11 Composing Meaning and Thinking

Roberto G. de Almeida

"...take the expression *to weigh one's words*. Now how can you weigh words? When I hear the word *weigh* I see a large scale—like the one we had in Rezhitsa in our shop, where they put bread on one side and a weight on the other. The arrow shifts to one side, then stops in the middle...But what do you have here—to *weigh one's words*?" —S., 1934 (In A. R. Luria, *The Mind of a Mnemonist*, 1968, p. 119)

Solomon V. Shereshevsky or simply "S." was a man with an extraordinary memory. He was able to encode long lists of words, addresses, digits, and events in no time. He was then able to recite them perfectly, even backwards. He would recall long and complex (sometimes random) mathematical formulas and tables full of numbers after just one reading. These virtually intact representations would last for many years. Shereshevsky described his memorization abilities as being entirely based on images. He would "see" everything told to him, or everything he read, as a sequence of images. Sometimes he would place them all on a "street" scene and would "walk" back and forth recalling them in any order. Infrequently missed items-he explainedwould have been mistakenly left in a dark alley, invisible to his mental walk. His interpretations of what was read or heard were illustrated by vivid images representing words, numbers, and phrases. Often he would use synesthesia-associating words or numbers with colors, yielding a comprehension process that was a constant barrage of Ideas in the form of images and colors. But the cost of doing so was a comprehension impairment of sorts: he could not understand most figurative expressions, nor could he make sense of abstract ideas or simple words used metaphorically and in the context of poems. He would struggle with such occurrences of normal language use to a point of being overwhelmed by them.

I read that 'the work got under way normally.' As for *work*, I see that work is going on ... there's a factory. But there's that word *normally*. What I see is a big, ruddy-cheeked woman, a *normal* woman... Then the expression *get under way*. Who? What is all this? You have industry... that is, a factory, and this normal woman—but how does all this fit together? How much I have to get rid of just to get the simple idea of the thing!" –*S*. (Luria, 1968, p. 128)

Putting aside all methodological caveats that interpreting cognitive "impairments" requires, Shereshevsky's case is illustrative of the task at hand: we seem to know what a sentence means when it is interpreted *literally*, but somehow the criteria for interpreting deviations from literal meaning (or simply "meaning") could potentially lead to havoc. Unlike Shereshevsky, however, it looks as though we make sense of such deviations: the absurdity of a metaphor is supposed to be quickly—and perhaps conventionally—corrected to some default interpretation. Indirect requests lead to efficient actions with little effort on the part of the listener. The apparent vagueness of an indeterminate sentence is filled in with appropriate semantic material. We even make sense of aphorisms and "deep thoughts", often peppered with hard-to-grasp analogies. But how do we do it? How do we successfully achieve our communicative goals by such imperfect ways?

It is now more than an article of faith in cognitive science that compositionality is a key characteristic of human cognitive architecture—that without compositionality there can't be productivity, and that without productivity we would be out to fend for ourselves producing or interpreting novel sentences, decoding objects,¹ or simply having thoughts we never had before. Far from being too dramatic, the picture that emerges of our linguistic and cognitive systems without compositionality is that of Shereshevsky's mind and its struggles with non-literal language, for without compositional meaning we would have just a not-so-vast memory likely incapable of "fitting it all together", even a "simple idea".

Granted these common semantic terms-the likes of "literal", "meaning", "compositionality" and many others-still linger in philosophy and cognitive science with little hope of finding consensus. And, by extension, so does a division of labor between semantics and pragmatics. Of more direct concern in the present essay is where exactly we can draw a line between semantics and pragmatics—or more specifically, between linguistic-driven computations and those of other cognitive systems bearing on meaning and utterance interpretation. With this goal in mind, I plan to examine Lepore and Stone's (2015) survey and hypotheses on semantics and pragmatics. They propose, in particular, that the language faculty contains numerous principles that yield linguistically motivated enriched interpretations, dispensing with much of the mind-reading job that is supposed to constitute pragmatic interpretation. Their analytic work provides a very fertile ground for the investigation of the algorithms and heuristics that guide interpretation towards semantic composition and beyond. My discussion of their approach will focus on a few cases that may pose a challenge to an enriched analysis of several types of linguistic expressions-keeping an eye on the intentional fallacy that plagues semantics. I will discuss so-called "bullshit" sentences, the "good-enough" approach to parsing, and in particular metaphors and so-called "indeterminate" sentences. The choice of topics is not merely for their central role in the debate on

¹ I assume that object representations rely largely on low-level compositional processes based on volumetric parts—as in models such as Biederman's (1987).

COMPOSING MEANING AND THINKING 203

where to draw the line between semantics and pragmatics, but also for convenient proximity, as they have been topics of my own investigation.

To anticipate a bit, my line will be traced close to where Lepore and Stone trace theirs, even if it may seem that their commitment to a variety of linguistic conventions as part of the language faculty appears to render the line blurred. I also take interpretation to be largely determined by information in the linguistic material (plus linguistic-perceptual computations; see below). I agree that linguistic principles are probably richer than we commonly think and that they play a key role in building or at least proposing interpretations. And, finally, I also agree that, to a large extent, "pragmatics merely disambiguates" (p. 94). But I differ from Lepore and Stone's approach in mood, if not in substance: I'm less optimistic about the breadth of linguistic rules that could successfully determine interpretations. Also, I take it that the main role of these rules is to compose meaning, and that, beyond meaning-composition, some heuristics might be applied to interpretation, but that largely holistic processes rule.

In section 1, I discuss some guiding assumptions on cognitive architecture, which constrain the nature of linguistic and cognitive representations and processes-and by implication the conception of the semantics/pragmatics divide I have to offer. The phenomena that I examine in section 2, relying on both linguistic arguments and experimental evidence, suggest that for certain constructions there is an early "literal" process of interpretation followed by a period of uncertainty, indicating that the early linguistic computations produce a "shallow" semantic representation, not a fully enriched one. The cases I discuss, culminating with metaphors and so-called indeterminate sentences, challenge the prowess of linguistic computations for resolving-even suggesting-interpretations. I provide evidence for the availability of true and false propositions computed from the meaning one attains about a sentence. In addition, I propose that "indeterminacy" cannot be resolved linguistically, not at least without appealing to an analytic/synthetic distinction—an appeal that should be avoided. I argue, in summary, that rules for converting linguistic utterances into mental representations bearing on meaning-semantic or pragmatic-are largely rules for disposing higher interpretive mechanisms with a rough compositional meaning, with potential semantic "gaps" being filled by abductive-inferential processes beyond compositionality.

1. Guiding Assumptions

I will start off with several guiding assumptions—some of which will be expanded on in later sections, where I discuss particular empirical cases, and others will simply be taken for granted as guiding my approach to sentence comprehension. A few of these assumptions might be fairly standard while some are certainly the object of numerous controversies in philosophy and cognitive science. I will start with the assumptions that are to be taken for granted throughout. As Napoleon would have said, "On s'engage et puis...on voit."

1.1 Symbols, Computations, and Propositions

I take mental representations to be symbolic, and mental processes to be computations over these symbolic representations. Symbolic representations can be simple or complex, and computations (viz., mental processes) are sensitive to both. A computation can be sensitive to a simple symbol (say, a symbol that stands for a morpheme in natural language) when the symbol can change the course of computations performed on its host symbolic expression. To put in more concrete terms, the difference between (1a) and (1b) might be captured by a system that takes verb types to be coded for their distributional properties—say, that *cut* can enter into both an *inchoative* form as in (1a) or a transitive form as in (1b).

- (1) a. This knife cuts easily
 - b. This meat cuts easily²

If so, the computations performed over these token sentences might be sensitive to the two possibilities, allowing for the parser to be committed either to just one analysis or to both. By contrast, linguistic computations might not be sensitive to distinctions such as those in (2), where both verbs and their arguments have the same distributional properties.

(2) a. The water frozeb. The ice melted

The difference between sentences in (2) is like the one observed in those in (3).

(3) a. The dress is blackb. The dress is blue

Even though their differences in content might be relevant in particular situations, they are, as Grimshaw (2005) observed, linguistically inactive. It is the content of these expressions, however, what matters to a pragmatist—say, the state of the weather in (2), and a difference of opinion or optics, as uttered in (3). As this initial assumption implies, the computations that the linguistic system performs are initially sensitive to structure but not the content of the utterances. Indeed, the semantic representation that, by hypothesis, the system outputs to higher cognitive mechanisms is largely dependent not on the content of utterances, but on their structure—a point to which we'll return shortly.

So far, we have seen examples in which computations are sensitive to particular symbols. Computations are also sensitive to complex symbolic expressions in cases where the structure of symbols might be ambiguous (such as in syntactic ambiguity), in cases where there are long-distance dependencies between elements in a structure,

 $^{^2}$ Just a reminder that (1b) is *transitive*, not intransitive as the linear surface form might imply. Although *meat* appears in the subject position, it is the object of *cut* (see de Almeida & Manouilidou, 2015, for discussion).

COMPOSING MEANING AND THINKING 205

such as between syntactically specified constituents and lexical ones, and in cases of analytic inferences in which, for instance, conjoined expressions entail their parts (e.g., $P\&Q \rightarrow P$).

Moreover, I assume that sentences convey propositions and that propositions are complex symbolic expressions that have constituents (simple symbols) and constituent structure (how symbols are arranged to represent a particular token sentence). Sentences in (1) by hypothesis convey propositions that are different in form, while those in (2) convey propositions that are similar in form. Interpretive processes can be, in principle, operations over propositions, computations over form, not content. Many forms of "enrichment" can be seen, thus, as computations over propositions instigated by linguistic principles encoded symbolically.

Another assumption is that understanding linguistic utterances involves a fair amount of symbolic computations and these computations at least at a linguistic level of analysis are hard-wired, based on rough-and-ready algorithms for outputting standard representations of linguistic utterances. But what happens after these algorithms perform their standard job—or in parallel to them—is at the heart of the distinction between semantics and pragmatics: operations over content are supposedly heuristic and contingent on a wide array of information sources, including *co-text*, wider utterance context, the listener's beliefs, background assumptions, and social norms.

Besides establishing a framework for the discussion that ensues, these explicit commitments to a symbolic-computational architecture help to clear the theoretical vocabulary from ill-posed terms such as "association", "activation", and the "strength" of a representation, all involving mental processes. And they also help clear the vocabulary of *representations*, distinguishing them from, say, "images". This terminological pruning also serves to constrain the nature of the products of linguistic computations and their relations to other representations. For instance, "context", "background knowledge", "common ground" and other such notions can be conceived as sets of propositions rather than unstructured units of activation. What a variety of tropes bring to mind—the thoughts that they evoke—are too to be conceived as propositions, and so are the contents of semantic memory about events and states.³ In summary, operations over token expressions—sentences and the likes of context—can be conceived in terms of algorithms and heuristics computed over propositions.⁴

There is yet another advantage in putting it all in terms of symbols and computations over symbols: the vocabulary of representations and processes can be seen as the same

⁴ This distinction can be misleading, for heuristic processes are also encoded in terms of algorithms the difference being that they do not have to have guaranteed end results. I will refer to "algorithmic" and "heuristic" simply to differentiate between fixed and malleable processes.

³ This propositional view of complex mental representations is certainly not new. Similar views have been around at least since Frege (though not explicitly *qua mental* representation) but in particular in works such as Pylyshyn (1973) and Kintsch (1974), in both incarnations of the "language of thought" hypothesis (Fodor, 1975, 2008), and in psychological theories of text interpretation (see, e.g., McKoon & Ratcliff's, 1992, "minimalist hypothesis" of propositional encoding). The differences between these approaches will not be of concern here.

206 ROBERTO G. DE ALMEIDA

for different input modalities and different sources of information. To wit, it is because we have a common vocabulary of representations for different cognitive systems that we can "talk about what we see," to use McNamara's (1978) expression. By the same token, we can use both representations from linguistically derived elements (e.g., names, pronouns, demonstratives) and representations of visually "grabbed" elements to *refer*. It is only because they employ a common code that both forms of reference can meet and talk to each other.

And finally, the commitment to a symbolic/computational view of how linguistic and cognitive processes run allows for compositionality, and compositionality is a property that only symbolic architectures instantiate.

1.2 Compositionality

Another assumption guiding the present work is that sentence meaning is obtained compositionally, that is, by the meaning of its constituents and how they are structured together to yield propositions. I take this to be the default, the null hypothesis. Much has been said about cases that are supposed to violate compositionality (see, e.g., Pelletier, 2004, for discussion). But the claim here is that compositionality applies to propositions, not necessarily to sentences. Thus, cases such as linguistic disfluencies (*uh, well, ... I... perhaps... you can go!*) can yield unambiguous and complete propositions. Idioms can be taken as lexicalized (thus, fixed) or they can be initially processed as non-idiomatic (*The maid kicked the bucket* is not necessarily parsed as an idiom, not even at the offset of *bucket*).⁵ Most cases that are taken to challenge a strict form of compositionality are deep down cases for which linguistic analyses may yield compositional representations. This is not to suggest that semantics should be bloated with ad hoc analyses so that compositionality is ultimately attained (see below on the perils of intentional fallacy). As I will argue in section 2.3, even copular metaphors (viz., those with a form *x is y*) are compositional.

1.3 Modularity

These architectural commitments are tied to the idea that language is modular—in principle, a cognitively impenetrable faculty of the mind. Here is how these ideas connect: the assumption that linguistic processes are computational depends on some of these computations being insensitive to contextual demands, background knowledge, and beliefs. Language processing relies on compositional meaning being obtained autonomously, at most as a relation between a proposition and information within its co-text. Intra-modular computations ought to be the ones that are algorithmic, the ones that are sensitive not to the content of the expressions it computes, but to their

⁵ Online (i.e., real-time) studies on idiom processing have suggested that idioms are treated as lexicalized forms (Swinney & Cutler, 1979), but also that idioms vary with respect to the point at which they are recognized as such (i.e., at the point at which the meaning of the expression is perceived not to correspond to the meanings of the parts put together), with some idioms being recognized early, while others being recognized late, only at the last word or after (see Cacciari & Tabossi, 1988).

COMPOSING MEANING AND THINKING 207

symbols and how they are structured in expressions. Quite likely, many computations performed by other systems, including higher cognitive ones, are "algorithmic" in that sense (e.g., solving a *modus ponens* problem). And it is also quite possible that some of the computations that the language module performs are heuristic, such as employing parsing strategies. But if language is a module, at least some of its main computations ought to be algorithmic, insensitive to the content of the symbols it computes.

It is an empirical issue as to which representations are computed within the modular system and which ones are not. If we regard many of the computations that are deemed pragmatic to be in fact rule-based syntactic or semantic computations performed by the language faculty, the line that divides modular input analysis from other cognitive systems is traced very high. Lepore and Stone (2015) have indeed assumed that many computations deemed "pragmatic" are products of the language faculty. But their approach also takes the position that this faculty is somewhat penetrable—that it adapts to new conventions over time and that linguistic symbols can affect computations in novel ways, catering to new conventions.

The view of modularity assumed here takes *semantic* computations to be part of the module (see de Almeida & Lepore, 2018, for extensive discussion). This is compatible with Fodor's view of modularity, as he laid out in his monograph:

[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. (Fodor, 1983, p. 90)

The sense of semantics that is taken to be part of the module is not a semantics that is concerned with idiosyncrasies of content, but one that is concerned with linguistically active properties of words and sentences, the ones that affect computation and the structure of propositions. As Chomsky puts it,

[The work often called "natural-language semantics" and "lexical semantics"] can be regarded as part of syntax, but oriented to a different interface and different aspects of language use. Insofar as the relation of rhyme that holds between "chase" and "lace" is based on properties of I-sound, and the relation of entailment that holds between "chase" and "follow" on properties of I-meaning, both fall under syntax, in a traditional sense.

Virtually all work in syntax in the narrower sense has been intimately related to questions of semantic (and of course phonetic) interpretation, and motivated by such questions. The fact has often been misunderstood because many researchers have chosen to call this work "syntax," reserving the term "semantics" for relations of expressions to something extra-linguistic.

(Chomsky, 2000, p. 174)

Along the lines of the constraints above (in 1.1 and 1.2), Chomsky sees semantics also operating over "properties and arrangements of symbolic objects" (p. 174). Of course, the empirical question is exactly which aspects of semantics ("syntax") might constitute part of the module and which ones might constitute "something extra-linguistic".

208 ROBERTO G. DE ALMEIDA

1.4 Processing Constraint

Finally, the distinction between semantics and pragmatics should be sensitive to the nature of the computations that unfold over time—which we can call a "processing constraint" (PC). Although somewhat trivial, this is rarely taken into account in philosophical discussions regarding sentence interpretation. The PC can be rather informally stated as in (4).

(4) Processing a sentence (or other unit of linguistic discourse) relies primarily on identifying constituents and building constituents' structural relations over time, with minimal commitments as to the hierarchical structure of the yet-to-come constituents and with as minimal revisions in structure as possible to attain a representation of the input.

There are several assumptions embedded in (4), of which I will highlight three: (i) The first is that obviously parsing and interpretation occur over time—with representations bearing linguistic and non-linguistic information built incrementally, on a millisecond-by-millisecond basis. This seldom needs further clarification in psycholinguistic circles, even if there is neither agreement on what sorts of representations are built, nor on the relative autonomy of linguistic analyses from non-linguistic ones. The time constraint is important because whatever one builds over time might ultimately influence what one takes to be the "meaning" that one attains. Consider a simple local ambiguity, in a typical "garden-path" sentence such as (5a) and its unambiguous pair (5b).

- (5) a. While Beyoncé was singing the song was playing on the radio
 - b. The song was playing on the radio while Beyoncé was singing

Taken as a whole, there is no ambiguity in either sentence. But over time, the representation of the syntactic structure at the point where the noun phrase *the song* is processed in (5a) might be committed to a transitive reading of *singing* and thus the structure that is built at that point might require syntactic revision. Alternatively, the parser might be committed to an intransitive reading of *singing* or even allow for both possibilities to be initially entertained.⁶ Nonetheless, the partial syntactic analysis may also yield a partial semantic representation, one that might be temporarily at odds with the nature of the event that the whole sentence describes. One may, at any given point, build partial or even full but false propositions compatible with the available (partial) stimulus. For instance, the proposition that *Beyoncé was singing the song* might be available at some point, possibly at the offset of the noun phrase *the song*.⁷

⁶ Most early evidence gathered from this type of sentence showed processing delays at the disambiguating point, suggesting that the parser is temporarily committed to one structure (see Ferreira & Clifton, 1986). More recent models have considered other factors such as thematic properties (by assumption, a *semantic* information) as well as more fine-grained lexical constraints on parsing (see Townsend & Bever, 2001).

⁷ A few observations about this example: First, there is nothing in either (5a) or (5b) that points to what Beyoncé was singing (could have been the national anthem, *Carmen*, or simply a song). Second, I refer to (5b) as being unambiguous, but if ambiguity arises as a function of the structures that the alternating verb

COMPOSING MEANING AND THINKING 209

Indeed Christianson et al. (2001) found some evidence for this, not during but after sentences similar to (5a): about 40 percent of the time, subjects report having read sentences that are compatible with the partial input but which are at odds with the full sentence (e.g., *Beyoncé was singing the song*) while also accepting true probes (e.g., *The song was playing on the radio*). And as we will see in section 2, both true and false propositions computed from the incremental analysis of a sentence might be available in memory for quite some time after the sentence offset.

(ii) A second assumption embedded in (4), which only partially overlaps with assumption (i) above, is that what parsing does in real time is to yield—as rapidly and as efficiently as possible—a representation of the input based on linguistic constituents and syntax. Recall that computations are sensitive to both, properties of the symbols and properties of symbolic expressions. While relations between constituents obey primarily the former, higher-order grammatical principles might also affect the course of computations. For instance, parsing commitments may be determined by principles such as the canonical *position* of constituents and, consequently, by the role these constituents play in the event or state described by the sentence.⁸ Parsing models may vary with regards to how much structure is built in advance and also how much parsing decisions depend on the nature of each constituent (see Townsend & Bever, 2001, for review). In all those circumstances, parsing computations are taken to be linguistically based—even if structural decisions are based on lexically encoded properties rather than on syntactic projection.

(iii) Another assumption embedded in (4) is that both processes of *building* and *revising* structure during potential misparses may also be linguistically driven. That is, potential misparses may be corrected by principles encoded within the parsing system. However, there is nothing in (4) that rules out other re-analyses being driven by higher cognitive mechanisms. For instance, the real-world plausibility of a particular token sentence is not something that the linguistic system is supposed to be checking. It is quite possible that the parser initially takes *meat* in (1b), *The meat cuts well*, to be the instrument or "logical subject" of *cut*. Revisions in assignment of thematic structure or revisions in the role that a constituent plays in a token sentence might as well be entirely driven by world knowledge (viz., that meats are cut and not used for cutting something else). But, what is perhaps most important, those revisions are not computations over the *sentence* itself but over the *proposition(s)*

to sing allows, it is also the case that (5b) is temporarily ambiguous at the offset of *singing*, although this won't have consequences for the parsing of the sentence. Finally, it is also possible that the proposition *Beyoncé was singing the song* is a pragmatic inference from the simpler proposition that *Beyoncé was singing* (see Brewer, 1977). The study by Christianson et al. (2001) also found that probes that are compatible with misparses have greater acceptance if they are pragmatically more plausible.

⁸ Alzheimer's patients show more difficulty with non-canonical [*Theme, Experiencer*] structures than with more canonical [*Agent, Theme*] or even [*Experiencer, Theme*] structures, suggesting that syntactic-semantic mapping may be sensitive to a thematic hierarchy, which takes *Agent* to be canonically assigned to the first noun phrase position. See Manouilidou, de Almeida, Schwartz, & Nair (2009).

that the sentence expresses and might be based on inferences that the proposition triggers—way beyond linguistic analyses.

Recovering the linguistic properties of a given utterance relies primarily on computing its input as such, but it also involves computing covert representations. As an analogy, consider object recognition processes. Recognition begins with a viewercentered representation, that is, a representation from the viewer's perspective. But in order to build an object-centered 3-D model, the viewer needs to rely on both, visible surfaces and encoded canonical representations of objects and parts of objects (see Biederman, 1987, 2013).⁹ This 3-D model is determined by visible properties that are assumed to be non-accidental, that is, unlikely to be an accident of viewpoint (e.g., continuity of lines, orientation of edges based on visible vertices, concave creases marking object parts) together with representations that are stored or computed anew. Occluded surfaces are largely determined by the volumetric properties of object parts.¹⁰

The meaning of a sentence relies to a large extent on what is "visible"—i.e., what is in the acoustic or visual (in case of reading) input, with its "occluded surfaces" emerging from the analysis that the input undergoes: they come in the form of the likes of co-reference assignment between nouns and pronouns, ellipsis of verb phrases, presuppositions, and other rules for building representations of what is heard or read. The processes involved are not that of *construction* but that of decoding plus applying rules and matching incoming constituents with stored representations—in real time.

It should be said that PC does not take compositional and non-compositional processes to be serially arranged; that we build non-compositional interpretations only *after* compositional processes end or halt. It is quite possible that compositional and non-compositional processes (e.g., inferences) work in parallel—that whichever thoughts the first incoming constituents of a sentence evoke go on to trigger other thoughts. But that does not preclude compositional analysis to proceed autonomously. Inferences about what one hears might be computed in real time, but as I will argue, these are not far from what a sentence means—in fact the original *un-enriched* proposition might linger, even if it competes with "enriched" forms of propositional representation.

Adopting PC, as roughly formulated in (4), is important also for designing and evaluating experiments investigating the time-course of events in sentence interpretation. Many studies on metaphors, for instance, rely on off-line experimental techniques such as asking participants to press a button "as soon as they understood each statement" (Bowdle & Gentner, 2005). Clearly, response times to such decisions will be affected not by the moment-by-moment processing but by more general processes of appreciation of the statement, providing little information on the nature of the computations

⁹ Although this view of object recognition is heavily contested, the alternative, view-based models (see Peissig & Tarr, 2007) have yet to determine how object representations are productive, if not by compositional processes akin to those proposed for language.

¹⁰ I am not claiming that these representations are visual *tout court*. They are, at some level, symbolic such as proposition-like expressions specifying geometrical properties of parts of objects and their relations (see, e.g., Pylyshyn, 2003).

COMPOSING MEANING AND THINKING 211

underlying the sentence. As we will see in the discussion of some experiments involving sentence comprehension, what matters is not what happens when, say, a metaphorical sentence has been appreciated in all its beauty or absurdity, but what goes on at the points in which the properties of its constituent symbols might affect the course of its computation.

2. Composing Meaning and Thinking (Fast)

Considerations regarding linguistic and cognitive architecture are important in guiding research on how different representations might enter semantic computations and beyond. In principle, semantic computations are sensitive to local—that is, within the sentence or co-text—properties of its symbols, but insensitive to non-local information such as context and different uses. However, what sort of information semantic computations are sensitive to, it turns out, is an empirical matter. There is no agreement on whether the information that linguistic constituents contribute to their host expressions is fixed, as there is no agreement on the range of possible "occluded" properties of a sentence. For instance, for Pustejovsky (1995), what a word contributes to a sentence is highly sensitive to other constituents as well as to properties of events that its host sentence describes. According to this view, the sense one gets of *window* might change whether one is talking about its *aperture* sense or its *physical object* sense, as in (6a) and (6b), respectively.

- (6) a. Mary jumped through the window
 - b. The window is rotting

Notice that, for Pustejovsky, lexical items encode a vast array of properties about their referents, including how they come into being, their physical properties, their purposes, etc. And different types of information encoded with the lexical item come to the fore in a variety of uses of the item, and communicate different aspects of its referent in the event or state it partakes. If it turns out that computations are sensitive to that level of specificity, there can't be a line between semantics and pragmatics, for computations would have to be sensitive to token utterances, not types. The idiosyncrasies of content (and intention) would have to play a part in determining the meaning of the message.

While the nuances of particular linguistic utterances play a significant role in linguistic communication and social interactions, this view of the semantics–pragmatics divide (or lack thereof) is hardly amenable to investigation. Quine (1953) had already alerted us to the perils of fiddling around with vague notions of meaning, but in particular, that a commitment to analyticity leads to a dead end for semantics.

In this section, I work through some challenging cases aiming to trace the line between semantics and pragmatics. In particular, I attempt to trace the line between context-insensitive (but *co-text* sensitive) expressions yielding propositions, on the semantics side, and largely abductive-inferential processes, on the pragmatics side.

True to the commitments in section 1, my focus will be on the primacy of symbolic computations that might yield a representation of the proposition that a sentence might convey. I will argue from the perspective of the null hypothesis: that compositionality rules, with the meaning of a sentence being composed from its "visible" and its "occluded" surfaces.¹¹ But the "occluded" surfaces, I will argue, have to be linguistically determined. I have little to say about what happens on the pragmatic side—and that is also why I differ from Lepore and Stone in mood. The phenomena that I examine suggest that for certain constructions there is an early "literal" process of interpretation—the building of propositions faithful to the input. The process is not optimal, for it often produces misparses (as we briefly saw above in the discussion of (1) and (5)) and thus yields wrong propositions corresponding to parsing commitments made along the way. Also true to the PC informally stated in (4), propositions are built over time and often (veridical or not) remain in memory past full sentence interpretation and pragmatic enrichment.

Separating linguistically determined computations from "extra-linguistic" ones is certainly difficult, requiring multiple methods, from linguistic analysis to experimental investigation. It is, as Lepore and Stone say, an empirical issue. So I will start by discussing a fallacy often committed in semantic analysis, but also one that is committed in other cognitive science areas, and I will work from there into other cases.

2.1 Intentional Fallacy and "Sneaky" Semantics

The processes of understanding people's utterances can be a lot like those involved in understanding a piece of literary work, even a poem or the lyrics of a song, with their inevitable appeals to figurative language. Like in a dialogue, such works may seem to carry information about the author's state of mind, social situation, and beliefs. Often times we know the author much like we know our interlocutors-their biographies or purposes—and this content can be confounded with the meaning of the message. But these confounding factors may not affect exclusively those who simply consume a literary piece. Even literary critics fall for them. In their attempt to draw a line between literary criticism and the task of reading authors' intentions, Wimsatt and Beardsley (1946) coined the term "intentional fallacy". They wanted literary criticism to be free of the critic's attempt to read the author's mind-what she might have intended to tell us with this or that jargon or verse. We are, rather, supposed to fend for ourselves, interpret the work as we wish-and so should the critic, for the work should stand on its own without appealing to what Wimsatt and Beardsley called the "author's psychology". Here are some relevant passages, which serve as an analogy with Lepore and Stone's project as well as with my goals presently:

Intention is design or plan in the author's mind. Intention has obvious affinities for the author's attitude toward his work, the way he felt, what made him write. (p. 469)

¹¹ See, for instance, Gillon (2008) for many cases of potentially "implicit" arguments, as linguistically motivated constituents.

COMPOSING MEANING AND THINKING 213

(...) One must ask how a critic expects to get an answer to the question about intention. How is he to find out what the poet tried to do? (p. 469)

(...) Judging a poem is like judging a pudding or a machine. One demands that it works. A poem can *be* only through its meaning—since its medium is words—yet it *is*, simply *is*, in the sense that we have no excuse for inquiring what part is intended or meant. (p. 469)

[The meaning of a poem] is discovered through the semantics and syntax of a poem, through our habitual knowledge of the language, through grammars, dictionaries, and all the literature which is the source of dictionaries, in general through all that makes a language and culture. (p. 477)

We could replace the role of "critic" with that of the listener, and "poet" with speaker, and we would have a similar formula: ultimately, what counts for meaning determination by the listener is the linguistic statement of the speaker, the "habitual knowledge of the language", its grammar and lexical properties ("dictionaries"). But that does not deny that literary works are forms of language *use*, just like many of our utterances in real-world situations are. One can certainly understand a poem in its literal form, but it's best appreciated beyond that, in whatever thoughts one might entertain that are triggered by the actual words and sentences that constitute the poem—that's the very point of a literary work (*le plaisir du texte*, as Roland Barthes put it).

Wimsatt and Beardsley called the critics who committed the fallacy "intentionalists". Authors (and artists in general) are often intentionalists, too, when they "explain" their own work. Goethe, as Wimsatt and Beardsley pointed out, proposed what he thought would be "constructive criticism", which included of course an evaluation of the author's intention and whether or not he succeeded in implementing it in the work. While we are often "intentionalists" in our linguistic exchanges—asking for clarifications, making (warranted or false) assumptions, reading into gestures, movements of the eyes and eyebrows—what we start off with are the sentences of the speaker: there can't be clarifications on what one hasn't evaluated *linguistically*.

The intentional fallacy also makes its way into cognitive science methods. In psychology, it has been called "stimulus error" (see Pylyshyn, 2003), and works like this:¹² the researcher assumes that a stimulus has particular properties—say, that a particular sentence has some postulated semantic material that is taken to determine its interpretation. The researcher in fact *knows* the (intended) meaning of the sentence; she knows the message that it is normally supposed to convey in the real world. Then, she attributes these properties to mental representations. It's a very common mistake to assume that the properties one knows (or believes) to exist in a stimulus are actually mentally represented as such. The fallacy is so widespread that it affects how experiments are carried out, as Pylyshyn has shown with regards to mental imagery, with many experimental results being determined not by the computations that the

¹² Although Pylyshyn (2003) refers to work on vision, I am adapting it to research on language.

participant performs on the stimulus properties, but by the knowledge and beliefs that the participant has about the stimulus. It is here that the intentional fallacy of the researcher meets the cognitive penetrability of the experiment (or analysis) she conducts. Shereshevsky's confusions—and Luria's recount—are typical of this situation: even if what Shereshevsky experienced were "vivid images", we cannot assume that these were the underlying culprits of his confusions with literal language—that images are the ultimate forms of mental representations.

Unfortunately, the intentional fallacy permeates work in semantics and, by extension, plays a key role in discussions on where the semantics–pragmatics line should be traced. In semantics, as we saw, the fallacy plays a similar role as that in psychology and involves the postulation of phonologically or syntactically null semantic constituents or analyses that are the product of the knowledge or intention of the semanticist. This *sneaky* semantic strategy, as Cappelen and Lepore (2004) once called it, can be demonstrated in the proposed analyses for sentences such as (7).¹³

- (7) a. One more can of beer and I'm leaving
 - b. A little water in the face and he goes indoors
 - c. His fists were clenched. A word, and he would lose his temper
 - d. A few days more of this and I'll go mad

For Culicover (1970), the sentences in (7) carry an implicit conditional, which does not account for the actual content of the events, but impose on them a logical relation. Culicover says that these sentences carry "considerable amount of semantic material which is unspecified" (p. 368). The supposedly missing semantic material has scope over the overt nominal in the "antecedent" clause of the implicit conditionals, as in (8) (other examples being parallel to these).

- (8) a. [If {somebody throws at me / you give me / I see / you drink / I crush / etc.} one more can of beer, {then} I'm leaving]
 - b. [If {somebody throws at him / he gets / etc.} a little water in the face, {then} he goes indoors]

While it is tempting to take (7a-b) to communicate the content in (8a-b), there is nothing in any of these sentences forcing us to come up with an implicit conditional interpretation; simply put, the first clause does not necessarily work as a logical antecedent of the second clause. One could as well conceive of countless situations that take, say, (7a) to be simply a case of conjunction with temporal order; to wit, its form is Pc Q, rather then Q as a consequence of P. Temporal order of conjoined events does not bring about logical implication, even when it appears to be the case that there is logical consequence. Compare (9a) with (9b).

- (9) a. Mary fell on the sidewalk and hit her headb. Mary hit her head and fell on the sidewalk
 - ¹³ Examples (7a)–(7b) are from Culicover (1970) and (7c)–(7d) are from Jespersen (1909).

COMPOSING MEANING AND THINKING 215

Observe here, again, that the imposed connection between events is simply that of a temporal order, not necessary a *logical* relation between antecedent and consequent. Consider (10).

(10) a. Mary had a car accident and burned her houseb. Mary had a car accident and her house burned

How are we supposed to connect the events in (10a)? Are we supposed to attribute the burning of the house to the car accident? To Mary's rage or frustration? While the story we impose on (10a) is of that kind, in (10b) this story needs to be changed, even if the conjoined events are the same as those in (10a):¹⁴ Should we now make up a story about how incredibly unlucky Mary is? If so, this story serves to account for (9) as well, because Mary may have suffered both incidents, falling on the sidewalk and hitting her head many hours apart on an unlucky day.

One way to conceive of sentences in (7) as carrying an implication, while also accounting for the intuitive differences between sentences in (9) and (10), is to assume that '*and*' is ambiguous. Johnson-Laird (1967, 1969) and others¹⁵ have pointed out that '*and*' could be ambiguous between the simple logical conjunction (*and*₁), a temporal (*and*₂; or '*and*₁ *subsequently*'), and a causal (*and*₃; or '*and*₁ *consequently*') sense. If so, sentences in (7)–(10) would be interpreted according to different senses—say *and*₃ for (7), *and*₂ for (9) and (10). But, as Johnson-Laird observes, '*and*' would still need a "setting" to disambiguate between its different senses. And even in cases such as (11a), where '*and then*' is overtly temporal (thus, equivalent to *and*₂), the conjoined events can be reversed, as in (11b) (examples from Johnson-Laird, 1969).

(11) a. The man was throwing the stick and then the dog was retrieving itb. The dog was retrieving the stick and then the man was throwing it

But clearly other factors are at play in (11). The aspectual, iterative properties of the conjoined events allow for an interpretation in which both events occur independently but linked. If so, there is no need to postulate an ambiguity for *'and'*. It might be the case that we have a univocal *'and'*, with the *content* of the conjoined events determining whether they are in temporal, causal, or simple constituency.

While it is possible to conceive of the semantic computations in sentences such as (7) to carry information about an implicit conditional, it does not seem to be the case that their enrichment is linguistically motivated, for they cannot dispense with an evaluation of their content and context. The "considerable semantic material" they require (such as in (8)), as Culicover put it, is not specified either. More likely, we impose on conjoined events a logical structure upon thinking about their possible

¹⁴ Notice that the second clause in both sentences may convey different propositions, even if in (10a) the event is causative and, in (10b), inchoative. In (10a) Mary may have burned her housed unintentionally, just as in (9b), but only in (9b) she may not have anything to do with the burning of her house.

¹⁵ According to Johnson-Laird, similar proposals have been around at least since XIX century lexicographers.

relations. Just as the examples in (10) suggest, the relations between events require a background story—one that is not semantic but pragmatic. If we are supposed to count on pragmatics to disambiguate the sentences, then there is no need to postulate an analysis that enriches sentences semantically.

2.2 "Bullshit" Sentences and "Good Enough" Interpretations

There are numerous cases in which pragmatic enrichment may seem to be required, cases for which there is no clear semantic path to the computation of propositions. Let us assume that the language faculty is truly rich in principles that, for most unmarked cases, determine the course of interpretation. A system this rich should be efficient in interpreting "prophetic" or "mystic" sentences following the same principles it deploys in the analysis of more mundane ones. Strictly speaking, the system that computes syntax and that, by hypothesis, computes the meaning of a sentence, should not be the one to judge whether or not a grammatical sentence is good or bad, conveying superficial ideas or deep thoughts. It shouldn't be able to determine whether or not words are being used metaphorically, except when processes of interpretation halt: as when a transitive verb is used intransitively, a pronoun lacks antecedent, a definite noun phrase suddenly appears in discourse, and many other cases that might derail local computations.

The case of metaphors appears to be one for which we have no way of determining interpretation from lower parsing analyses. It appears—as Lepore and Stone (2015) suggest—that appealing to linguistic conventions won't do. In section 2.3, I discuss cases of simple copular metaphors (*My lawyer is a shark, Juliet is the sun*). But before I do that, let's look at sentences that express a juxtaposition of aphorisms, metaphors, analogies, as well as literal statements: cases of so-called "pseudo-profound bullshit" (henceforth, BS), such as the statements in (12) (Pennycook et al., 2015).

- (12) a. Every material particle is a relationship of probability waves in a field of infinite possibilities. You are that.
 - b. Matter is the experience in consciousness of a deeper non-material reality.

c. Our minds extend across space and time as waves in the ocean of the one mind.

d. We are non-local beings that localize as a dot then inflate to become non-local again. The universe is mirrored in us.

These statements—from the Twitter account of Deepak Chopra—together with others, randomly generated by sites that use similar vocabulary to produce "Chopra-like" statements, were given by Pennycook et al. to a group of subjects to rate for "profoundness" (defined as "of deep meaning; of great and broadly inclusive significance"). These statements were considered on average 55 percent profound (in their 1 to 5 scale, that's 2.77 or in between "somewhat profound" and "fairly profound"). Perhaps unsurprisingly, participants who showed greater BS acceptance were also

COMPOSING MEANING AND THINKING 217

more religious, had greater propensity to make ontological confusions (e.g., relations between the material and the immaterial world), and were more susceptible to "epistemically suspect" claims (such as the existence of angels).

Clearly, accepting BS as "deep" is related to having the talent to suspend reality or to endorse an "alternative reality". But what is interesting about these examples is that they further the idea that appreciation of BS—perhaps along the lines of several forms of figurative expressions—is a cognitive exercise. However, this exercise seems to be triggered by a failure to interpret what lower-level linguistic computations deliver to higher-level interpretive systems. Crucially, the failure or success in understanding deep BS must rely on some prior form of linguistic analysis. Although Pennycook et al. did not investigate real-time comprehension of BS statements, one can surmise that participants computed the propositions that the statements initially yielded—that, as in (12d), *We are non-local beings*, that *we localize as a dot*, that *we inflate*... etc. This is necessary to obtain the representations that participants are then able to elaborate on possible connections with other propositions to make a "profoundness" judgment. It is only by composing the propositions that the expressions yield that participants are free to think whatever they want about the expressions.

The idea that BS sentences might allow for multiple veridical and false interpretations, however, might be challenged by some pervasive effects in the literature, which suggest that rather than "deep", interpretations are superficial. Consider the questions in (13), discussed below.

(13) a. How many animals of each kind did Moses take on the ark?b. What was the nationality of Thomas Edison, inventor of the telephone?

For a question such as (13a), 81 percent of the participants who had the knowledge that the tale was about Noah responded "two" (Erickson & Mattson, 1981). A question such as (13b) produced a weaker but still significant effect (44 percent). While it is possible that this type of error might be due to a misinformation effect, it is clear that subjects composed the meanings of the questions, perhaps not attending to important points about their content. However, this effect—known as the "Moses illusion"—led many to believe that what is going on is a form of "shallow" processing: subjects are not engaging in full compositional process and thus interpret the questions without actually linking the pieces of meaning and grammar with the knowledge of the events that are presupposed.

A more subtle effect of this general illusion can be obtained with simple cases such as the one in (14a).

- (14) a. The dog was bitten by the man
 - b. The man was bitten by the dog
 - c. The man bit the dog
 - d. The dog bit the man

Ferreira and colleagues (e.g., Ferreira, 2003; Christianson et al., 2001) conducted several studies showing that when people are presented with sentences such as (14a), about 25 percent of the time they respond that the "doer" of the action is the dog—an error rate that is much higher than with other sentences such as (14b) or (14c). The theory that Ferreira has been advancing assumes that, for many cases, there is a form of normalization that has its origins not on a full parsing of the sentence but on a superficial one. Ferreira and colleagues term this sort of processing effect "good enough".

However, a recent study by Riven (2017; see also Riven & de Almeida forthcoming a) found that when the task (identifying the "doer" in a sentence such as (14a)) involves a memory load interference (subjects have to keep in mind a series of digits during sentence presentation and response), native-language (L1) speakers of English do better than second-language (L2) English speakers (native speakers of French) performing the same task. Without the memory load interference, L2 speakers do better than L1 speakers. One hypothesis is that L1s make more errors than L2s in the simple version of the task (as in Ferreira's 2003) because L2s rely on explicit, metalinguistic knowledge of sentence properties. Conversely, L1s do better in the task with the memory load interference because they are able to tap into their native, implicit interpretive mechanisms to make a judgment, while L2s' conscious judgments are disrupted by the memory task. What the original task probes, then, is not low-level compositional processes but metalinguistic judgments, which are more prone to interference and, hence, errors. If this is right, the compositional representation stays intact and is not fooled by the "pragmatic normalization" (Fillenbaum, 1974) that might occur in evaluating what the man did to the dog in (14a).

In summary, it is quite possible that compositional representations are fully processed but that inattention or processing bottlenecks might lead to further interpretation errors. This effect might also help us understand acceptance of BS sentences: underlying compositional mechanisms deliver consistent, grammatical compositions of the sentence, which are disrupted by their odd content. These errors of interpretation are not necessarily foreign to what one attains of a sentence (as per PC). False propositions remain in long-term memory, long after a sentence is processed—and so do *un-enriched* propositions, as we will see in 2.4.

2.3 Minimal Metaphors

BS statements represent the extreme case of figurative expressions, requiring more than analogical reasoning to appreciate the intricacies they appear to convey. But they are not far from copular metaphors such as (15) in expressing blatantly false or absurd relations between constituents.

- (15) a. Roads are snakes
 - b. My lawyer is a shark.

COMPOSING MEANING AND THINKING 219

c. Memory is a heap of broken images where the sun beats and the dead tree gives no shelter¹⁶

In fact, some BS statements such as (12a) and (12b) have the same *x* is *y* form as those in (15) and could be treated just as a case of metaphor. In contrast with BS statements, however, the metaphor kind I will discuss here is more conventional in expressing a relation between a "topic" (e.g., *lawyer*) and a "vehicle" (*shark*). The question I am interested in is whether there are *linguistic* principles driving metaphorical interpretation and, beyond that, whether there are particular cognitive principles deployed in metaphor appreciation. This is perhaps where I differ from Lepore and Stone's approach in substance rather than just in mood. For them, "there are no linguistic cues and no linguistic reflexes for the insights that the speaker is offering" (p. 164). I will argue, however, that there is, at a minimum, syntactic clues suggesting how a metaphor ought to be interpreted initially. Moreover, I will suggest that there are further variables that affect how the content of a metaphorical expression is explored. The caveat, however, is that the "clues" are fragile and might not be, after all, true conventions in the sense they adopted.

While there is general agreement that lawyers are not sharks, we are far from any agreement on what one does with an expression like (15) during or after its first parse. Some researchers have postulated that interpretation of metaphors might be "direct" (Gibbs & Colston, 2012; Glucksberg, 2003), but this postulation is based mostly on off-line tasks indicating that metaphors don't differ in acceptance time from similarly structured but literal statements. In addition, there is no agreement on how exactly the metaphor produces its effect. Some have assumed that the process of interpretation itself is a process of "domain mapping" (Bowdle & Gentner, 2005) creating analogies, or that it involves novel categorizations (Glucksberg, 2003), or the creation of *ad hoc* concepts (Carston, 2010).

Lepore and Stone (2015) assume that metaphor interpretation is not determined linguistically, but they also seem to agree with a view such as that of Bowdle and Gentner (2005) for whom metaphor interpretation ultimately relies on processes of building analogies between predicates. As Lepore and Stone propose, "speakers and their audiences explore these analogies in open-ended ways" (p. 163). But the building of analogies must rely on content upon which analogies are formed. One could think of sets of inferences that different entities trigger—like meaning postulates (see de Almeida, 1999). But in Bowdle and Gentner's perspective, metaphor interpretation relies on specific mechanisms involving "alignment" between predicates and the "projection" of elements from the vehicle to the topic. Their mapping processes are based, to a large extent, on semantic decomposition: one decomposes the meaning of both topic and vehicle and compares the properties they yield by aligning these properties

¹⁶ See Katz et al. (1988) and Roncero & de Almeida (2015) for corpora of metaphors and their semantic properties.

220 ROBERTO G. DE ALMEIDA

in terms of corresponding predicates. Thus, an expression such as (16a) would yield an interpretation roughly such as (16b), which could be read as in (16c).

(16) a. Socrates was a midwife

b. [HELP [Socrates [PRODUCE [Student, Idea]]] and [HELP [Midwife [PRODUCE [Mother, Child]]] & [GRADUALLY [DEVELOP [WITHIN [Child, Mother]]]]]].¹⁷

c. "Socrates did not simply teach his students new ideas but rather helped them realize ideas that had been developing within them all along." (Bowdle & Gentner, 2005, p. 196)

It seems, however, that this is not the kind of analogy Lepore and Stone would want to be committed to—for one, they would have to handle the heavy analytic baggage that Bowdle and Gentner carry. It is important to distinguish the view that metaphors create analogies (perhaps among many other kinds of propositional content) from the view that the *semantic* representation of a metaphorical expression is something like a simile, as in (16).

- (17) a. Roads are like snakes
 - b. My lawyer is like a shark

c. Memory is like a heap of broken images where the sun beats and the dead tree gives no shelter

Notably, metaphors are not similes (contra Aristotle), for they convey different meanings. For Bowdle and Gentner, metaphors have a career, starting as similes—and, thus, being interpreted, relying on analogies such as in (16)—but later turning into metaphors as their vehicles become more conventionalized.¹⁸ Thus, the more *shark* is used in different expressions to denote some set of figurative properties, the more conventionalized it becomes and the more it occurs in metaphors. But a recent study of written posts on the Internet, suggests that metaphors and similes with the same topic–vehicle pairs (such as (15a) and (17a)) occur with the same frequency. Moreover, these posts use explanations significantly more with similes than with metaphors (*Time is like money—because only retired executives have a lot*) suggesting that they are used to convey ideas that call for a supporting co-text (see Roncero et al., 2016; Roncero, Kennedy, & Smyth, 2006).

These novel interpretive explorations are triggered before the explanations come into place, before the absurdity of a metaphor is detected. We have put forth a possible linguistic difference between metaphors and similes that lines up with the idea that

¹⁷ This is my notation capturing their "parallel connectivity" graphs. For simplicity, I omitted other ontological categories such as "instrument", "action", and "object".

¹⁸ The notion of "convention" used by Bowdle and Gentner (2005) is different from the linguistic conventions of Lepore and Stone (2015).

COMPOSING MEANING AND THINKING 221

linguistic conventions might be at work in determining the interpretation of a simile or a metaphor expression containing the same topic–vehicle pairs (de Almeida et al., 2011). One hypothesis is that an expression of the form *x is y* is predicative—simply put, that *y* predicates something of *x*, with a form in which *e* (an *entity*) is taken to be the topic to which the predicative type $\langle e, t \rangle$ applies, as in (18a).

(18) a. BE (e (<e,t>))
b. BE (LIKE (e (e)))

This hypothesis requires committing to an ontology of semantic types, something one might not be willing to do. But at a minimum one does not have to resort to typeshifting operations (Partee, 1986) to account for the difference between the two expressions, for the type of internal argument is determined structurally: the predicate determines the nature of its internal argument. Crucially, if one is committed to exhausting linguistic resources before committing to a pragmatic-level interpretation, this view assumes that the interpretation of copular metaphors is linguistically determined and that their computations might be different than those involved in interpretation of similes.

Besides their linguistic differences—their potentially different semantic types (*entity* v. *predicate*)—similes and metaphors also seem to be processed differently in real time. In an eye-tracking study involving similes and metaphors such as those in (15a–b) and (17a–b) (Ashby et al., 2018), we found that metaphors take longer to read at the vehicle position (see Figure 11.1), and that metaphors trigger twice as many regressive saccades towards the vehicle from their accompanying explanations than similes do.



Figure 11.1. Reading times on the vehicles (*shark*) of metaphors (*lawyers are sharks*) and similes (*lawyers are like sharks*) in the study by Ashby et al. (2018, Experiment 1). Bars represent go-past time (total time reading a region before moving to the next one); error bars are standard errors.

222 ROBERTO G. DE ALMEIDA

These eye-tracking data, together with data from Internet searches and linguistic arguments, seem to indicate that similes are not metaphors. And if they are different linguistic expressions, yielding different types of propositions, interpretive mechanisms that operate on their linguistic outputs might also differ. The key point I am making is that copular metaphor interpretation might be *linguistically determined*—not in all its beauty or absurdity, but minimally in the predicative relation between its arguments, as its form suggests.

In summary, the interpretation of metaphors might rely on an early parsing that promotes the predication of the vehicle to the topic—predicates *sharkness* of lawyers. This is what "syntax" does, allowing for further interpretive mechanisms to elaborate on this predicative relation—perhaps by calling for "properties" (viz., other predicates) that might be related to *shark*. Many of these properties might be driven by what Roncero and colleagues have called "aptness"—how well the predicates computed from a vehicle apply to the topic (see Roncero et al., 2006; Roncero & de Almeida, 2014). The interpretive task might simply rely on that, after the linguistic system has done its job. In the next section, I discuss a case in which semantics produces an unenriched interpretation that lingers, even in strongly biasing utterance contexts.

2.4 Indeterminacy

The final case I would like to discuss is that of sentences such as those in (19).

(19) a. I finished the chapterb. Mary began a book

The phenomenon that these sentences characterize is usually called "coercion", "typeshifting", or "type-coercion" in the semantics literature (see de Swart, 2011, for a review). I will adopt the more neutral term "indeterminacy"¹⁹ because coercion is tied to a particular *hypothesis* on how these sentences might be enriched—namely the hypotheses that complement noun phrases are "coerced" to be interpreted as an event, or "coerced" to provide information about an event that fills-in the sentence interpretation. This is so because, intuitively, these sentences are taken to be interpreted as in (20).

- (20) a. I finished [writing/reading/typing] the chapter
 - b. Mary began [writing/reading/reviewing/eating] a book

The phenomenon of interest here is that, even though it is not clear what sort of activity Mary began doing with a book, we seem to assign default enriched semantic representations to this kind of sentence (e.g., reading or writing) amongst several possible—but perhaps less plausible—interpretations (e.g., eating, burning), such as in (20).

This type of sentence has received some attention in the theoretical literature in cognitive science (e.g., Briscoe, Copestake, & Boguraev, 1990; de Almeida & Dwivedi, 2008;

¹⁹ Although the label "indeterminate" might be applied, it should be clear that these sentences are fully grammatical and that they allow for a truth judgment (viz., they are true if, say, Mary began doing anything with a book).

COMPOSING MEANING AND THINKING 223

de Almeida & Lepore, 2018; Fodor & Lepore, 2002; Pustejovsky, 1995; see also de Swart, 2011, for review). And it has also been the subject of numerous experiments involving a great variety of psycholinguistic techniques and neuroimaging methods (e.g., de Almeida, 2004; de Almeida et al., 2016; Husband, Kelly, & Zhu, 2011; Katsika et al., 2012; McElree et al., 2006; McElree et al., 2001; Pickering, McElree, & Traxler, 2005; Traxler, Pickering, & McElree, 2002; Pylkkänen & McElree, 2007). Given the range of issues involved, I will focus on one particular theoretical issue and limit the discussion of empirical results to one study (see de Almeida et al., 2016, for a recent review).

The key issues under dispute include the nature of semantic composition and the linguistic and cognitive resources involved in resolving (or attempting to resolve) indeterminacy. Perhaps the first to discuss this issue in print was Culicover (1970) who wrote that sentences such as those in (19) contain an "infinite ambiguity" (p. 368) and, similarly to those in (7), require a "considerable amount" of semantic material that is neither overtly expressed in the sentence, nor linguistically (i.e., syntactically) motivated.

In Pustejovsky's (1995) theory, cases such as (19b) typify the context sensitivity of the verb-complement composition: *begin* is supposed to require *book* to provide information about what *begin the book* might be about. This information is supposed to come from what is stored with the lexical entry for the complement noun. As mentioned above, in his theory, nouns carry information such as *roles* or *functions*, how they *come into being*, what they are *made of*, etc.²⁰ This information enters into the compositional process yielding (20b), given (19b).

While some early experiments have shown processing delays in post-verbal constituents (say, after *began* in (19b)), compared to a "non-coerced" sentence (e.g., *Mary read a book*; see McElree et al., 2001), the nature of this difference has been explained by different theories. Our position—*contra* coercion—is that this type of sentence is indeterminate with regards to the event, and it keeps its "infinite ambiguity" at the linguistic level of analysis. However, its pragmatic enrichment—if any—comes not from lexical decomposition as proposed by Pustejovsky, nor from some form of typeshifting (see Pylkkänen, 2008), but by its own more complex syntactic structure—as in (21) (see de Almeida & Dwivedi, 2008; and de Almeida & Riven, 2012 for linguistic evidence).

(21) $[VP [V^0 began [V^0 e [OBJ NP]]]]$

The proposal is that the syntactically specified "gap" is tagged at logical form (LF) and that it is a "trigger" for pragmatic inferences bearing on the content of the event.

It is the job of the context—or even the *co-text*—to propose information that may help "disambiguate" the sentence. We found evidence for these sentences tapping extra

²⁰ In Pustejovsky's theory these are "quale" structures, which are part of the lexical entry and constitute the basic information about the content of a word and thus how it might vary its contribution to contexts.

224 ROBERTO G. DE ALMEIDA

processing resources and, crucially, engaging brain areas that are involved in pragmatic processes of conflict resolution (such as the anterior cingulate cortex and temporal structures in the right-hemisphere; see de Almeida et al., 2016).

Finally, there is evidence that these sentences—although possibly enriched past their initial linguistic analysis—yield dual-representations, one proposition that is true to the initial indeterminate token sentence, and one that is enriched over time, with both propositions lingering several seconds after the initial presentation. Evidence for this hypothesis comes from a study in which subjects are presented aurally with long discourse passages such as (22a) preceding a sentence such as (22b) (Riven & de Almeida forthcoming b).

(22) 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,

b. Lisa began the book.

At the offset of the sentence or at about 25 seconds later (with more intervening neutral discourse presented aurally), subjects are presented visually with either the original sentence (22b) or one of two foils, one that is consistent (23a) and another, inconsistent (23b) with the context.

- (23) a. Lisa began reading the book
 - b. Lisa began writing the book

Subjects perform a sentence probe task (responding "yes" if the presented sentence matches the original—in this case, (22b)—or "no" otherwise). Results showed that while at the sentence offset point subjects perform nearly at ceiling, at the delayed probe point (25 seconds later), they reject the inconsistent probe (23b), but accept the consistent probe (23a) at the same rate as they accept the original sentence (22b): that is, both *began the book* and *began reading the book* in context are accepted at the same rate. Moreover, response times show that subjects hesitate much more in accepting the enriched (23a) than they do accepting the original (22b) sentence. We hypothesized that both acceptance rates and longer response times suggest that propositions computed from the original presentation linger and compete with the contextually consistent foil. These results, in fact, are compatible with other studies on memory for sentences (Brewer, 1977; Brewer & Sampaio, 2006), supporting a dual-propositional representation for linguistic material, a true memory and a false one (Reyna et al., 2016).

We take these results further to support the idea that sentences are *not* enriched when they are initially perceived and that they produce propositions that remain encoded, coming to compete in recognition memory with contextually supported but false propositions. It is quite possible that acceptance of the consistent false probe (23a)

COMPOSING MEANING AND THINKING 225

is warranted by the context and that it is part of what subjects understand of the sentence. But underlying that representation, the true proposition lingers.

3. Coda: Indeterminacy Rules

Just like Shereshevsky's comprehension troubles suggest, when words and sentences deviate from what they actually *mean* it may be difficult to fit them all together. But we don't possess many of Shereshevsky's traits. For one, we always find a way of making sense of what we hear—if not by the workings of the compositional linguistic system alone, by thinking about its products. We make them fit with all the available information at the time of the utterance—a virtually unlimited and unconstrained array of information.

In summary, after all is said and done about linguistic analysis, indeterminacy about what is not resolved early on will rule. It is even possible that there are actual principles—more likely heuristic ones—that take care of what is not clear, or what is considered anomalous, or what is pseudo-profound "BS", "good enough", or indeterminate, because linguistic analyses do not exhaust the thoughts that a sentence might provoke in the listener. What people ponder about utterances draw a lot from context, but no matter how rich or poor the context is, the literal representation lingers.

As the dialogue from *Seinfeld*, in (24), well depicts, reading intentions is not necessarily the only way people interpret utterances: a representation of the literal interpretation is always under consideration, even if all the clues for alternative interpretations suggest otherwise.

(24) George:—She invites me up at 12 o'clock at night for "coffee"... and I don't go up. "No, thank you, I don't want coffee, it keeps me up. It's too late for me to drink coffee." I said this to her. People this stupid shouldn't be allowed to live. I can't imagine what she must think of me.

Jerry:-She thinks you are a guy that doesn't like coffee.

George:-She invited me up. "Coffee" is not coffee; "coffee" is sex.

Elaine:-Maybe "coffee" was coffee!

George:—"Coffee" is coffee in the morning, not at 12 o'clock at night!

Finally, as we discussed elsewhere (de Almeida & Lepore, 2018) Fodor rightly thought that a semantics committed to lexical decomposition was necessarily holistic—because there is no account of the analytic/synthetic distinction. Alternatively, there is a view of semantics that keeps it tidy within the module: it's a semantics that is atomistic and operates in consonant with lexical-structural properties and syntactic principles. If so, the language faculty is rich indeed, and much of what is deemed pragmatics is already pre-packaged linguistically. Crucially, what it does is to compose meaning—and then there is thinking.²¹

²¹ Research for this article was supported by a grant from the National Sciences and Engineering Research Council of Canada (NSERC). I'm grateful to Caitlyn Antal and Linmin Zhang for comments.

References

- Ashby, J., Roncero, C., de Almeida, R. G., & Agauas, S. J. (2018). The early processing of metaphors and similes: Evidence from eye-movements. *The Quarterly Journal of Experimental Psychology*, 71(1), 161–8.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94(2), 115–47.
- Biederman, I. (2013). Human object recognition: Appearance vs. shape. In S. J. Dickinson & Z. Pizlo (Eds.), Shape perception in human and computer vision (387–97). London: Springer-Verlag.
- Bowdle, B. F. & Gentner, D. (2005). The career of metaphor. Psychological Review, 112(1), 193-216.
- Brewer, W. F. (1977). Memory for the pragmatic implications of sentences. *Memory & Cognition*, 5(6), 673-8.
- Brewer, W. F. & Sampaio, C. (2006). Processes leading to confidence and accuracy in sentence recognition: A metamemory approach. *Memory*, 14, 540–52.
- Briscoe, T., Copestake, A., & Boguraev, B. (1990). Enjoy the paper: Lexical semantics via lexicology. In *Proceedings of 13th International Conference on Computational Linguistics* (pp. 42–7). Helsinki: Association of Computational Linguistics.
- Cacciari, C. & Tabossi, P. (1988). The comprehension of idioms. *Journal of Memory and Language*, 27, 668-83.
- Cappelen, H. & Lepore, E. (2004). Insensitive semantics. Oxford: Blackwell Publishers.
- Carston, R. (2010). Lexical pragmatics, ad hoc concepts and metaphor: A Relevance Theory perspective. *Italian Journal of Linguistics*, 22, 153–80.
- Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.
- Christianson, K., Hollingworth, A., Halliwell, J. F., & Ferreira, F. (2001). Thematic roles assigned along the garden path linger. *Cognitive Psychology*, 42, 368–407.
- Culicover, P. (1970). One more can of beer. Linguistic Inquiry, 1(3), 366-9.
- de Almeida, R. G. (1999). What do category-specific semantic deficits tells us about the representation of lexical concepts? *Brain and Language*, 68(1), 241–8.
- de Almeida, R. G. (2004). The effect of context on the processing of type-shifting verbs. *Brain* and Language, 90, 249–61.
- de Almeida, R. G. & Dwivedi, V. (2008). Coercion without lexical decomposition: Type-shifting effects revisited. *Canadian Journal of Linguistics*, 53, 301–26.
- de Almeida, R. G. & Lepore, E. (2018). Semantics for a module. In R. G. de Almeida & L. Gleitman (Eds.), On concepts, modules, and language: Cognitive science at its core (pp. 113–38). Oxford: Oxford University Press.
- 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: Springer.
- de Almeida, R. G. & Riven, L. (2012). Indeterminacy and coercion effects: Minimal representations with pragmatic enrichment. In A. M. DiSciullo (Ed.), *Towards a biolinguistic understanding of grammar* (pp. 277–301). Amsterdam: John Benjamins.
- de Almeida, R. G., Manouilidou, C., Roncero, C., & Riven, L. (2011). Three tales of semantic decomposition: Causatives, coercion, and metaphors. In A. I. França & M. Maia (Eds.), *Papers in psycholinguistics* (pp. 172–90). Rio de Janeiro: Imprint.

COMPOSING MEANING AND THINKING 227

- de Almeida, R. G., Riven, L., Manouilidou, C., Lungu, O., Dwivedi, V., Jarema, G., & Gillon, B. (2016). The neural correlates of indeterminate sentence interpretation: An fMRI study. *Frontiers in Human Neuroscience*, *10*, 614. https://doi.org/10.3389/fnhum.2016.00614.
- de Swart, H. (2011). Mismatches and coercion. In C. Maienborn, K. von Leusinger, & P. Portner (Eds.), Semantics: An international handbook of natural language meaning (pp. 535–74. Berlin: Mouton de Gruyter.
- Erickson, T. A. & Mattson, M. E. (1981). From words to meaning: A semantic illusion. *Journal of Verbal Learning and Verbal Behavior*, 20, 540–52.
- Ferreira, F. (2003). The misinterpretation of noncanonical sentences. *Cognitive Psychology*, 47, 164–203.
- Ferreira, F. & Clifton, C. (1986). The independence of syntactic processing. Journal of Memory and Language, 25, 348–68.
- Fillenbaum, S. (1974). Pragmatic normalization: Further results for some conjunctive and disjunctive sentences. *Journal of Experimental Psychology*, 102, 574–8.
- Fodor, J. A. (1975). The language of thought. Cambridge, MA: Harvard University Press.
- Fodor, J. A. (1983). The modularity of mind. Cambridge, MA: MIT Press.
- Fodor, J. A. (2008). LOT 2: The language of thought revisited. Oxford: Oxford University Press.
- Fodor, J. A. & Lepore, E. (2002). The emptiness of the lexicon: Reflections on Pustejovsky. In J. A. Fodor & E. Lepore, *The compositionality papers* (pp. 89–119). Oxford: Oxford University Press.
- Gibbs, R. & Colston, H. (2012). *Interpreting figurative meaning*. New York: Cambridge University Press.
- Gillon, B. (2008). On the semantics/pragmatics distinction. Synthese, 165, 373-84.
- Glucksberg, S. (2003). The psycholinguistics of metaphor. *Trends in Cognitive Science*, 7, 92–6. Grimshaw, J. B. (2005). *Words and structure*. Stanford, CA: CSLI Publications.
- Husband, M. E., 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, 3254–66.
- Jespersen, O. (1909). A modern English grammar on historical principles. London: Allen and Unwin.
- Johnson-Laird, P. N. (1967). Katz on analyticity. Journal of Linguistics, 3, 82.
- Johnson-Laird, P. N. (1969). '&'. Journal of Linguistics, 6, 111-14.
- Katsika, A., Braze, D., Deo, A., & Piñango, M. M. (2012). Complement coercion: Distinguishing between type-shifting and pragmatic inferencing. *Mental Lexicon*, 7, 58–76.
- Katz, A. N., Paivio, A., Marschark, M., & Clark, J. M. (1988). Norms for 204 literary and 260 nonliterary metaphors on 10 psychological dimensions. *Metaphor and Symbol*, 3, 191–214.
- Kintsch, W. (1974). *The representation of meaning in memory*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Lepore, E. & Stone, M. (2015). Imagination and convention: Distinguishing grammar and inference in language. Oxford: Oxford University Press.
- Luria, A. R. (1968). *The mind of a mnemonist: A little book about a vast memory*. Cambridge, MA: Harvard University Press.
- Manouilidou, C., de Almeida, R. G., Schwartz, G., & Nair, N. P. V. (2009). Thematic roles in Alzheimer's disease: Hierarchy violations in psychological predicates. *Journal of Neurolinguistics*, 22, 167–86.

228 ROBERTO G. DE ALMEIDA

- McElree, B., Pylkkänen, L., Pickering, M. J., & Traxler, M. J. (2006). A time course analysis of enriched composition. *Psychonomic Bulletin & Review*, 13(1), 53–9.
- McElree, B., Traxler, M. J., Pickering, M. J., Seely, R. E., & Jackendoff, R. (2001). Reading time evidence for enriched composition. *Cognition*, *78*, B17–B25.
- McKoon, G. & Ratcliff, R. (1992). Inference during reading. Psychological Review, 99, 440-66.
- McNamara, J. (1978). How can we talk about what we see? Unpublished manuscript. McGill University.
- Partee, B. (1986). Noun phrase interpretation and type-shifting principles. Reprinted in B. Partee (2004), *Compositionality in formal semantics* (pp. 203–30). Oxford: Blackwell.
- Peissig, J. J. & Tarr, M. J. (2007). Visual object recognition: Do we know more now than we did 20 years ago? *Annual Review of Psychology*, 58, 75–96.
- Pelletier, P. H. (2004). The principle of semantic compositionality. In S. Davis & B. Gillon (Eds.), *Semantics: A reader* (pp. 133–56). Oxford: Oxford University Press.
- Pennycook, G., Cheyne, J. A., Barr, N., Koehler, D. J., & Fugelsang, J. A. (2015). On the reception and detection of pseudo-profound bullshit. *Judgment and Decision Making*, 10(6), 549–63.
- Pickering, M. J., McElree, B., & Traxler, M. J. (2005). The difficulty of coercion: A response to de Almeida. *Brain and Language*, 93, 1–9.

Pustejovsky, J. (1995). The generative lexicon. Cambridge, MA: MIT Press.

- Pylkkänen, L. (2008). Mismatching meanings in brain and behaviour. Language and Linguistic Compass, 2, 712–38.
- Pylkkänen, L. & McElree, B. (2007). An MEG study of silent meaning. Journal of Cognitive Neuroscience, 19, 1905–21.

Pylyshyn, Z. (2003). Seeing and visualizing: It's not what you think. Cambridge, MA: MIT Press.

- Pylyshyn, Z. W. (1973). What the mind's eye tells the mind's brain: A critique of mental imagery. *Psychological Bulletin*, 80, 1–24.
- Quine, W. V. O. (1953). *From a logical point of view* (2nd ed.). Cambridge, MA: Harvard University Press.
- Reyna, V. F., Corbin, J. C., Weldon, R. B., & Brainerd, C. J. (2016). How fuzzy-trace theory predicts true and false memories for words, sentences, and narratives. *Journal of Applied Research in Memory and Cognition*, 5(1), 1–9.
- Riven, L. (2017). The metalinguistic buffer effect: When language comprehension is good but linguistic judgment is only "good enough". Unpublished PhD dissertation, Concordia University.
- Riven, L. & de Almeida, R. G. (forthcoming a). *Metalinguistic switch-costs drive "good enough" misinterpretations of passives*. Manuscript submitted for publication.
- Riven, L. & de Almeida, R. G. (forthcoming b). *Context breeds enriched interpretations of indeterminate sentences*. Manuscript submitted for publication.
- Roncero, C. & de Almeida, R. G. (2014). The importance of being apt: Metaphor comprehension in Alzheimer's disease. *Frontiers in Human Neuroscience*, 8. https://doi.org/10.3389/ fnhum.2014.00973.
- Roncero, C. & de Almeida, R. G. (2015). Semantic properties, aptness, familiarity, conventionality, and interpretive diversity scores for 84 metaphors and similes. *Behavior Research Methods*, 47, 800–12.
- Roncero, C., de Almeida, R. G., Martin, D., & de Caro, M. (2016). Aptness predicts metaphor preference in the lab and on the internet. *Metaphor & Symbol*, *31*, 31–46.

COMPOSING MEANING AND THINKING 229

- Roncero, C., Kennedy, J. M., & Smyth, R. (2006). Similes on the Internet have explanations. *Psychonomic Bulletin & Review*, 13(1), 74–7.
- Swinney, D., & Cutler, A. (1979). The access and processing of idiomatic expressions. *Journal of Verbal Learning and Verbal Behavior*, 18, 523–34.
- Townsend, D. J., & Bever, T. G. (2001). Sentence comprehension: The integration of habits and rules. Cambridge, MA: MIT Press.
- Traxler, M. J., Pickering, M. H., & McElree, B. (2002). Coercion in sentence processing: Evidence from eye movements and self-paced reading. *Journal of Memory and Language*, 47, 530–47.
- Wimsatt, W. K., & Beardsley, M. C. (1946). The intentional fallacy. *The Sewanee Review*, 54(3), 468–88.