or pure Devittian conceptualism. In lieu of the picture that has been implicitly embraced by many philosophers throughout the twentieth century, on which monolithic thought has sweeping priority over equally monolithic language, we should instead embrace a model on which both human cognition and natural language involve many distinct, potentially dissociable abilities functioning together in a way that is significantly but not entirely integrated. Both thought and talk do involve a systematic predicative core. But in neither case is there a clean division between this core and the rest of cognition or language. Nor is there good reason to privilege that core as what makes us distinctively human, whether by nature or by enculturation.

References

Allwein, G., and J. Barwise, eds. 1996. *Logical reasoning with diagrams*. Oxford: Oxford University Press.

Asher, N., and A. Lascarides. 2003. *Logics of conversation*. Cambridge: Cambridge University Press.

Bach, K. 1999. The myth of conventional implicature. *Linguistics and Philosophy* 22: 367–421.

Bach, K., and R. Harnish. 1979. *Linguistic communication and speech acts*. Cambridge, MA: MIT Press.

Bermudez, J.L. 2003. *Thinking without words*. Oxford: Oxford University Press.

Brandom, R. 1983. Asserting. *Noûs* 17 (4): 637–650.

Brock, J. 2007. Language abilities in Williams syndrome: A critical review. *Development and Psychopathology* 19: 97–127.
Camp, E. 2007. Thinking with maps. *Philosophical Perspectives* 21 (1): 145–182.
———. 2009a. A language of baboon thought? In *The philosophy of animal minds*, ed. R. Lurz, 108–127. Cambridge: Cambridge University Press.
———. 2009b. Putting thoughts to work: Concepts, systematicity, and stimulus-independence. *Philosophy and Phenomenological Research* 78 (2): 275–311.
———. 2013. Slurring perspectives. *Analytic Philosophy* 54 (3): 330–349.
———. 2015. Logical concepts and associative characterizations. In *The conceptual mind: New directions in the study of concepts*, ed. E. Margolis and S. Laurence, 591–621. Cambridge, MA: MIT Press.
———. 2016. Conventions' revenge: Davidson, derangement, and dormativity. *Inquiry* 59 (1): 113–138.
———. 2018a. Why cartography is not propositional. In *Non-propositional intentionality*, ed. A. Grzankowski and M. Montague, 19–45. Oxford: Oxford University Press.
———. 2018b. Slurs as dual-act expressions. In *Bad words*, ed. D. Sosa, 29–59. Oxford: Oxford University Press.
Camp, E., and E. Shupe. 2017. Instrumental reasoning in non-human animals. In *The Routledge handbook of philosophy and animal minds*, ed. J. Beck and K. Andrews, 100–108. London: Routledge.
Carruthers, P. 1998. Thinking in language? Evolution and a modularist possibility. In *Language and thought*, ed. P. Carruthers and J. Boucher, 94–119. Cambridge: Cambridge University Press.
———. 2003. On Fodor's problem. *Mind and Language* 18 (5): 502–523.
Casati, R., and A. Varzi. 1999. *Parts and places: The structures of spatial representation*. Cambridge, MA: MIT Press.
Charlow, N. 2015. Prospects for an expressivist theory of meaning. *Philosophers' Imprint* 15: 1–43.
Cheney, D.L., and R.M. Seyfarth. 2007. *Baboon metaphysics: The evolution of a social mind*. Chicago: University of Chicago Press.
Clark, H.H., and J.E. Fox Tree. 2002. Using 'uh' and 'um' in spontaneous speaking. *Cognition* 84: 73–111.
Davidson, D. 1986. A nice derangement of epitaphs. In *Truth and interpretation: Perspectives on the philosophy of Donald Davidson*, ed. E. Lepore, 433–446. New York: Blackwell.
De Toffoli, S. 2017. 'Chasing' the diagram: The use of visualizations in algebraic reasoning. *The Review of Symbolic Logic* 10 (1): 158–186.
Devitt, M. 2006. *Ignorance of language*. Oxford: Clarendon Press.
Dummett, M. 1989. Language and communication. In *Reflections on Chomsky*, ed. A. George, 192–212. Oxford: Oxford University Press.
———. 1994. *Origins of analytical philosophy*. Cambridge, MA: Harvard University Press.
Elwert, F. 2013. Graphical causal models. In *Handbook of causal analysis for social research*, ed. S.L. Morgan, 245–273. New York: Springer.
Embick, D., A. Marantz, Y. Miyashita, W. O'Neil, and K.L. Sakai. 2000. A syntactic specialization for Broca's area. *PNAS* 97 (11): 6150–6154.
Evans, J.St.B.T., and D.E. Over. 1996. *Rationality and reasoning*. Hove: Psychology Press.
Fiddick, L., L. Cosmides, and J. Tooby. 2000. No interpretation without representation: The role of domain-specific representations and inferences in the Wason selection task. *Cognition* 77: 1–79.
Fodor, J. 1987. Why there still has to be a language of thought. In *Psychosemantics: The problem of meaning in the philosophy of mind*, ed. J. Fodor, 135–154. Cambridge, MA: MIT Press.
———. 2007. The revenge of the given. In *Contemporary debates in philosophy of mind*, ed. B.P. McLaughlin and J.D. Cohen, 105–116. Oxford: Blackwell.
Fodor, J., and Z. Pylyshyn. 1988. Connectionism and the cognitive architecture of mind. *Cognition* 28: 3–71.

Franconeri, S.L., J.M. Scimeca, J.C. Roth, S.A. Helseth, and L.E. Kahn. 2012. Flexible visual processing of spatial relationships. *Cognition* 122: 210–227.

Giardino, V., and G. Greenberg. 2014. Introduction: Varieties of iconicity. *Review of Philosophical Psychology* 6 (1): 1–25.

Gopnik, A., C. Glymour, D.M. Sobel, L.E. Schulz, T. Kushnir, and D. Danks. 2004. A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review* 111 (1): 3–32.

Green, M.S. 2007. *Self-expression*. Oxford: Oxford University Press.

Grice, H.P. 1975. Logic and conversation. In *Syntax and semantics volume 3: Speech acts*, ed. P. Cole and J.L. Morgan, 41–58. New York: Academic Press.

Groenendijk, J., and M. Stokhof. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14 (1): 39–100.

Hauser, M.D., N. Chomsky, and W.T. Fitch. 2002. The language faculty: What is it, who has it, and how did it evolve? *Science* 298: 1569–1579.

Heim, I. 1983. File change semantics and the familiarity theory of definiteness. In *Meaning, use and interpretation of language*, ed. R. Bäuerle, C. Schwarze, and A. von Stechow, 164–189. Berlin: De Gruyter.

Horn, L. 2014. Information structure and the landscape of (non-)at-issue meaning. In *The Oxford handbook of information structure*, ed. C. Féry and S. Ishihara, 108–128. Oxford: Oxford University Press.

Israel, M. 2011. *The Grammar of polarity: Pragmatics, sensitivity, and the logic of scales*. Cambridge: Cambridge University Press.

Johnson, K. 2004. On the systematicity of language and thought. *The Journal of Philosophy* 101 (3): 111–139.

Johnson-Laird, P. 2005. Mental models and thought. In *The Cambridge handbook of thinking and reasoning*, ed. K.J. Holyoak and R.G. Morrison, 185–208. Cambridge: Cambridge University Press.

Kehler, A. 2004. Discourse coherence. In *Handbook of pragmatics*, ed. L.R. Horn and G. Ward, 241–265. Oxford: Basil Blackwell.

Koechlin, E., and T. Jubault. 2006. Broca's area and the hierarchical organization of human behavior. *Neuron* 50: 963–974.

Krifka, M. 1995. The semantics and pragmatics of polarity items. *Linguistic Analysis* 25: 209–257.

Lackey, J. 2007. Norms of assertion. *Noûs* 41 (4): 594–626.

Ladusaw, W. A. 1979. *Polarity sensitivity as inherent scope relations*. PhD dissertation, University of Texas, Austin.

Lemon, O., and I. Pratt. 1997. Spatial logic and the complexity of diagrammatic reasoning. *Machine Graphics and Vision* 6 (1): 89–108.

Locke, J. 1689. *An essay concerning human understanding*. London: Thomas Bassett.

Ludlow, P. 2009. Review of Devitt's i*gnorance of language*. *Philosophical Review* 118 (3): 393–402.

MacEachren, A. 2004. *How maps work: Representation, visualization, and design*. New York: Guilford Press.

MacFarlane, J. 2011. What is assertion? In *Assertion*, ed. J. Brown and H. Cappelen, 79–96. Oxford: Oxford University Press.

Marchette, S., J. Ryan, and R. Epstein. 2017. Schematic representations of local environmental space guide goal-directed navigation. *Cognition* 158: 68–80.

McCready, E. 2008. What man does. *Linguistics and Philosophy* 31 (6): 671–724.

———. 2010. Varieties of conventional implicature. *Semantics and Pragmatics* 3 (8): 1–57.

McGeer, V., and P. Pettit. 2002. The self-regulating mind. *Language and Communication* 22: 281–299.

Morgan, L., S. MacEvoy, G. Aguirre, and R. Epstein. 2011. Distances between real-world locations are represented in the human hippocampus. *The Journal of Neuroscience* 31 (4): 1238–1245.

Moro, A., M. Tettamanti, D. Perani, C. Donati, S. Cappa, and F. Fazio. 2001. Syntax and the brain: Disentangling grammar by selective anomalies. *NeuroImage* 13: 110–118.
Murray, S. 2014. Varieties of update. *Semantics and Pragmatics* 7 (2): 1–53.
Murray, S., and W. Starr. 2020. The structure of communicative acts. *Linguistics and Philosophy*. https://doi.org/10.1007/s10988-019-09289-0.
Musolino, J., and B. Landau. 2012. Genes, language, and the nature of scientific explanations: The case of Williams syndrome. *Cognitive Neuropsychology* 29 (1–2): 123–148.
Musso, M., A. Moro, V. Glauche, M. Rijintjes, J. Reichenbach, C. Büchel, and C. Weiller. 2003. Broca's area and the language instinct. *Nature Neuroscience* 6 (7): 774–781.
Neale, S. 1999. Coloring and composition. In *Philosophy and linguistics*, ed. K. Murasugi and R. Stainton, 35–82. Boulder, CO: Westview Press.
Nowak, M.A., J.B. Plotkin, and V.A. Jansen. 2000. The evolution of syntactic communication. *Nature* 404: 495–498.
Patel, A.D. 2003. Language, music, syntax and the brain. *Nature Neuroscience* 6: 674–681.
Pearl, J. 2000. *Causality*. Cambridge: Cambridge University Press.
———. 2009. Causal inference in statistics: An overview. *Statistics Surveys* 3: 96–146.
Pietroski, P., and S. Crain. 2012. The language faculty. In *The Oxford handbook of philosophy of cognitive science*, ed. E. Margolis, R. Samuels, and S.P. Stich, 361–381. Oxford: Oxford University Press.
Pinker, S. 1999. *Words and rules: The ingredients of language*. New York: Basic Books.
Potts, C. 2005. *The logic of conventional implicature*. Cambridge, MA: MIT Press.
———. 2007. The expressive dimension. *Theoretical Linguistics* 33 (2): 165–198.
Pratt, I. 1993. Map semantics. In *Spatial information theory: A theoretical basis for GIS lecture notes in computer science*, ed. A.U. Frank and I. Campari, 77–91. Berlin: Springer.
Recanati, F. 2004. *Literal meaning*. Cambridge: Cambridge University Press.
Rescorla, M. 2009. Predication and cartographic representation. *Synthese* 169: 175–200.
Rey, G. 1995. A not 'merely empirical' argument for the language of thought. *Philosophical Perspectives* 9: 201–222.
Rice, C. 2011. Massive modularity, content integration, and language. *Philosophy of Science* 78 (5): 800–812.
Roberts, C. 1996/2012. Information structure in discourse: Toward an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5: 1–69.
———. 2018. Speech acts in discourse context. In *New work on speech acts*, ed. D. Fogal, D. Harris, and M. Moss, 317–359. Oxford: Oxford University Press.
Rossen, M.L., E.S. Klima, U. Bellugi, A. Bihrle, and W. Jones. 1996. Interaction between language and cognition: Evidence from Williams syndrome. In *Language, learning, and behavior disorders: Developmental, biological, and clinical perspectives*, ed. J.H. Beitchman, N. Cohen, M. Konstantareas, and R. Tannock, 367–392. New York: Cambridge University Press.
Schmidhuber, J. 2015. Deep learning in neural networks: An overview. *Neural Networks* 61: 85–117.
Searle, J., and D. Vanderveken. 1985. *Foundations of illocutionary logic*. Cambridge: Cambridge University Press.
Shin, S. 1994. *The logical status of diagrams*. Cambridge: Cambridge University Press.
Siegel, M. 2002. Like: The discourse particle and semantics. *Journal of Semantics* 19: 35–71.
———. 2006. Biscuit conditionals: Quantification over potential literal acts. *Linguistics and Philosophy* 29: 167–203.
Simons, M., D. Beaver, J. Tonhauser, and C. Roberts. 2010. What projects and why. *Proceedings of SALT* 20: 309–327.
Sloman, A. 1978. *The computer revolution in philosophy: Philosophy, science and models of mind*. Atlantic Highlands, NJ: Humanities Press.

Smith, N., I.-A. Tsimpli, and J. Ouhalla. 1993. Learning the impossible: The acquisition of possible and impossible languages by a polyglot savant. *Lingua* 91: 279–347.

Sorensen, R. 2007. Bald-faced lies! Lying without the intent to deceive. *Pacific Philosophical Quarterly* 88 (2): 251–264.

Stalnaker, R. 1978. Assertion. In *Syntax and semantics*, Vol. 9: Pragmatics, ed. P. Cole, 315–332. New York: New York Academic Press.

Szabo, Z. 2012. The case for compositionality. In *The Oxford handbook of compositionality*, ed. W. Hinzen, E. Machery, and M. Werning, 64–80. Oxford: Oxford University Press.

Tenenbaum, J., C. Kemp, T. Griffiths, and N. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331: 1279–1285.

Tettamanti, M., I. Rotondi, D. Perani, G. Scotti, F. Fazio, S.F. Cappa, and A. Moro. 2009. Syntax without language: Neurobiological evidence for cross-domain syntactic computations. *Cortex* 45 (7): 825–838.

Thomas, M.S.C., and A. Karmiloff-Smith. 2005. Can developmental disorders reveal the component parts of the language faculty? *Language Learning and Development* 1: 65–92.

Tufte, E. 1983. *The visual display of quantitative information*. Connecticut: Graphics Press.

Veltman, F. 1996. Defaults in update semantics. *Journal of Philosophical Logic 25*: 221–261.

von Fintel, K. 1999. NPI-licensing, Strawson-entailment, and context-dependency. *Journal of Semantics* 16: 97–148.

Wittgenstein, L. 1921/2001. *Tractatus logico-philosophicus (translated by D. Pears and B. McGuinness)*. New York: Routledge.