de Almeida, R. G. & Lepore, E. (2018). Semantics for a module. In R. G. de Almeida & L. R. Gleitman (Eds.) On Concepts, Modules, and Language. New York: Oxford University Press.

# Semantics for a Module

#### ROBERTO G. DE ALMEIDA AND ERNIE LEPORE

Modularity is a hypothesis about a nomological distinction between perceptual, input-driven computations and background knowledge. It hinges on the very nature of representations and processes computed by input systems—and, crucially, on what input systems deliver to higher cognition. Perceptual computations are said to be encapsulated or have only relative—principled—access to background knowledge in the course of running its default algorithms. Moreover, perceptual computations are supposed to deliver "shallow" representations of transduced inputs. This is where we begin, and this is pretty much where *The Modularity of Mind* (Fodor, 1983) left off: the theoretical and empirical research programs were—and still are—to determine the scope of perceptual computations and their degree of autonomy; and, more broadly, to search for the line that divides perception from cognition, hardwired computations from contingent and malleable operations.

Of course, *Modularity* was not only about the encapsulation of some psychological capacities. It advanced an epistemological thesis about the distinction between observation and inference in the acquisition of knowledge—or the fixation of belief. In the present chapter, we are concerned with the psychological, rather than the epistemological, thesis that *Modularity* advanced. We tie together two threads bearing on sentence representation and processing: one is that sentence perception is, to some extent, computationally encapsulated; and the other is that sentence *meaning* is, to some extent, context insensitive, or at least its sensitivity is rule-governed.

These threads come together in the following way: we argue that the output of sentence encapsulation is a minimally and highly constrained, context-sensitive

propositional representation of the sentence, built up from sentence constituents. Compatible with the original *Modularity* story, we argue that the output of sentence perception is thus a "shallow" representation—though it is *semantic*. The empirical cases we discuss bear on alleged cases of sentence indeterminacy, and how such cases might (a) be assigned (shallow) semantic representations, (b) interact with context in highly regulated ways, and (c) whether and, if so, how they can be enriched. In the course of our discussion, we will advance and defend a proposal for a semantic level of representation that serves as output of the module and as input to other systems of interpretation, arguing for a form of modularity or encapsulation that is minimally context sensitive provided that the information from context—whatever it may be—is itself determined nomologically, namely, by linguistic principles.

## THE CASE FOR MODULARITY

The Modularity of Mind (Fodor, 1983) raised fundamental questions about the architecture of perception and cognition, and, in particular, about linguistic and visual computations—whether they are to some degree encapsulated from back-ground knowledge, beliefs, and expectations. These questions have long been the focus of inquiry in cognitive science, with implications not only for standard cases of sentence parsing and early visual perception but, with language, for related questions such as whether there is a semantics/pragmatics distinction, and the nature of compositionality. In this chapter, we explore these latter topics and their relevance for the general hypothesis of language modularity. In particular, we discuss which type of *semantic* representations might be computed within a linguistic module, or, rather, serve as the *output* of computations performed by the module.

Pertinent to our general goals is the following question (Fodor, 1983, p. 88):

[W]hich phenomenologically accessible properties of an utterance are such that, on the one hand, their recovery is mandatory, fast, and relevant to the perceptual encoding of the utterance and, on the other, such that their recovery might be achieved by an informationally encapsulated computational mechanism?

Although Fodor does not offer a precise answer to this question, he suggests a research program, which we plan to elaborate on:

[W]hile there could perhaps be an algorithm for parsing, there surely could not be an algorithm for estimating communicative intentions in anything like their full diversity. Arguments about what an author meant are thus able to be interminable in ways in which arguments about what he said are not. (p. 90)

The research program lies, of course, in determining where the line should be drawn between sentence parsing and the recovery of speaker intentions that go

#### Semantics for a Module

beyond communicating information the sentence conventionally contributes. We contend that somewhere between parsing (viz., syntactic analysis) and the points a speaker intends to get across lies an encapsulated semantic representation—if you like, the specification of the proposition the uttered sentence expresses—which we take to be an excellent candidate for what the language module delivers to higher cognitive systems.

Our case for a *semantic* (propositional) output of the module will be made along the following lines: first, we present general assumptions about the cognitive architecture that underlies our view about the workings of a language module-in particular, about how the module computes sentential (and perhaps even discourse) meaning. We then discuss how a semantic representation might be seen as a "shallow" representation, which the module outputs. We support our view by presenting theoretical and experimental evidence for socalled cases of indeterminacy-specifically, sentences which, when tokened, are alleged to only express propositions after enrichment via pragmatic operations or lexical-semantic decomposition. We claim, contra the mainstream, that these discourses are arguably, by default, output by the module as complete, unenriched propositions, and that these propositions are only minimally contextsensitive (viz., sensitive to antecedents within a discourse), and further, that the recovery of whatever information theorists claim is required to understand their utterances is the product of functions performed by other cognitive systems of interpretation-that is, distinctly not linguistic ones.

### MODULARITY AND COGNITIVE ARCHITECTURE

Most discussions on the properties of a language faculty come infused with assumptions about the nature of mental representations and processes-more precisely, what kinds of representations the language system encodes, and how the system performs its functions. We want to make our assumptions on these issues explicit because what we will say about modularity—and, in particular, about the output of the language module-requires a clear understanding of the type of cognitive architecture we assume supports the workings of the language system and its interfaces with other domains of cognition. The need for an explicit commitment to a particular type of cognitive architecture also derives from the kind of thesis we advance here concerning the very nature of the semantic output of the language module: in short, we will postulate that a semantic output is not only plausible but perhaps even necessary vis-à-vis a sharp distinction between operations of an encapsulated linguistic system and operations contingent on world knowledge, beliefs, expectations, and so forth. A question, to which we will return, is what sort of semantic system that might be. But we begin with rather orthodox cognitive-architecture commitments.

We are committed to two closely related guiding assumptions on cognitive architecture and modularity: first, we take representations to be *symbolic* insofar as the codes that serve for perceptual and cognitive computations are symbols—simple and complex. We will see that in order to distinguish between semantic representations and other types of knowledge and beliefs we need to adopt a symbolic architecture; what the system outputs is the product of its semantically interpretable computations—namely, a symbolic expression encoding a proposition. A key feature of this representation is that it accounts for the compositionality of whatever linguistic expression that expresses the proposition. Compositionality is a fundamental property of symbolic systems and, as far as we know, of symbolic systems alone. Although we take these symbols to be physically instantiated as patterns of neuronal code or, perhaps, as Gallistel (2018; Gallistel & King, 2010) has proposed, as molecular changes within neurons, the assumption that representations are symbolic is largely independent of hypotheses on the actual form of neuronal implementation. This is so because we take representations and processes to constitute, in principle, an autonomous explanatory level of how cognitive systems work.<sup>1</sup> So much for our first assumption.

Our second assumption is that mental processes—that is, how information is perceived, analyzed, transformed, and so forth—are *computational* and, to a large extent, realized algorithmically. That is, many of the operations that the mind performs—and, in particular, the ones that ought to be taken as *intra*modular—follow rule-like symbol manipulation. These processes, to a first approximation, are operations performed entirely in a mechanical fashion, akin to logical or mathematical proofs. Now, it is certainly a matter of empirical investigation which cognitive processes can be cast in terms of *algorithms*, which ones follow *heuristic* principles, and which ones are entirely subject to contingencies.<sup>2</sup> Indeed, it is perhaps this second guiding assumption—the extent to which certain processes are algorithmic—that constitutes the main overarching hypothesis bearing on the modularity of mind: that at least some perceptual computations are fixed and operate without interference from background knowledge and beliefs.

Symbolic representations and computational processes are well known guiding assumptions, adopted by a great number of cognitive scientists, in particular, those who subscribe to some version of the Representational/Computational Theory of Mind.<sup>3</sup> It is important, though, to recite them because much of our discussion on the nature of modularity and the semantic system that we take to be internal to—or as an output of—the language module rests on there being *algorithmic* processes, that is, fixed rule-like computations over symbolic representations in the course of analyzing a sentence.<sup>4</sup> We should further qualify these assumptions regarding the nature of representations and processes, mainly because the symbolic-computational view has generated numerous controversies (and misunderstandings), particularly on the semantic nature of "amodal" symbols and whether or not symbolic computations are "meaningless" (e.g., Searle, 1980).

Although we assume that the computations of a module are formal, the symbols the language system computes, according to us, specify, inter alia, truth conditions, and therefore, the system must distinguish *types* of representations

in the course of performing its computations. In other words, the computations must be sensitive to *symbol types* and symbolic structures, and so, the relevant semantic distinctions between representations must be coded in the symbols or symbolic structures. Thus, while operations are performed in a mechanical fashion, semantic distinctions ought to be coded symbolically.<sup>5</sup> For example, sentence-parsing computations must be sensitive to distinctions between verb *types*, which in turn must determine the course of processing routines bearing on the *types* of arguments a verb takes and, hence, interpretation should be sensitive to the nature of symbols and how they are structured.<sup>6</sup> We will return to this issue below, in the context of examining so-called indeterminate sentence interpretation.

The upshot of this brief discussion is that a module ought to perform its computations algorithmically, with computations being sensitive to type/token distinctions that are supposed to be encoded in the elementary symbols and symbolic expressions. We will turn now to how computations performed within the module—especially within the hypothetical language module—might be carried out, and what sort of output the module might produce.

## THE BOUNDARIES OF LINGUISTIC MODULARITY

Fodor (1983) discussed, at length, criteria for what he called "vertical faculties" (e.g., information encapsulation, domain-specificity, neurological specialization, genetic determination, fast and mandatory computations). So, we won't further exegete here. Our focus, instead, will be primarily on two criteria that bear more directly on our main point: information encapsulation of modular computations, and "shallow" outputs that modules are supposed to produce. Arguments for information encapsulation, simply put, turn on the degree of functional autonomy of a given perceptual domain (e.g., language) with regards to other "vertical" faculties (e.g., vision), general "horizontal" faculties (e.g., long-term memory) and, more broadly, the central system (the "Quineian" holistic space), where encyclopedic and episodic knowledge, beliefs and so forth reside. Arguments for "shallow" outputs turn on the types of representations a module produces: by hypothesis, they do not produce full analyses but, rather, minimally, translations of post-transduced inputs that preserve the nature of the distal stimulus (viz., relevant properties of linguistic structure and lexical denotations). To a certain degree, we assume that whatever the language system computes-the operations on its natural-kind linguistic data-is to a large extent encapsulated. As Fodor puts it,

... data that can bear on the confirmation of perceptual hypotheses includes, in the general case, considerably less than the organism may know. That is, the confirmation function for input systems does not have access to all of the information that the organism internally represents; there are restrictions upon the allocation of internally represented information to input processes. (Fodor, 1983, p. 69)

The question becomes what exactly "the system might know" in order to yield a sufficiently rich representation of the input without being contingent on "all of the information" the organism might internally represent. The standard answer, in the course of the modularity debates in psycholinguistics, has been to focus on syntactic parsing (see, e.g., Townsend & Bever, 2001, for a review). And the research strategy has been to show that syntactic parsing might be unsettled by semantic (i.e., "background") knowledge. Crucially, this strategy rests on the assumption that syntactic analyses are immune to semantic variables—thus, any demonstration of semantic influence on syntactic analyses ought to be a violation of the key encapsulation principle. But while this research strategy has proven fruitful, producing an enormous amount of data (as Fodor says, "that's why we have careers"), it seems to us that it also misses the mark. This is so because what is modular is entirely dependent on the sort of fixed linguistic information the input system might encode. It may turn out that some "semantic" data bearing on structural analyses is encoded in the symbolic rules that govern the workings of the parsing routines, and thus, might have influence on what sort of parsing choices a syntactic system might make. To wit, it may be, for example, that thematic/"semantic" roles assigned by verbs to their arguments are part of the database that the parsing system consults in making its choices; it may also be that properties of thematic/"semantic" structure enter into determining what sort of representation the input system might produce. For instance, it has been demonstrated that, traditionally, "semantic" information encoded in different verb classes are affected selectively in cases of brain traumas or diseases (see, e.g., Piñango, 2006; Thompson & Lee, 2009). Data from Alzheimer's patients, in fact, suggest that verb-thematic hierarchy plays a significant role in patients' preferences for how the arguments of a verb map onto syntactic structure (see Manouilidou, de Almeida, Nair, & Schwartz, 2009; Manouilidou & de Almeida, 2009). Patients have no difficulty with canonical sentences in which the noun phrase in subject position is assigned an Agent role, as in (1a). However, patients have difficulty understanding the sentence when the subject position is occupied by an Experiencer, as in the case of subjects of psychological verbs, as in (1b). Moreover, when the verb assigns the role of *Experiencer* not to the noun phrase in subject position but to the one in object position, as in (1c), patients show a much greater impairment in comprehension. This effect is independent of voice, that is, it is obtained even in the passive form of the same sentences, where the linear order of constituents is inverted.

- (1) a. The gardener cultivated the carrots (Agent, Theme)
  - b. The public admired the statue (*Experiencer*, *Theme*)
  - c. The statue fascinated the public (Theme, Experiencer)

It is quite plausible that parsing operations rely on more than the to-besaturated arguments of a verb and their structural arrangements: decisions might also take into account the role the arguments play in the semantic representation of a sentence. This view, of course, does not commit us to a particular ontology of thematic roles, but simply suggests that semantic information can enter into decisions the parser makes. A parsing model such as Bornkessel and Schlesewsky's (2006), although "incremental," seems to be entirely governed by principles that include algorithmic and heuristic rules for determining structural choices concerning verb-argument relations. This seems to be a "modular"—input-driven—system whose choices are not dependent on background information but on encoded syntactic and semantic principles.

We end this section with a summary of our guiding assumptions, and how they relate to our view that a semantic analysis is the output of the language module. We assume modular systems operate as computations over posttransduced symbolic expressions. Moreover, we assume modular computations are sensitive to semantic distinctions among symbolic expressions, and thus, the input to linguistic analysis could well be guided by encoded—fixed—semantic principles. As we saw, there is a case to be made for semantic representations being the determinants of intra-modular decisions: thematic-role assignment is just one, enriching the nature of the computations that the input system for language computes. In the next two sections, we elaborate, first, on what sort of representation serves as output for the module. We aim to show that the module computes shallow semantic information on the assumption that the input system knows "considerably less" than what the sentence is *about*. We then focus on a particular case: sentences whose propositional contents are alleged to require enrichment in order to explain what their uses can communicate. We aim to show that whatever this sort of enrichment includes is a function of contextual information that goes beyond input analysis.

### **OUTPUT: A SHALLOW PROPOSITION**

The proposal that the language module outputs a type of semantic representation suggests that one function of the perceptual system is to analyze utterance content. But is the idea of intra-modular semantic representations and processes in conflict with sentence encapsulation? Encapsulation, after all, requires that semantics not be served by background knowledge and other systems of interpretation, and this requires a clear distinction between semantic properties that are encapsulated (thus, algorithmic) and other knowledge systems that are not. This is all true, but still, we will defend a view of semantics that is compatible with modularity, where *semantic* representation is recoverable from what is expressed overtly by sentence constituents (viz., lexical, morphological, and functional constituents) and syntactic (and discourse) arrangement. Our proposal, more specifically, is that the symbolic expression the language module outputs carries all the relevant information for (further) elaboration by higher cognitive systems. We take, in short, the proposition that a sentence conveys to be recoverable from its constituents, its structure, and its linguistic relations to other sentences in a discourse. More importantly, we argue that these broader contextual effects-always lurking as a threat to modularity-are either, as a

matter of fact, intra-modular (viz., linguistically determined by semantic or syntactic operations inside the module), or are post-modular (i.e., higher cognitive effects on modular outputs).

Determining the nature of the semantic output of the module depends fundamentally on what one takes to be "semantic." In *Modularity*, Fodor often mentions that, among the tasks of an encapsulated language system is that of producing a *logical form*. Although this is not explored in detail, if we take the symbols the module computes to be distinguished by semantic properties, then the logical form that the module outputs has many of the ingredients interpretative processes require in order to perform their basic functions. In other words, if symbolic expressions carry semantic properties that distinguish them from one another, we can assume that much of what the input system does is to produce the semantically relevant representation that symbolic combinations yield.

This view is a bit more explicit in Fodor (1990):

[W]e are committed to the traditional assumption that the outputs of parsers are representations of utterances that determine their meaning. Roughly, a parser maps wave forms of (sentence-length) utterances onto representations of the proposition that they express. (p. 4)

Significantly, what the parser outputs determines what the sentence means; its (output) representation is mapped onto the proposition that the sentence expresses. Fodor adds to this "a position that is quite appealing":

Parsing recovers the lexico-syntactical structure of an utterance, but *no* semantic-level representation (no representation that generalizes across synonymous utterances; hence, for example, no model theoretic representation). (p. 8)

Fodor's main reason for keeping semantics out of the module's computational tasks is his belief that, in order to perform any sort of semantic computation "the speaker/hearer would have to know the semantics of their language in the way that they indisputably have to know, say, its lexical inventory." (p. 8)<sup>7</sup>

It is not clear exactly what sort of semantics the speaker/hearer "would have to know" nor is it clear what sort of representations they would need to encode in order to compute sentential meaning while preserving modularity. We assume something akin to lexical denotations (or pointers to lexical denotations; viz., morphemes) and whatever apparatus yields a logical form would suffice. Fodor (1990), in fact, leaves the door open to some form of semantic representation within the module by proposing the following:

 We will use 'parsing' to name a process that starts with acoustic specifications and ends with meaning representations.  We assume that meaning representations are probably distinct from, and probably constructed on the basis of, syntactic structural descriptions; (...) (pp. 8–9)

The proposal that parsing "ends with meaning representations" entails some form of representation that might be available to other systems of "interpretation." This is clearer in the model Fodor proposes, encompassing the following serial stages:

(a) acoustics → (b) structural description (syntax) → (c) meaning representation →
(d) discourse model → (e) real world model<sup>8</sup>

With the exception of the last stage (e), Fodor leaves it open which operations might be encapsulated. It might be safe to assume that the first transformation, from (a) to (b), follows from the initial transduction of linguistic codes devoid of semantic content: the operations, by hypothesis, are transformations over symbols or symbolic expressions representing the likes of grammatical categories. That is where the standard view of modularity assumes information encapsulation ends, and that is where many studies have suggested there is penetration by meaning representations or even the discourse model. But Fodor's (1990) revised version of modularity traces the line higher, admitting that meaning representation and even a discourse model could be computed by the modular parser. In fact, faced with many studies suggesting that parsing might be influenced by local context (e.g., Crain & Steedman, 1985; Altmann & Steedman, 1988), Fodor assumes the level of discourse representation (d) might provide intra-modular constraints to decisions made at (b), the building of syntactic representations.

This last point is crucial for our understanding of a modular mechanism that is semantically shallow at the output level, while assuming that outputs are representations of the propositional content of input sentences. We take "propositional content" to include denotations of what is overtly specified in the sentence—namely, its lexical or morphological constituents—as well as how these constituents are structurally arranged. Moreover, we also take this propositional content to include the specification of the linguistically active but perhaps phonologically null elements that constitute the sentence. These elements are like the nominal antecedents of pronouns, elliptical verb phrases, and other linguistically specified elements (including cross clausal and sentential mechanisms for establishing discourse anaphora).

Thus far, this amounts to assuming that propositional content is compositional—or, to be more precise, obeys "classical" compositionality—for even in cases where propositional content might be attributed to phonologically null elements, they must be linguistically (i.e., syntactically) licensed. However, there are cases in which the elements called for by phonologically null and overt elements are outside the scope of the sentence proper. One can imagine contexts in which pronouns have their antecedents in the immediate discourse, or cases of indexicals (e.g., "there" and "now"), which pick up their contents from already referenced discourse elements, or even from the visual context. Below, we discuss in more detail experimental evidence for linguistically determined discourse enrichment of sentences.

We assume that these cases can be accommodated by a theory that takes the building of a proposition to be conventionally governed by contextual factors linguistically determined and part of the local discourse. One case is presupposition. Assuming that what speaker A says to B makes reference to what is in their common ground, one can take pieces of the common ground to aid in the proposition-building process. To be clear, what is presupposed is, to a large extent, linguistically determined. Thus, it might be part of the "discourse model" that Fodor refers to, providing intra-modular constraints on the types of information that constituents of the proposition pick up.

Consider now reference in a visual context. Imagine that upon referring to a particular person on the scene, speaker A says to B: "That is the girl I told you about." Indexicals such as "that," "I" and "you" as well as what has been talked "about" (the girl) constitute elements of discourse that may enter into the propositions that A and B coordinate on during their linguistic exchange. What A and B talk about, or refer to, in the context are not sentential, but rather (local) discourse, constituents that hold a special relation to both the sentence A utters and the information B exploits in building a propositional representation of what A says. A modular output might build the proposition that the sentence expresses taking into account the elements that are within the immediate discourse (discourse referents).

For another case in point consider a discourse like (2):

(2) A man walked in. He sat down.

On one of its readings, 'He' is interpreted as co-varying with 'A man.' (2) is all true just in case some man walked in and sat down. The pronoun resolution for (2) is guided by an implicit organization that knits together information in discourse. On this anaphoric reading, the discourse begins with a description involving 'A man' and proceeds directly to develop a narrative: accordingly, 'He' is interpreted as dependent on 'A man.' This information is entirely encapsulated to this discourse. Confronted with this discourse, without any attendant pointings or other sorts of gestures, speakers know automatically how to interpret it and resolve its pronoun. (For many other examples and for a general defense of the claim that all context sensitivity in resolved in this rule governed conventionalized manner, see Stojnic, Stone & Lepore, 2017).

Our view of Fodor's revised program for a module assumes that the basic representation the module outputs is sensitive to contextual information but always in a rule-governed fashion. The elements upon which B builds propositions corresponding to A's contribution to the conversational record begins with what A says with an utterance and might include what they both take to be common ground as well as other conventional contributions. In keeping with modularity, our proposal, then, is that the output encodes "considerably less" than what an individual takes away from another's utterance. It encodes what is linguistically determined.

We are now ready to turn to cases where the interface between modular computations and semantic/discourse interpretations are in play: resolution of (alleged) indeterminacy.

#### SENTENCE INDETERMINACY

What happens when the linguistic contribution of an uttered sentence underdetermines what its speaker is able to get across with her utterance? For example, on hearing in isolation an utterance of (3a), we might infer reference to an event in which a man began *reading* a book. But obviously, *reading* is not the only possible inference; why not *eating*? And if nothing is off limits, shouldn't we conclude (3a) is *indeterminate* with regards to the event it refers to, that is, indeterminate with regards to what some man began doing with some book?

- (3) a. A man began a book
  - b.  $\exists x (= \text{man}), \exists y (= \text{book}) (begin (x, y))^9$

To be clear, indeterminacy issues from the activity the aspectual verb *begin* scopes over. There is, however, a default interpretation for this type of sentence—that a man began doing something with a book—which is the proposition (3a) expresses, *ceteris paribus*. If we assume that what the module computes is minimally a logical form that captures the proposition that the sentence expresses, then (3a) ought to output something like (3b).

We have an initial observation about the encoding of (3a) as (3b). Clearly, (3b) does not exhaust everything that an utterance of (3a) can get across, nor is it supposed to. What (3b) might specify is how the transduced symbols of the proximal stimulation of (3a) is to be encoded. (3b) is a proposed *symbolic* output of stimulus (3a). (As we mentioned, other symbols might enter into the shallow semantic representation for (3a). For instance, if thematic roles are encoded, they might enter into the representation the module outputs.<sup>10</sup> Crucially, what (3b) might account for is the logical structure of (3a) together with the translation of its lexical/morphological constituents.)

There are linguistic-theoretic and experimental treatments of (3a) that assume it gets enriched even at the linguistic level of analysis by a default operation of *coercion*. The key idea is that a verb such as *begin* (and many others within the same class) require an event complement, the absence of which triggers, roughly, a change in the nature of the internal argument (in (3a), *book*) to make it fit with the requirements of the verb. One proposal, we call *coercion with interpolation*, hypothesizes that the supposed mismatch between the event-taking verb (*begin*) and an entity-type object (*book*) is resolved by, first, extrapolating an event information from the lexical entry *book*; and, second, by interpolating this event information into the semantic composition (deemed *enriched*) of the sentence (see Pustejovsky, 1995, and Jackendoff, 2002, about variants of this process). Another coercion proposal, we call *type-shifting*, makes no direct claims for interpolation, nor does it assume any form of extrapolation of information from the lexical entry: what it proposes is that the entity denoting a book is supposed to change its type to an *event*, to respect the requirements of the event-selecting verb (e.g., Pylkkanen, 2008).

These two proposals agree that (3a)'s being enriched linguistically in the "classical" way does not work, for what's needed is a way to make the arguments "fit" with their selecting verbs. The proposals differ about the source of enrichment and, consequently, about their commitments vis-à-vis the nature of semantics. Type shifting rests on an ontology of semantic *types* that has not been established—one we are not prepared to adopt. The idea that the alleged verbargument mismatch is resolved by changing the semantic type of the complement strikes us as affirming the consequent. But our main reason for suspicion that enrichment is obtained via type shifting is that it requires postulating at least one of two things: that token nouns are loaded with information about their possible types and their modes of combination with their host verbs; and that semantic principles are informed about these modes of combination.

The assumption that items are informed about their possible semantic types entails that lexical items are polysemous between diverse types they can be coerced into. One has to assume that this is true of all lexical forms. The assumption that semantic-type combinations are driven by rules also assumes that the rules ought to be informed about the default types of token items. Either version of this approach to coercion relies on internal analyses of token items to yield appropriate combinations as well as to reject anomalous ones. But how should the linguistic system be informed about such semantic properties (viz., semantic types and their appropriate combinations) without also being informed about putatively holistic world facts-arguably determinants of plausibility? To put it simply, our main point is that to know that *book* can be read on the basis of the linguistic system as an event performed with a book requires knowing that the noun *book* allows for such an *event* reading upon the implausibility generated by its *entity* reading. Simply postulating that the *entity* $\rightarrow$ *event* shift is demanded by the verb does not work because type shifting relies on analyzing the noun default type before triggering the shifting operation.

Coercion with interpolation runs into its own problems. Most notoriously, it rests on an analytic/synthetic distinction, as pointed out elsewhere (see Fodor & Lepore, 2002; de Almeida & Dwivedi, 2008; de Almeida & Riven, 2012). Semantic interpolation requires a vastly rich, encyclopedic lexicon, whose properties are supposed to provide filler information—what can be interpolated in the resulting semantic composition. For instance, we would need to know a great deal about books in order to find out what is possible or likely for one to begin doing with them, in order to select an appropriate filler to enrich (3a). And besides, the lack of a principled analytic/synthetic distinction leaves us wondering what sorts of properties—the ones which are supposed to constitute the lexical-semantic information—are to be regarded as part of the "semantic lexicon" and which are not. Moreover, there is no evidence that nouns such as *book* are actually constituted by properties or features, let alone that this process takes place at a linguistic level of analysis, a level independent from general knowledge.

One issue common to both the coercion proposals we have sketched is that they assume *entity* arguments don't fit with verbs such as *begin*. A linguistic test used to support the alleged oddness of the *begin-entity* combination rests on showing that so-called *event* nominals do not require coercion, such as in (4).

(4) The general began the war

This argument, however, has little validity, for sentences with *begin-event* forms might also require enrichment. Simply put, if *x begins entity* calls for an event interpretation, so does *x begins event*: (4) is neither synonymous with (5a), nor does it entail (5a). Rather, (4) can be roughly paraphrased by something like (5b) because one can begin a war without fighting it (see de Almeida & Riven, 2012, for further discussion on this issue; see also Cappelen & Lepore, 2005, for a defense of slippery slope arguments of this sort).

- (5) a. The general began fighting the war
  - b. The general caused the war to begin

An alternative to coercion assumes that (3a) remains linguistically "indeterminate" with respect to what sorts of enrichments its tokenings might admit of. Crucially, this view assumes enrichment is beyond the linguistic level of analysis—it comes from post-linguistic processes. Such processes are most likely abductive, for they take into account what might be contextually appropriate, what might be most probable, etc. Essentially, this view, assumes what the module outputs is based on what the sentence expresses (at least in a discourse) without lexical decomposition (*a la* coercion with interpolation) or type-shifting. Moreover, this view assumes the enrichment of (3a) is linguistically motivated by the *syntax* of VP. Key here is that verbs such as *begin* are represented by an argument structure that specifies a syntactic position for the filler event, as in (6).

(6)  $\left[ _{VP} \left[ _{V0} began \left[ _{V0} e \left[ _{OBI} NP \right] \right] \right] \right]$ 

There are many distributional arguments for the *linguistic* reality this gap within VP (see de Almeida & Dwivedi, 2008; and de Almeida & Riven, 2012). For instance, it is within VP in the second clause of (7) where verb ellipsis is realized—that is, where the second-clause *reading* is syntactically determined to re-appear.

(7) I started reading with my contacts but finished  $[[_{VP} [_{V0} e]]_{PP}$  with my glasses]]]

What is important about this proposal is that the gap (*e*) might serve as the trigger for the inferences that would ultimately enrich (3a). It may be that the proposition (3a) expresses, then, allows for the gap that we suggest occurs in VP of a verb such as *begin* (but see note 9). The key point we want to register is that whichever form this representation takes, (a) it does not specify *how* (3a) is enriched (i.e., it does not determine a default content for (3a)); and (b) it provides a linguistically-motived basis for enrichment without committing to a type-shifting analysis of the complement. In both cases, the syntactic gap analysis provides a linguistically-motivated source for ulterior enrichment thus avoiding the problems that afflict the different views of coercion.

We now turn to experimental evidence, first in support of coercion views, and then, against coercion.

Experimental work supporting coercion is slim. Earlier studies (e.g., McElree, Traxler, Pickering, Seely, & Jackendoff, 2001) have shown that (8a) takes longer to process at post-verbal positions when compared to (8b). This extra time was assumed to be due to the process of semantic interpolation.

- (8) a. The secretary began the memo before it was due
  - b. The secretary typed the memo before it was due

Obtaining longer reading times at post-verbal positions need not constitute support for "interpolation." Alleged indeterminate sentences such as (3a) or (8a) differ syntactically from fully determined ones such as (8b) (see de Almeida & Dwivedi, 2008, for linguistic analyses). Thus, longer RTs could be due to structural differences between them. Besides, results obtained by McElree et al. (2001) could not be replicated by de Almeida (2004, Experiment 1), employing a similar experimental paradigm and conditions. And while Pickering et al. (2005) have attempted to replicate McElree et al.'s results, most effects were statistically weak or reflected relatively late processes (e.g., second-pass reading), compatible with post-parsing enrichment. Replicability is of the essence for establishing a given phenomenon. But even if those results were to be consistently replicated, they only suggest there are differences between sentence types, without exactly accounting for what yields those differences; more specifically, they cannot be taken to support "interpolation" or "type-shifting" forms of coercion directly.

Similarly, experiments involving ERPs (event-related potentials) have shown processing differences between sentences such as (8a) and (8b), but without determining how these sentences differ (see, e.g., Kuperberg et al., 2010; Baggio et al., 2010). MEG (magnetoencephalography) experiments have also suggested that processing sentences such as (8a) and (8b) yield different magnetic patterns, but they too have not accounted for the source of the difference.

Most studies that claim support for coercion have in fact served two purposes: either they have been specifically designed to show that these sentences behave differently (thus, supposedly supporting some form of coercion), or they have focused on determining the anatomical source of the difference, on the assumption that coercion is necessarily at play. On both accounts, they are compatible with the view that indeterminacy is attributable to sentential structural properties. At the very least, they have shown that differences in processing are manifestations of structural differences. At most, they have shown that these sentences call for different enrichment processes, coercion or something else. More directly related to our concerns, a view that takes indeterminate sentences to be initially analyzed based on their constituents and syntactic form—thus, initially without enrichment—stands as the default.

Experiments using brain imaging-in particular, functional magnetic resonance imaging (fMRI)—can be further illuminating with regards to the source of enrichment, on the assumption that different anatomical sites might be engaged in processing different kinds of stimuli. There are, however, a few caveats regarding the use of fMRI to determine the nature of linguistic and nonlinguistic processes involved in indeterminacy resolution. First, accounting for differences between sentence types in terms of anatomical resources involved requires having a clear understanding of "where" or even "how" different kinds of syntactic, semantic, and pragmatic processes take place in the cortex or even in subcortical areas of the brain. Lacking such a clear understanding of the mapping of language (and post-linguistic processes) leaves us with a fair amount of speculation. This is akin to finding reading-time, ERP, or MEG differences without knowing the source of these differences. Second, even if we were to have a firm foundation upon which to build our neuroanatomical hypotheses, it is quite possible that similar networks might be deployed to achieve functionally different ends. While this is certainly a strong argument against a strict physicalist explanation, it is also a call for keeping the spotlight on the very theories that underlie the anatomical predictions. And third, there are numerous constraints on the analysis of fMRI data, which relies for the most part on set parameters of what is to be considered "activated" in the course of processing a given stimulus. At the voxel (unit of activation) level, this means determining a significance parameter; at the cluster level this means determining a particular number of contiguous voxels (the "regions") while leaving lower quantities at bay (the heap paradox comes to mind: why 100 voxels and not 99?); and, overall, establishing activation levels often requires leaving unreported networks that do not reach a given threshold but which, nonetheless might be engaged in processing the stimuli. Despite these general constraints on the use and interpretation of fMRI data, this technique can be used to complement both linguistic analyses of indeterminate sentences as well as studies employing behavioral and electrophysiological techniques.

Thus far, the main anatomical sites involved in the resolution of alleged indeterminacy (or in the *attempt* to resolve it) have been elusive. MEG studies (e.g., Pylkkanen & McElree, 2007) have suggested that the main area involved in interpreting sentences such as (8a) compared to (8b) is the ventro-medial prefrontal cortex (vmPFC). This area was activated following an initial bilateral temporal activation, though the estimate that the vmPFC is the main "site" of coercion is, at this juncture, highly speculative given the involvement of other areas. Also, the advantage that MEG has over fMRI in terms of

temporal resolution, it lacks in spatial resolution. Employing fMRI, Husband et al. (2011) found no evidence of vmPFC activation but greater activation at the left inferior frontal gyrus (IFG), suggesting that this region "supports the application of coercion processes" and "the detection of semantic type mismatch between the verb and its complement" (pp. 3260-3261). While these results are consistent with the idea that indeterminacy might involve a structural-gap detection, the claims that Husband et al. make go far beyond that. For them, activation of the IFG suggests "the mismatch and its repair only affect semantic composition and do not recruit other processes for repair or rejection" (p. 3262). Their idea is that semantic composition incorporates mechanisms of detection of anomaly and repair, though it is not clear on what grounds semantic anomaly is detected, or how repair is obtained. The only way to assume this is happening is to assume that—as we discussed above the semantic composition system is informed about world contingencies. This seems to be the position they take: "Assuming event meanings for nouns are also stored in the lexicon (Pustejovsky, 1995), IFG may function to select and retrieve the noun's event-related meaning" (Husband et al, 2011, p. 3262). But of course, this cannot be achieved unless there is an account of the analytic/ synthetic distinction for lexical-semantic representations-which, as far as we know, nobody has.

The first caveat above—regarding the lack of clear neuroanatomical parameters for linguistic and post-linguistic processes—requires us to investigate phenomena that are poorly understood by taking a broad stance. The most parsimonious approach is to map out the process, typically reporting its neuronal correlates, by contrasting several variables. For instance, contrasting sentences such as those in (9), representing a wide spectrum of normal and abnormal constructions allow us to dissociate indeterminate sentences such as (9a) from sentence types such as those that are determinate (9b) or semantically/pragmatically anomalous (9c), or even syntactically anomalous (9d). Underlying this approach is the assumption that differences and similarities in terms of regions, activation levels, or even number of activated voxels obtained between these sentences are indicative of the nature of the resources involved in the processes of parsing and interpretation.

- (9) a. The author started the book.
  - b. The author wrote the book.
  - c. The author drank the book.
  - d. The author yawned the book.

In the fMRI study conducted by de Almeida, Riven, Manouilidou, Lungu, Dwivedi, Jarema, and Gillon (2016), employing sentences such as those in (9), the neuronal correlates of indeterminacy resolution were found to be somewhat different from those in previous studies.<sup>11</sup> Indeterminate sentences such as (9a) were found to activate a wide network, in particular, the left and right IFG, both



Figure 5.1 Partial results from de Almeida et al.'s (2016) fMRI study. Areas within ellipsis represent some the main regions activated in the contrast between "indeterminate" (such as (9a)) and "determinate" (9b) sentence types. Activation maps represent (a) the right hemisphere, superior temporal gyrus (Talairach +45), (b) medial right hemisphere (+4), with activation of the anterior cingulate cortex (ACC), and (c) the left hemisphere, superior temporal gyrus (-48) regions. For color figures and more details, see http://journal.frontiersin.org/article/10.3389/fnhum.2016.00614/full.

temporal lobes and the anterior cingulate cortex (ACC). While other sentences also showed above set threshold activation in so-called "language areas" (left superior temporal lobe and L-IFG), indeterminate sentences surpassed other sentences in all those regions. Figure 5.1 shows data for the contrast between indeterminate and determinate (control) sentences—(9a) and (9b), respectively.

In addition, as Figure 5.2 shows, the number of voxels activated for indeterminate sentences by far surpasses those activated for other sentences in (9)—even in cases of blatant semantic and syntactic violations, such as in (9a) and (9b).

While these data do not completely rule out coercion, they point to a different perspective, one compatible with the one we proposed: greater activation beyond traditional linguistic areas for indeterminate sentences allied to overall greater



**Figure 5.2** Number of whole-brain significantly activated voxels for sentences in (8). From de Almeida et al. (2016).

number of activated voxels bilaterally suggest that indeterminate sentences trigger a search for a resolution, consistent with a state of uncertainty—more so than with a default intra-linguistic semantic coercion.

## MODULARITY AND CONTEXT SENSITIVITY

So-called indeterminate sentences are supposed to constitute a challenge to modularity: if they are resolved during initial parsing, they ought to be resolved based on knowledge that traditionally lies outside the module. The lack of an analytic/synthetic distinction defers the resolution of indeterminacy to post-parsing mechanisms of interpretation. We have assumed that the output of the module is something akin to a proposition, but one *un*enriched by local lexical-semantic processes. In our proposal, however, syntactic (and discourse) triggers work to signal higher interpretive processes where enrichment might be due. And if the syntactic analysis presented in section (5) holds, the trigger is within the VP. The widespread activations that indeterminate sentences cause suggest that there is at least an attempt to resolve indeterminacy.

In principle, cases such as A man began a book appear to be well resolved or enriched because they come embedded in utterance contexts. Nobody relatively sane addresses you with this sentence without having first established a frame of reference or presuppositions, outside of a common ground. Few experimental studies have attempted to manipulate the role of context in processing indeterminate sentences (de Almeida, 2004; Traxler et al., 2005), and the results have been inconsistent. We have assumed that so-called indeterminate sentences are indeterminate only in isolation, that no enrichment takes place by default, not at least by coercion. But we have also speculated that these sentences harbor a syntactic position that might serve as a "trigger" for processes of enrichment down the stream. We have also suggested that there are ruled governed (that is, conventionally determined) resolutions for some (much?) of what goes under the general rubric of "indeterminacy," for example, as in the cases of pronoun resolution in a discourse, as in (2) above. This is entirely consonant with Fodor's (1990) revised modularity model, which takes the scope of modularity to be not the sentence but, more broadly, what he called "discourse model."

This discourse model, to be clear, is also local, for it relies on linguistically determined links among sentence and clauses, and various discourse elements, such as pronouns, tenses, elliptical verb phrases and the like (cf., Stojnic, Stone & Lepore, 2017; Stojnic, 2016a, 2016b). The very use of indefinite article is taken to presuppose the introduction of a novel discourse referent. By calling for "a man," "a book," etc., one grounds their interpretation to elements not yet established in the prior discourse. Perhaps, more directly related to our immediate concerns is the idea that the VPs of some so-called indeterminate sentences carry a trigger, as in (6), above. We take it that the role of this trigger, in the absence of a supporting context, is to generate inferences—some abductive—that will attempt to put an end to any appearance of indeterminacy. But uttering such sentences within a discourse, allows the trigger to operate locally, picking out elements that have

been either clearly established or hinted, and thus part of the propositions that the preceding context generates.

Effects of a preceding discourse on the processing of an alleged indeterminate sentence have been investigated in only two studies, with results somewhat consistent. De Almeida (2004, Experiment 2) found that a "context" such as (10a) facilitated the processing of a sentence such as (10b)—a contextually preferred sentence (following norms)—compared to less appropriate sentences (10c) and (10d) which took equally longer to process at the complement NP (*the memo*) than (10a). While this does not constitute facilitation of (10c), the relevant findings here are that (i) both (10c) and (10d) were less contextually appropriate and (ii) there was no extra cost associated with indeterminacy when context provided a relevant (local) filler for the indeterminate sentence (say, *working* on the memo).

- (10) a. The secretary would always be sure to work ahead of schedule. She was asked to work on a memo.
  - b. The secretary typed the memo long before it was due.
  - c. The secretary began the memo long before it was due.
  - d. The secretary read the memo long before it was due.

An eye-tracking study by Traxler et al. (2005) was closer to obtaining a real facilitation effect of an indeterminate sentence by its local discourse. It is perhaps in their Experiment 3 where we can find clearer results.<sup>12</sup> They presented "contexts" such as (11a) or (11b), which were followed by "target" sentences such as (12a) or (12b) in a factorial design.

- (11) a. The student read a book in his dorm room.
  - b. The student started a book in his dorm room.
- (12) a. Before he read the book about the opium trade, he checked his email.
  - b. Before he started the book about the opium trade, he checked his email.

While they found differences in reading times between "context" sentences (11), they found no differences between "target" sentences (12). We have seen that indeterminate sentences in isolation can produce longer reading times—though not consistently so. That's the case of their "context" sentences, which precede their targets. Also, we have seen that the cost associated with indeterminate sentences compared to controls can be accounted for by differences in syntactic structure, as in (6). Thus, here again, coercion cannot be the only explanation. More importantly, the null effects they obtained in the "target" sentences in (12) can be seen as an effect of attenuation of target by context. First, it is expected that (11a) primes (12a) by virtue of repetition of the VP *read a/the book*. The same can be said of the pair (11b) and (12b). When "context" and "target" types are crossed, however, having "read a book" in the context, as in (11a), does not speed

up reading of the indeterminate (12b) any more than having "begin a book" facilitates "begin the book."

To summarize, processing an indeterminate sentence in a biasing discourse one that provides a potential filler event—facilitates resolution. Key here is that sentences—or propositions they yield—are sensitive to information within the "discourse model." That does not constitute a violation of modularity, for the information that so-called indeterminate sentences seek are within the local context and do not depend on analytic lexical decompositions. Put somewhat differently, they do not violate modularity because the resolutions are entirely predictable (because these discourse resolutions are entirely conventionally (linguistically) governed).

It is important to highlight that what we are here calling enrichment goes beyond the local discourse in the relevant sense. Sentences in discourses that dictate resolutions are not enriched. Rather, the effects of prior discourse in the enrichment of indeterminate sentences unfolds across the discourse (in a rule governed fashion). A study from Riven and de Almeida (2017) might be taken to support this view. Participants heard biasing contexts such as (13a) and, either immediately after the critical clause *Lisa began the book*, or 25 seconds after it (with intervening neutral discourse), they were presented visually with one of the probe sentences (13b)–(13d). Participants were asked to press a button indicating whether probe sentences were identical to segments they heard.

- (13) a. Lisa had been looking forward to the new Grisham novel ever since it came out. She had finally managed to set aside some time this weekend and made sure to make her home library nice and cozy. First thing Saturday morning, Lisa curled up on the sofa in her library with a blanket and a fresh cup of coffee. With everything in place, *Lisa began the book*. [Immediate probe point; discourse continues for 25 seconds]
  - b. Lisa began the book (identical/indeterminate)
  - c. Lisa began reading the book (biased foil)
  - d. Lisa began writing the book (non-biased foil)

This procedure is similar to one employed by Sachs (1967, 1974) showing the effect of propositional (or "gist") encoding of sentences in memory, with quick loss of verbatim representation. Crucial in our manipulation, however, was the effect that the context would have on participants' acceptance of the biased foil as if it was the original sentence. Here, contrary to previous studies, there is nothing in the context providing a clear event for enriching the indeterminate sentence other than suggestions that Lisa was about to *read* a book. In the case of (13a), the context is much closer to "hinting" about what is happening than providing a filler event. As can be seen in Figure 5.3, results show a clear effect of enrichment of the indeterminate sentence over time.



**Figure 5.3** Proportion of correct recognition for sentences (13b)–(13d) following contexts such as (13a) at two probe points: immediate and delayed by 25s, with intervening neutral discourse. From Riven & de Almeida (2017).

The biased foil (13c) is accepted as much as the original sentence heard in context. Confidence ratings collected after each trial, in fact, show that subjects are more confident that (13c) is the sentence they heard than they are of the true stimulus (13b). But these effects only obtain at the later probe point, not at the early one.

Overall, the results suggest that a sufficiently rich context might create a false memory—an effect of enrichment of the proposition—that is not driven by the local "discourse model" but one that comes from what Fodor referred to as "real world model." The line between the two, as we suggested, is thin, but one that makes a crucial distinction between encapsulated and unencapsulated processes: the former relies on linguistically determined enrichment, the latter not. While the discourse model provides a local source for antecedents of determinate noun phrases, pronouns, and the like, the "real world" hints on what is possibly the best way for enriching a sentence but without providing the actual information.

To summarize these results: the suggestion is that operations between sentence-level representations and local context are obtained within module, that is *local* enrichment is modular, for it is driven by linguistic processes. It is in this sense that sentences can be said to be mildly context-sensitive. In particular, local context or *co-text* provides the fillers that linguistic elements (syntactic gaps, pronouns, etc.) call for.

## CONCLUSION

In this chapter, we have tried to advance the view that sentence perception is largely computationally encapsulated; and, more surprisingly, that sentence meaning is context insensitive, or at least its sensitivity is rule-governed. The way these two work together is that while the output of sentence encapsulation is a minimally and highly constrained, context-sensitive representation of the sentence composed from its constituents, it remains semantic. The long-term challenge to a semantic output from a language module has been the alleged cases of interpretive indeterminacy. However, we showed how to assign semantic representations to such cases, and that they interact with context in highly regulated ways. We did not deny that such cases admit of enrichment of some sort or other, but we argued that these issues go well beyond anything concerning the language module itself. In short, we have defended a proposal for a semantic level of representation that serves as output of the module and as input to other systems of interpretation, arguing for a form of modularity or encapsulation that is minimally context sensitive provided that the information from context-whatever it may be-is itself determined nomologically, namely, by linguistic principles.

## AUTHORS' NOTE

Research for this article was supported by a grant from the Natural Sciences and Engineering Research Council of Canada (NSERC). We are thankful to Caitlyn Antal for comments on an earlier draft.

#### NOTES

- 1. We endorse the view that representations and processes are autonomous *qua an explanatory level*, but also that a full account of a cognitive system cannot dispense with a proper characterization of the biological implementation of its functions. See, e.g., the tri-level hypothesis of Pylyshyn (1984); and Gallistel (2018, this volume) for specific implementation proposals.
- 2. In common usage, *algorithms* are used for computational procedures that are guaranteed to produce a certain outcome, while *heuristic* rules are incomplete computational procedures. As noted by several authors (e.g., Pylyshyn, 1984; Haugeland, 1981), all computations involve algorithms and the distinction amounts to the end result—whether or not it is guaranteed to produce a given outcome. We will assume, in the present discussion, that algorithms—say, semantic ones—computed by the module are deterministic, while procedures on modular *outputs* (e.g., computing something akin to implicatures) are generally heuristic. These two options, of course, do not exhaust the range of possible cognitive mechanisms—including the possibility that some mechanisms might be entirely contingent on the individual's belief states.
- 3. By contrast one could postulate a connectionist type of system, with representations being nodes in a network and processes being activation patterns over those

nodes. Such a system could in principle be "modular" as long as the operations it performs are *encapsulated*, that is, are not subject to influences from other systems, and, in particular, from background knowledge (see below). In this case, the nodes at the "encapsulated" part of the system would have to be severed from feedback from higher-up nodes, in particular, those at the hypothetically *un*encapsulated part of the network. But this contrasts sharply with what connectionist networks stand for: that patterns of activation at lower levels are in large part constrained by patterns of activation at higher levels. Also, even if this could be fixed, the system would not operate *algorithmically*, nor would it be *compositional*, thus, lacking key architectural features to which our proposal adheres.

- 4. This certainly does not entail that modules (or "input analyzers," as Fodor reluctantly calls them) are the only systems to operate algorithmically, but they are the ones that compute algorithms on *post-transduced* symbols and, so, autonomously. Moreover, this does not entail that modules operate only algorithmically. It is possible to conceive of modular operations that are heuristic, as long as the choices it makes in the course of its computations are internal to the module; that is, encapsulated from general knowledge.
- 5. We take Pylyshyn (1984) to be rather clear about this: "formal symbolic structures mirror all relevant semantic distinctions to which the system is supposed to respond" (p. 74). See Pylyshyn (1984) for extensive discussion on symbols and their interpretation.
- 6. See, e.g., de Almeida and Manouilidou (2015) for a review on verb argument structure and on the content of arguments.
- 7. One reason Fodor keeps semantics out of the module is that semantic descriptions often appeal to lexical decomposition; and semantic theories that do so patently have to rely on an analytic/synthetic distinction. Fodor's rejection of this distinction implies that the module is open to all the possible beliefs the speaker/hearer might have encoded, which, of course, is exactly what modularity denies. But as we will show, there is a sense in which semantic representation need not invoke semantic decomposition, and thus, can constitute the level of representation the module outputs.
- 8. The final "stage" in this model, *the real world model*, is distinguished from *the discourse model* on the assumption that one ought to construct a representation of the (linguistic) discourse before checking it against the hearer's knowledge or an "aggregate picture representation of how things are" (p. 9). The discourse model is the wider-scope linguistic representation of the sentence, which prevails even when it conflicts with real-world knowledge. We assume that information contained within the discourse model can be conceived as being intra-modular while the real world model cannot.
- 9. Notice that on one analysis (Davidson, 1967) of (3a), the verb *begin* introduces a variable—say, *w*—which in (3a) ranges over not the action/event *x* began doing with *y*, but *begin* itself, thus yielding something like ∃*w* (*begin* (x, y, w)).
- 10. A view similar to this one has been proposed by Parsons (1990; see also Pietroski, 2015).
- 11. It is important to note that de Almeida et al. (2016) employed a different method, materials, and analyses, thus it did not constitute an attempt to replicate Husband et al. (2011) also because the data collection predates the publication of this latter study.

12. Traxler et al.'s (2005) findings, however, are difficult to interpret given the inconsistent results between and within experiments—both in terms of region, where effects are found and in terms of eye-tracking measures that yield the effects. Moreover, many of their statistical analyses—including some that are taken to support their views—are "tendencies," not statistically significant results. And although their results are offered in support to coercion, they can also be claimed to support the perspective we take.

#### REFERENCES

- Altmann, G., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, *30*(3), 191–238.
- Baggio, G., Choma, T., Van Lambalgen, M., & Hagoort, P. (2010). Coercion and compositionality. *Journal of Cognitive Neuroscience*, 22(9), 2131–2140.
- Bornkessel, I., & Schlesewsky, M. (2006). The extended argument dependency model: A neurocognitive approach to sentence comprehension across languages. *Psychological Review*, *113*(4), 787–821.
- Cappelen, H., & Lepore, E. (2005). Insensitive semantics: A defense of semantic minimalism and speech act pluralism. Oxford, England: Wiley-Blackwell.
- Crain, S., & Steedman, M. (1985). On not being led up the garden path: The use of context by the psychological parser. In D. R. Dowty, L. Karttunen, & M. Zwicky (Eds.), *Natural language parsing* (pp.320–358). Cambridge, England: Cambridge University Press.
- Davidson, D. (1967). The logical form of action sentences. In D. Davidson (Ed.), *Essays* on actions and events (pp. 105–122). Oxford, England: Oxford University Press.
- de Almeida, R. G. (2004). The effect of context on the processing of type-shifting verbs. *Brain and Language*, *90*(1-3), 249–261.
- de Almeida, R. G., & Dwivedi, V. (2008). Coercion without lexical decomposition: Type-shifting effects revisited. *Canadian Journal of Linguistics*, 53(2), 301-326.
- de Almeida, R. G., & Manouilidou, C. (2015). The study of verbs in cognitive science. In R. G. de Almeida & C. Manouilidou (Eds.), *Cognitive science perspectives on verb representation and processing* (pp. 3–39). New York, NY: Springer.
- de Almeida, R. G., & Riven, L. (2012). Indeterminacy and coercion effects: Minimal representations with pragmatic enrichment. In A. M. Di Sciullo (Ed.), *Towards a biolinguistic understanding of grammar: Essays on interfaces* (pp. 277–302). Amsterdam, Netherlands: John Benjamins.
- de Almeida, R. G., Riven, L., Manouilidou, C., Lungu, O., Dwivedi, V. D., Jarema, G., & Gillon, B. A. (2016). The neuronal correlated of indeterminate sentence in interpretation: An fMRI study. *Frontiers in Human Neuroscience*, *10:614*. doi: 10.3389/ fnhum.2016.00614.
- Fodor, J. A. (1983). The modularity of mind. Cambridge, MA: MIT Press.
- Fodor, J. A. (1990). On the modularity of parsing: A review. Unpublished manuscript, Rutgers University.
- Fodor, J. A., & Lepore, E. (2002). The emptiness of the lexicon: Reflections on Pustejovsky. In J. A. Fodor & E. Lepore (Eds.), *The compositionality papers* (pp. 89–119). Oxford, England: Oxford University Press.

- Gallistel, C. R. (2018). The neurobiological bases for the computational theory of mind. In R. G. de Almeida & L. Gleitman (Eds.), *On Concepts, Modules, and Language*. Oxford, England: Oxford University Press.
- Gallistel, C. R., & King, A. P. (2010). *Memory and the computational brain: Why cognitive science will transform neuroscience*. New York, NY: Wiley-Blackwell.
- Haugeland, J. (1981). Semantic engines: Introduction to mind design. In J. Haugeland (Ed.), *Mind design: Philosophy, psychology, artificial intelligence* (pp. 34–50). Cambridge, MA: MIT Press.
- Husband, E. M., Kelly, L. A., & Zhu, D. C. (2011). Using complement coercion to understand the neural basis of semantic composition: Evidence from an fMRI study. *Journal of Cognitive Neuroscience*, 23(11), 3254–3266.
- Jackendoff, R. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford, England: Oxford University Press.
- Kuperberg, G. R., Choi, A., Cohn, N., Paczynski, M., & Jackendoff, R. (2010). Electrophysiological correlated of complement coercion. *Journal of Cognitive Neuroscience*, 22(12), 2685–2701.
- Manouilidou, C., & de Almeida, R. G. (2009). Canonicity in argument realization and verb semantic deficits in Alzheimer's disease. In S. Featherston & S. Winkler (Eds.), *The fruits of empirical linguistics* (pp. 123–149). Berlin: Mouton de Gruyter.
- Manouilidou, C., de Almeida, R. G., Schwartz, G., & Nair, N. V. (2009). Thematic roles in Alzheimer's disease: Hierarchy violations in psychological predicates. *Journal of Neurolinguistics*, *22*(2), 167–186.
- McElree, B., Traxler, M., Pickering, M., Seely, R., & Jackendoff, R. (2001). Reading time evidence for enriched composition. *Cognition*, 78(1), B17–B25.
- Parsons, T. (1990). *Events in the semantics of English: A study in subatomic semantics.* Cambridge, MA: MIT Press.
- Pickering, M. J., McElree, B., & Traxler, M. J. (2005). The difficulty of coercion: A response to de Almeida. *Brain and Language*, 93(1), 1–9.
- Pietroski, P. M. (2015). Lexicalizing and combining. In R. G. de Almeida & C. Manouilidou (Eds.), Cognitive science perspectives on verb representation and processing (pp. 43-66). New York, NY: Springer.
- Piñango, M. (2006). Thematic roles as event structure relations. In I. Bornkessel, M. Schlesewsky, & A. Friederici (Eds.), Semantic Role Universals and Argument Linking: Theoretical, Typological, and Psycholinguistic Perspectives. Berlin, Germany: Mouton.
- Pustejovsky, J. (1995). The generative lexicon. Cambridge, MA: MIT Press.
- Pylkkänen, L. (2008). Mismatching meanings in brain and behavior. *Language and Linguistics Compass*, 2(4), 712–738.
- Pylkkänen, L., & McElree, B. (2007). An MEG study of silent meaning. *Journal of Cognitive Neuroscience*, 19(11), 1905–1921.
- Pylyshyn, Z. W. (1984). Computation and cognition: Toward a foundation for cognitive science. Cambridge, MA: MIT Press.
- Riven, L., & de Almeida, R. G. (2017). *Context breeds enriched interpretations of indeterminate sentences*. Manuscript submitted for publication.
- Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, 3(03), 417–424.
- Stojnić, U. (2016a). One's modus ponens: Modality, coherence and logic, *Philosophy and Phenomenological Research*, doi:10.1111/phpr.12307.

- Stojnić, U. (2016b). *Context-sensitivity in a coherent discourse*. Unpublished PhD Dissertation. Rutgers University, Piscataway, NJ.
- Stojnić, U., Stone, M., & Lepore, E. (2014). *Discourse and Logical Form: Pronouns, Attention and Coherence*. Linguistics and Philosophy, doi:10.1007/s10988-017-9207-x.
- Thompson, C. K., & Lee, M. (2009). Psych verb production and comprehension in agrammatic Broca's aphasia. *Journal of Neurolinguistics*, 22, 354–369.
- Townsend, D. J., & Bever, T. G. (2001). Sentence comprehension. Cambridge, MA: MIT Press.
- Traxler, M. J., McElree, B., Williams, R. S., & Pickering, M. J. (2005). Context effects in coercion: Evidence from eye movements. *Journal of Memory and Language*, 53(1), 1–25.