

How Can Semantics Avoid the Troubles with the Analytic/Synthetic Distinction?



Roberto G. de Almeida and Caitlyn Antal

One failure a week is just bracing and good for you.
—Allen Newell (1991).

Abstract At least since Quine (From a logical point of view. Harvard University Press, Cambridge, MA, 1953) it has been suspected that a semantic theory that rests on defining features, or on what are taken to be “analytic” properties bearing on the content of lexical items, rests on a fault line. Simply put, there is no criterion for determining which features or properties are to be analytic and which ones are to be synthetic or contingent on experience. Deep down, our concern is what cognitive science and its several competing semantic theories have to offer in terms of solution. We analyze a few cases, which run into trouble by appealing to analyticity, and propose our own solution to this problem: a version of atomism cum inferences, which we think it is the only way out of the dead-end of analyticity. We start off by discussing several guiding assumptions regarding cognitive architecture and on what we take to be methodological imperatives for doing semantics within cognitive science—that is a semantics that is concerned with accounting for mental states. We then discuss theoretical perspectives on lexical causatives and the so-called “coercion” phenomenon or, in our preferred terminology, indeterminacy. And we advance, even if briefly, a proposal for the representation and processing of conceptual content that does away with the analytic/synthetic distinction. We argue that the only account of mental content that does away with the analytic/synthetic distinction is atomism. The version of atomism that we sketch accounts for the purported effects of analyticity with a system of inferences that are in essence synthetic and, thus, not content constitutive.

R. G. de Almeida (✉) · C. Antal
Department of Psychology, Concordia University, Montreal, QC, Canada
e-mail: roberto.dealmeida@concordia.ca

C. Antal
e-mail: caitlyn.antal@mail.concordia.ca

Keywords Analytic/synthetic distinction • Compositionality • Causatives • Indeterminacy • Coercion • Intentional fallacy • Concepts • Atomism • Cognitive architecture

Our concern in this paper is, on the surface, not new. For long—at least since Quine (1953) in modern times, to say little of Kant’s “cleavage” problems way back then—it has been suspected that a semantic theory that rests on defining features, or on what are taken to be “analytic” properties bearing on the content of lexical items, rests on a fault line. Simply put, there is no criterion for determining which features or properties are to be analytic and which ones are to be synthetic or contingent on experience. But that is just the glossy if old shell of our concern. Deep down, our concern is what cognitive science and its several competing semantic theories have to offer in terms of solution, if any at all. With this in mind, we analyze a few cases, which run into trouble by appealing to analyticity, and propose our own solution to this problem: a version of atomism cum inferences. We are aware that the proposal we have to offer is at odds with widely held views, but we think it is the only way out of the dead-end of analyticity, if one is not to be burdened with producing an analytic/synthetic criterion. We start off by discussing several guiding assumptions regarding cognitive architecture and on what we take to be methodological imperatives for doing semantics within cognitive science—that is a semantics that is concerned with accounting for mental states. We then discuss theoretical perspectives on a range of seemingly disconnected phenomena—in particular lexical causatives and the so-called “coercion” phenomenon or, in our preferred terminology, indeterminacy. And we advance, even if briefly, a proposal for the representation and processing of conceptual content that does away with the analytic/synthetic distinction. We will argue that the only account of mental content that does away with the analytic/synthetic distinction is atomism. The version of atomism we will sketch accounts for the purported effects of analyticity with a system of inferences that are in essence synthetic and, thus, not content constitutive.

1 Semantics and the Architecture of Cognition

It is not uncommon for cognitive scientists working in semantics to mix their metaphors regarding how they envision the nature of mental representations and processes. Perhaps they do so inadvertently, but the price is a lack of clarity on what one takes to be the very nature of the representation of *content* and the computational processes that are content-bearing. And if there is one issue that research in semantics needs to be clear about, it is how it conceives content representation and processing. As an example, consider sentence (1).

- (1) Mary began a book.

Imagine now that the issue at hand is how a sentence such as (1) might be interpreted. The proposal quoted in (2) is apropos the sorts of psychological events carried out during the comprehension process of (1). The semantic issues underlying this proposal will be dealt with a little later, but we start off with the commitments of this proposal vis-à-vis cognitive architecture.

- (2) “(a) When encountering the noun *book*, comprehenders access the word’s lexical entry and attempt to integrate various stored senses of this word into the evolving semantic representation of the sentence.
- (b) The mismatch between the verb’s selectional restrictions and the stored senses of the noun triggers a coercion process.
- (c) Comprehenders use salient properties associated with the complement noun and other relevant discourse elements (including but not necessarily limited to the agent phrase) to infer a plausible action that could be performed on the noun.
- (d) Comprehenders incorporate the event sense into their semantic interpretation of the VP by reconfiguring the semantic representation of the complement, converting [_β began[_α the book]] into [_β began[_α reading the book]]. (Conceivably, this could also require reconfiguration of an associated syntactic representation.)” (Traxler et al., 2005, p. 4)

We use this as a convenient example of the kinds of constraints—or lack thereof—that may drive semantic proposals within the language processing literature. As we will see, similar proposals abound in semantic theory.

To begin, our commitments unequivocally reside with the view that representations are symbolic, with processes over these representations being computational. These general commitments come with numerous caveats. First, it is not clear whether the nature of computations performed over symbolic representations involve hard-wired algorithmic, intra-modular kinds of principles, or heuristic, perhaps malleable principles. This difference is important for semantics because, by hypothesis, it marks the boundary between linguistically-driven computations bearing on “shallow” meaning (viz., a logical form), and those deemed pragmatic or based on world-knowledge, contingent on experience. We mentioned “intra-modular” computations because our proposal relies on there being a modular level of linguistic computations whose output is a form of compositional semantic representation, a shallow one nonetheless (see Fodor, 2001; and de Almeida, 2018; and de Almeida & Lepore, 2018, for recent discussion).

Postulating that linguistic processes are computations over symbolic representations is crucial to our take on what sorts of knowledge representation enter into tasks such as understanding a sentence or having a thought. This is so because we assume that some of these processes are executed in virtue of the formal properties of the expressions that are computed, including properties of its constituent symbols, while others are entirely dependent on the content of token symbols—or the content that token symbols point to. Furthermore, we assume that semantic units—or

concepts—are the very elements of higher-level representations and processes, not only of linguistic representations proper. That is, thoughts have concepts as their most elementary parts, and those happen to be the same elements one recovers in the process of understanding a sentence; they are the same we ought to use in semantic analysis. As such, we assume that in order to account for the nature of these cognitive processes—that is, in order to account for the nature of those thoughts—it is crucial we not only understand the nature of the elementary parts, but also how they combine to yield the meaning that the thought carries.

Moreover, we think that to entertain a thought is to entertain something like a *proposition* whose basic elements are concepts. We take a proposition to be a mental object, a symbolic expression standing for the meaning of a sentence or other higher cognitive representation. Thus, we argue that *any* complex representation carrying content is propositional, barring cases in which ideas are incomplete (viz., arguments are not saturated) or when representations refer to individuals.¹ Thinking, thus, entails combining all the elementary concepts into series of propositions, which are most likely represented as something akin to a logical form specifying the relations between conceptual constituents (see Kintsch, 1974; and McKoon & Ratcliff, 1992, for early propositional theories). This view also applies to the process of language comprehension: understanding a sentence requires recovering the meanings of words/morphemes in the context of the *proposition* that the sentence expresses. Propositions are thus the mental objects whose referents are states and events in the world (and ideas about events and states in the imaginary world, if you will). In order for propositions to refer, or in order for propositions to stand for the events and states whose contents they represent, they have to compose, and in order for them to compose they require a syntax.

Much of what we talk about in the present chapter, thus, has a particular notion of compositionality lurking in the background: namely, one that takes lexical and functional constituents and how they are combined syntactically to determine sentence-level meaning. Clearly, any position one takes on the analytic/synthetic distinction (or lack thereof) has direct consequences for the kinds of elements that enter into the composition of meaning. For instance, let us assume that one holds an enriched form of compositionality, as proposed by Pustejovsky (1995) and Jackendoff (2002)—a proposal to which (2) above adheres. Leaving details aside, enriched compositionality takes the meaning of a sentence to rely on the interpolation of some features or ontologically primitive properties stored within lexical entries. Such a view is burdened with establishing an analytic/synthetic distinction. In principle, by appealing to the internal analyses of lexical items, compositionality cannot hold, for analyticity is necessarily unbounded, thus holistic. Furthermore, assuming that our thoughts are productive, and that productivity requires compositionality, then thoughts ought to be compositional. Thus any theory on the basic elements of meaning necessarily needs to account for the compositionality of thoughts (see Fodor, 1998, for a similar

¹We could argue that general or singular terms carry a property, viz., that ‘ $\exists x$ (MARY = x)’ is about being Mary. But we will eschew this issue and assume that complex representations include at a minimum singular terms and their predicates.

point). We think, in summary, that holding on to a strict notion of compositionality is imperative for determining which concepts theory prevails. However, as we will see in Sect. 3, there are different approaches to compositionality and this issue interacts with the position one takes with regards to the analytic/synthetic distinction.

So far, this general view of the nature of complex representations strikes us as standard, though by no means consensus. But before we move on to discuss analyticity in semantics, we have two other brief methodological observations to make regarding semantics research in cognitive science. The first methodological observation is this: since we are realists and naturalists about mental representations—semantic or otherwise—we contend that to *do* semantics one needs to appeal to all tools of cognitive science, bar none. We take it that linguistic methods may take precedence over others, for crosslinguistic generalizations and distributional properties of expressions often provide us with rich data, supporting arguments for the reality of particular types of semantic algorithms. But by the same token, we take the experimental tools employed in cognitive psychology and neuroscience to be crucial to advance theory, rather than simply supporting linguistic postulates. As Fodor, Fodor, and Garrett (1975) once suggested, native speakers' intuitions are psychological data; and if we are tasked to investigate the realm of psychological data, experimental evidence might be at par with crosslinguistic and distributional evidence. This is important to mention here because what we are about to discuss requires analyzing certain phenomena not only in light of theoretical arguments, but also relying on the results of empirical observations typically obtained in experiments.

The second methodological observation we want to make regards how semantics research often proceeds. We take it that the fault line of the analytic/synthetic distinction, which we will address in the next section, has caused some other cracks in the foundations of semantics. Virtually all attempts to develop a theory of features has taken place by appealing to what one knows to be true about referents—objects and events—in the world, which are not necessarily the kinds of information one represents in mind about these objects and events. Appeals to intuitions here can only go so far. We surmise, however, that much of what drives the proposal for feature sets as constituents of concepts relies on what has been called the “intentional fallacy”. In a nutshell, the intentional fallacy arises when the particular properties that one assumes to be part of a stimulus are attributed to its mental representation. In psychology, this is sometimes referred to as the “stimulus error”, after Titchener (1909). The intentional fallacy permeates work in semantics, for any semantics that appeal to features has the burden of establishing the criteria for what is to be taken as true properties of a stimulus (whatever those may be) from properties that may result from one's knowledge or beliefs about that particular stimulus. To put it simply, what the researcher knows to be true about a referent is not necessarily true of its mental representation. The consequences of this fallacy are pervasive, crucially affecting the discussion on what is analytic and synthetic, and by extension, where the line should be drawn between semantics and pragmatics (for further discussion, see de Almeida, 2018). As we will see, a key issue—in line with what we see in proposal (2)—is the idea of “coercion”. We turn to these matters now.

2 The Analytic/Synthetic Distinction and Semantic Theories

We start off by briefly revisiting the problem of analyticity and why it poses a challenge for semantic theories—at least semantic theories that share our architectural commitments—in particular the key issue of compositionality. We do so aware that these issues are far from new. But at the same time, we are concerned that they are rarely, if ever, addressed in the semantics literature.²

The analytic/synthetic distinction has been like a dark cloud over semantics ever since Quine wrote his *Two dogmas* paper. Quine was interested in debunking a kind of semantics—in particular Carnap’s—that appealed to what Carnap called logically true (or *L-true*) as opposed to “indeterminate” or factual (*F-true*) statements. The distinction goes back at least to Kant’s attempt at separation between analytic (*L-true*) and synthetic (*F-true*) (see Carnap, 1956, Chap. 1). But as Quine showed, there were no firm criteria for establishing this difference: in essence, *L-true* and *F-true* were sourced from the same data, even if on the surface some statements appear to be true in virtue of the meaning of their constituents (the likes of *A dog is an animal*). It should be clear, before we advance discussion, that our concern is not with truly analytic statements such as those in which a conjunction entails its parts. These are run over form—something like $P \& Q \rightarrow P$. The first case is obviously compatible with the architecture we adopt: in fact it is essential to algorithmic cognitive processes that they run over form, not content, such that it is always the case that $P \& Q \rightarrow P$ or $P \& Q \rightarrow Q$, no matter what *P* and *Q* stand for. Thus, analyticity of form holds. Our concern is with other, often subtler, forms of analyticity, common to lexical-semantic theories as well as theories of composition relying on certain types of semantic operations such as “coercion”. And, more broadly, our main concern is with the shaky ground upon which all of semantics that appeal to analytic *features* stands.

There are, we think, roughly three ways to conceive how a concept might enter into—i.e., contributes content to—a proposition. (i) The first is by contributing its full content, whatever that may be. If one believes concepts to be composed of particular sets of features, then the content that a given concept contributes to a proposition must necessarily be that particular set of features—nothing more, nothing less. (ii) Another way in which a concept might contribute content to a proposition is by contributing some, but not necessarily all, of its features. If one believes a concept to be made up by a set of features, then, the kinds of features that a concept might contribute to a particular proposition is relative to the particular context of the proposition—that is, it is sensitive to other constituent concepts, perhaps to the wider discourse, and perhaps to the syntax of the expression. And (iii) the third way in which a concept

² An anonymous reviewer was right at pointing out, among other problems, that the analytic/synthetic issue that we are trying to “reawaken” is “not new”. This, of course, is not an argument against our view. If anything, this is an embarrassment for semantic theories. We believe that the two case studies we discuss below, though limited in scope, are representative of a widespread practice in semantics. It should be noted that the kind of a/s issue we are raising is about mental representation, not linguistic analysis.

can contribute content to a proposition is somewhat similar to (i), but does away with analyticity: concepts contribute all their content, except that, according to this view, a concept has no features. In the present section, we will discuss (i) and (ii); the case for (iii) will be further advanced in Sect. 3.

We cannot possibly be exegetic in our evaluation of semantic theories that are committed to analyticity (see, e.g., Engelberg, 2011a, for review). Our goals here are to illustrate the state of the art and thus motivate our proposal for moving away from analyticity—namely, to make the case for our brand of atomism. And we will substantiate our case by discussing work from two particular semantic phenomena, one involving the representation of causative verbs, and one involving the representation of what we call “indeterminate” sentences, which in some circles is known as “coercion”. These two cases are illustrative for two reasons. The first, and perhaps most important one, is because both cases expose the root of the problem we want to shed light on: the problem of analyticity in semantics. The nature of the representation of causative verbs has long been the focus of disputes in linguistics and lexical semantic theories at least since the time of generative semantics (e.g., McCawley, 1972). The case of indeterminate sentences such as (1) has also received some attention early on (see Culicover, 1970). As we will see, these two topics are representative of how intuitions about meaning can lead to the intentional fallacy trap. And both represent challenges to the classical way of conceiving compositionality. But as we will see, in Sect. 3, we offer a parsimonious treatment of these two cases with the type of atomism *cum* inferences we propose and the classical notion of compositionality it entails. The second reason we focus on these two cases is, not coincidentally, that they have been topics of our own research—so we conveniently stay close to familiar cases to make a point we deem fundamental for investigating semantics in cognitive science, more broadly.

2.1 Causatives

Most theories of lexical semantic representation are committed to a form of analyticity that takes lexical meaning to be represented in terms of a cluster of features, usually expressed in the form of templates filled with variables and predicates. Causative verbs are the paradigm example as they have been the topic of many disputes between camps. A typical case is (3a), whose meaning is represented in (3b).

- (3) a. John_x broke the vase_y
 b. [[*x* ACT] CAUSE [BECOME [*y* *BROKEN*]]]

A representation such as in (3b), in the notation of lexical semantics (Levin & Rappaport Hovav, 2005) is nonetheless representative of other approaches such as conceptual semantics (Jackendoff, 1990, 2002), cognitive semantics (Croft, 2012), frame semantics (e.g., Fillmore & Baker, 2009), to cite a few. These theories differ in terms

of the types of information that enter into meaning representation, how features are combined, the nature of the primitive bases (viz., ontological categories upon which concepts are built), as well as the level, whether it be linguistic or conceptual, at which these representations are entertained.³ But their commonalities, by far, surpass their differences, for they all seem to appeal to hidden predicates and other analytic properties to account for the semantic representation of lexical constituents and their carrier sentences.

We assume that semantic templates such as (3b) are intended to represent the *propositional* content of (3a) specifying its form and key elements of meaning.⁴ The evidence corroborating this view either comes from distributional data or from experiments suggesting that complex templates are more difficult to process than simplex ones (i.e., they engender longer reading times; McKoon & Macfarland, 2000) or involve more “connections” (Gentner, 1981) between other simpler concepts in memory and are thus better recalled. We won’t repeat the review of the arguments and experimental studies supporting predicate decomposition, here (see de Almeida & Manouilidou, 2015; also Engelberg, 2011b): there seems to be widespread agreement of decompositional views, which spares us from a more thorough review. Our mission is rather to call attention to the evidence *against* decomposition, which also comes from distributional evidence and experiments—but which enjoy much less acceptance.

The first kind of evidence pertains to the lack of synonymy between sentences that are supposed to be semantically represented by the same constituents.⁵ Take (4a) and (4b) as examples. These sentences, by hypothesis, yield the same semantic representation, as in (4c): while (4a) involves the lexical causative, (4b) involves its periphrastic counterpart. Unless the periphrastic *cause x to die* does not mean what is in (4c), the idea is that the two sentences are synonymous—hence that the template in (4c) should hold for *both* (4a) and (4b).

- (4) a. John killed the cat
 b. John caused the cat to die
 c. [[x ACT] CAUSE [BECOME [y <DEAD>]]]

But as Fodor (1970) argued sentences such as (4a) and (4b) do not denote the same events, for one can cause the cat to die on Saturday by poisoning his food on Thursday,

³We are assuming throughout that these theories all postulate that template structures are representations of psychological objects, as in Jackendoff (1983), similar to representations in a language of thought, though this is not always explicit in the works we cite.

⁴Although most of our discussion focuses on a theory such as Levin and Rappaport Hovav’s (2005), we assume that the main points we make apply to all theories we mentioned.

⁵An anonymous reviewer pointed out that, “Most people don’t assume that in order for there to be synonymy (and thus, analytic truths), the expressions in question need to be psychologically perfectly equivalent. For instance, it is standardly accepted that a correct analysis can be highly non-obvious.” We fail to understand what “most people” assume, for we do take synonymous sentences in natural language to be expressions of “perfectly equivalent” mental states (viz., propositions).

but one cannot kill the cat on Saturday by poisoning his food on Thursday. The distribution of time adverbials suggests that these are not similar events.⁶

Along similar lines, there are diverse experiments suggesting that causatives do not decompose, for they do not exhibit complexity effects (e.g., de Almeida, 1999a; Fodor et al., 1975, 1980; Kintsch, 1974; Manouilidou & de Almeida, 2013; Rayner & Duffy, 1986; Thorndyke, 1975; see de Almeida & Manouilidou, 2015, for review). These studies have employed numerous techniques—from judgment to reading times—and have been consistent in pointing to the lack of decomposition effects. More recently, data from Alzheimer’s patients have also landed support to this camp. For instance, if verbs are represented by semantic templates, we should expect the pattern of deficits to reflect the purported effect of semantic complexity—with more complex concepts being harder to retrieve. Notice also in passing that the more predicates a template carries, the greater the chances that the concept might be impaired. But as we have recently shown (de Almeida, Mobayyen, Antal, Kehayia, Nair, & Schwartz, 2021), when Alzheimer’s patients are asked to name video clips of events and states which depict classes of verbs with varying complexity (e.g., causatives, motion, and perception/psychological), these patients’ naming pattern does not line up according to the predicted complexity. Causatives, which contain hypothetically more predicates are not affected as severely as psychological verbs, which contain less predicates. The pattern of results suggests that categorical deficits are not along the lines of semantic template complexity, but rather along the lines of thematic structure, with verbs assigning an *Experiencer* role to the subject position being harder to name. We assume that thematic roles are “psychologically real”: they affect the composition of a sentence in the mapping between syntax and the logical form, viz., by assigning roles to constituents based primarily on their syntactic positions and following the structural specifications of the predicate (see also Manouilidou, de Almeida, Nair, & Schwartz, 2009, for compatible results).

Crucially, the properties that enter into templates are far from well justified, for neither their ontological status has been determined, nor has the selection of features been principled.⁷ At first, it may seem like a daunting task to think of a concept without thinking about the constituent parts we know (or more like *think*) to be true of that particular stimulus. For instance, it may be difficult to think of DRINK without entertaining thoughts such as LIQUID, or MOUTH. But entertaining these thoughts, as a function of entertaining DRINK does not necessarily entail that the likes of LIQUID and MOUTH are to be taken as *constituent features* of DRINK. Furthermore, if these features are taken to be constituents of DRINK, then, we can conclude that they too carry content themselves which are expressed in terms of

⁶This is perhaps old news but to our knowledge, with few exceptions (e.g., Jackendoff, 1990, 2002; Harley, 2012), it has not been addressed in the literature.

⁷As Jackendoff (2002, p. 377) puts it, lexical-semantic decomposition “... is a richly textured system whose subtleties we are only beginning to appreciate (...). It does remain to be seen whether all this richness eventually boils down to a system built from primitives, or if not, what alternative there may be.” While we take this position seriously, our point here is that the *a/s* distinction stands as the main obstacle to the empirical prospects of lexical semantics.

other features. The consequence of this is holism about content. And holism is the antithesis of semantics—as Quine had first suggested.

As a further example of this state of affairs, consider the distinction between so-called “externally caused” and “internally caused” change of state verbs such as those in (5a) and (5b) respectively.

- (5) a. The cement crumbled
b. The apple rotted

Although much of this distinction bears on the realization of predicate-arguments (e.g., externally caused verbs usually do not enter into transitive forms), a critical issue is how the distinction is made in *semantic analysis*. For Levin and Hovav (1995), internally caused change of state verbs denote events brought about naturally in the object, while externally caused change of state verbs “imply the existence of an ‘external cause’ with immediate control over bringing about the eventuality described by the verb: an agent, an instrument, a natural force, or a circumstance” (p. 92).

The way the difference between these verb classes is presented appeals to our (perhaps naïve) knowledge of physics. But even that might fail us for we are not certain whether what makes something *rot* is internal or external, that is, whether atmospheric variables are the triggers of rotting, or alternatively if an object—say, an apple—rots entirely on its own. The same can be said of cement crumbling. The physics baggage is heavy. And we suspect this case lines up with classical cases of intentional fallacy plaguing semantics: even if the rot/crumble distinction can be determined solely on linguistic (viz., structural) principles, it is an entirely different claim to attribute the difference to mentally represented properties of the two types of events. Understanding the properties of the world will not help us fix the properties of semantic representations.

The point we are making, in summary, is one we have briefly touched upon in the previous section: just because one knows a stimulus or phenomenon to be composed of certain properties, it does not entail that these properties are encoded as *mental representations* of the stimulus or phenomenon. This is precisely the perennial effect of the intentional fallacy on semantic theorizing.

Before we further explore this issue, in contrast to atomism in Sect. 3, we would like to address rather briefly a second semantic phenomenon—*coercion*—one for which appeals to analyticity are also quite evident.

2.2 Indeterminacy (or “Coercion”)

The term “coercion” (or type-coercion, or type-shifting) is identified with particular hypotheses on how sentences such as (1) are interpreted—among which is the proposal presented in (2). We refer to these sentences as “indeterminate” because the actual action that Mary performed with the book is not determined, although the sentence is grammatical and a truth value judgment can be made (namely, it is true

if Mary began to do anything with the book); so much for terminology. The “coercion” hypothesis assumes that the proposition expressed by sentences such as (1) are necessarily enriched along the lines of what is exemplified in (2), but in particular proposal (2d), which we repeat here for convenience.

- (2) (d) Comprehenders incorporate the event sense into their semantic interpretation of the VP by reconfiguring the semantic representation of the complement, converting [_β began[_α the book]] into [_β began[_α reading the book]]. (Traxler et al., 2005, p. 5)

This processing hypothesis largely follows the theory of type coercion proposed by Pustejovsky (1995). The essence of coercion is that the alleged mismatch between the verb’s selectional restrictions and the nature of the internal argument. By assumption, the verb *begin* selects for an *event*, though the noun *book* is an *entity*. This mismatch triggers the search for a “plausible action” that would yield an enriched semantic composition, by interpolating a semantic constituent such as *reading* into the final form. But as we briefly alluded to in Sect. 1, a commitment to such a process entails a commitment to determining which, among all possible *senses*, are the ones to be interpolated into the resulting representation.

There is perhaps some confusion here between *meaning*, *sense*, and *use*—damage that unfortunately Wittgenstein cannot come back to repair. If we tell you that it is *hot* today, in Montreal, when actually it is -20°C , we are most likely being sarcastic. It does not entail, now, that the concept HOT includes COLD, among its *senses*. We are certainly *using* the word *hot* to convey something else entirely, to provoke you or, as Davidson (1978) would say, to invite you to think, just like we would do with a metaphor. And even if we were to admit that *senses* are represented in close proximity (by some metric) with the original concept, as a function of extensive *use*, there is no saying on how a *sense* is to be accessed, other than via its actual host concept. Thus, to make a simple point: it is HOT that needs to be accessed such that COLD can be entertained.

It is clear that hypotheses committed to multiple layers of properties supposedly stored with token items are simply question begging: which sorts of elements are the ones to be chosen, and how are they to be chosen? As we will argue in Sect. 3, a different explanation can be offered in cases of conceptual tokening: inferences driven by synthetic relations are the ones that yield the *effects* which decomposition-*alists* claim to be effects of constituency. We will, thus, offer a more parsimonious analysis of this phenomenon, doing away with analyticity and placing the burden of interpretation on the identification of gaps, at the syntactic and logical-form representation of sentences, with most interpretation post-logical form being inferential, not relying on analytic properties of lexical concepts.

3 Alternative: Atomism and Inferences

What is, then, our proposal for doing away with analyticity? We should warn you that the proposal might be disappointingly simple, and our presentation of the theory will be somewhat constrained by the scope of the present chapter. Here is how we proceed. We start off by connecting our view of concepts with what we envision to be the architecture of cognition, as briefly presented in Sect. 1. Then, we discuss two main issues: (i) the representation of concepts according to our brand of atomism; and (ii) how concepts might be causally connected to each other—viz., as inferential relations. And, throughout, we tailor our discussion of atomism and inferences to the analysis of the two phenomena we discussed in Sect. 2.

We have mentioned that we are committed to symbolic representations and to computational processes. Patently, we take symbols that stand for content to be atomic, not molecular representations. And we take these symbols to compose into complex structures the classical way: complex symbolic expressions get their meaning as a function of the meaning of their constituent symbols and how they are arranged in propositions. Symbols then carry (or point to) information about the things (and events) they refer to. We do not establish a lower limit on the content that the simplex symbols convey—or more properly on the very content that they individuate—but we suggest that they are properties, predicates, and “particulars”, as Russell (1913) once put it. We assume that, for the most part, atoms are expressed by the simplex bound and free morphemes of natural language. And since we take concepts to be the very symbols of (again, Russell) our “experience”, we assume that they enter into different cognitive processes via computations.

So much for linking our view of conceptual representation and processes to the architecture we presented in Sect. 1. As for the nature of conceptual representation, if concepts are “atoms”, they are simply individuated by the kinds of things they refer. One quick note should suffice to address the problem of reference here: while we take concepts to be pointers to objects (in a very broad sense, including properties like patches of color) and events, they are also representations of things for which there is no referent (or, again, as Russell put it, in the “past, present, or not in time at all”, p. 5).

Two further observations are in order. The first is that it is likely that the things concepts individuate are full objects—the midsize things that populate scenes like chairs and pencils—or full events. But they can be just fractions of these: there is nothing in the system we suggest that ties the tokening of concepts to these ontological categories. And, to our knowledge, there is no clear line demarcating parts and objects, or objects and scenes (to wit, HORIZON is an “object” for all practical purposes; and so are DOG and TAIL). Second, a related issue: it is quite plausible to take “particulars” to be the tokening elements upon which one arrives at a given concept. For instance, it is well known that events have no fixed boundaries, that is, that the meaning of the verb *to kill*, say, does not pick up particular time and space properties, with well determined beginning and end points. Not even the property of being dead marks the endpoint of *kill*, for *to die* also lacks clearly perceptually

marked boundaries. Moreover, it is not the case that having *kill* entails having *dead*. In our system, the relation is inferential, not one of dependency.⁸ If so, most likely the kinds of “particulars” that the conceptual system locks into may be the very entry points to the sets of inferences one runs in conceptual processing. This may become clearer with an example.

Take (6) to be the referential relation that obtains between the word (or the object) *dog* and its concept.

(6) *dog* → DOG

The locking mechanism that affords DOG out of the word or object is a mechanism that in principle is tokened by whole objects, assuming that the visual attentional mechanism locks into full objects (see Fodor & Pylyshyn, 2015; Jackendoff, 2002). But it may well be the case that what one gets are *parts* of objects. Thus, getting TAIL tokened is what gets one to eventually entertain DOG. Notice that in order for this system to work, there ought to be a system of relations between concepts. As we mentioned above, we are committed to having conceptual relations that are *not* necessary; that is, to use the example, it is not the case that tokening TAIL necessarily causes DOG; only *tail* causes TAIL, but we suggest that one might get to the host object via its parts, not because they are conceptually dependent, but because they are inferentially connected.

We owe you, of course, a bit more clarity on how the system might work regarding these *non-analytic* inferences. We propose to work with the two phenomena we discussed in Sect. 2, beginning with causatives and, soon after, with the comprehension of indeterminate sentences. Along the way, we make a few observations regarding the less developed parts of our proposal.

3.1 Back to Causatives

Although we take Carnap’s commitment to analyticity in semantics to be misguided—just like Quine put it—the *tools* we inherited from him are of particular importance for conceiving psychological inferences bearing on meaning. Enter meaning postulates (henceforth MPs), which are quasi-logical inferences. We say quasi-logical only in the sense that they are not proper inferences whose consequent is *by necessity* entailed by the antecedent. And while this is a common tool in semantics, we take the kinds of MPs that run between concepts to be the very inferences

⁸We note in passing that, although this would take us far afield, what counts for us as a perceptual boundary for, say, *to die*, is tied to observation, not to the actual act of dying which is independent of observation. To wit, consider the end point of the verb *to break* as in *John broke the vase*: would it be when all physical particles of said vase cease moving? The concept BREAK is not determined by the actual physical phenomenon, by Newtonian laws (those are not “in the head”; cf. the intentional fallacy) but by when *break* causes BREAK.

that give rise to a myriad of relatedness *effects* found in the empirical literature and in other frameworks committed to analyticity.

Consider causatives. As we discussed above, voices in unison claim that causatives decompose. But there is strong evidence—from experiments and arguments—that causatives might not decompose. How, then, can one account for the pervasive effects obtained in the relations between arguments of the verb? How can one account for the pervasive effect of relations between transitive and intransitive variants of the same root verb? One way to conceive the relation between concepts—such that KILL and DIE or BOIL-transitive and BOIL-intransitive are related—could be by running inferences such as in (7).

- (7) a. $\forall x \forall y [\text{BOIL}(x, y)] \rightarrow [\text{BOIL}(y)]$
 b. $\forall y [\text{BOIL}(y)] \rightarrow [100^\circ\text{C}(y)]$
 c. $(\forall x [P(x)] \rightarrow [Q(x)])_n$

We can cast this proposal in simple predicate logic, by attributing properties to individuals and by linking predicate relations as inferences. We can only highlight a few of the characteristics of this system—the ones that are in direct contrast with compositional views discussed in Sect. 2. Notice also that the relation between transitive and intransitive variants of the same core concept can be accounted for by the entailment between arguments of the verb. But our suggestion is that beyond those entailments—which are in essence argument-structure driven—“properties” of the event denoted by the verb are also attained by these relations. We won’t extend this account of causatives here much further (but see de Almeida, 1999a, b, for early versions of this proposal). Suffice it to say that these inferences are not content-constitutive, thus, that it is not the case that the content of an utterance or a thought somehow depends on the “appropriate” inferences being computed. To us, the inferences that are typically run when concepts are tokened are *synthetic*, thus their actual content cannot be accounted for by semantic analysis.

We also acknowledge that even those with whom we share the main tenets of atomism have argued against adopting MPs for they are too unconstrained and thus cannot be used as an account of semantic inferences (Fodor, 1998). We part ways here. While we agree that they are unconstrained, our goal is not to model the very content tokened by a concept such as KILL or BOIL, but the inferences that might ensue that are taken to account for the conceptual content in all sorts of psychological effects (from priming to prototypicality to semantic-memory impairments). In summary, we suggest that inferences such as (7b) are entirely contingent on experience. And we suggest (7c) to be a basic law of how inferences run over predicates. To assume that those inferences constitute the representation of lexical content is, in principle to incur in the intentional fallacy.

3.2 Back to “Coercion”

We turn now to the other phenomenon, that of the comprehension of indeterminate sentences such as (1). To ease discussion and comparison with (2)—we will cast our proposal rather informally as in (8).

- (8) (a) Every incoming token lexical item (i) maps onto its corresponding concept (*book* → BOOK), (ii) contributes its syntactic information to the evolving syntactic tree (*book* _N), and (iii) contributes logical information to an evolving semantic composition (viz., a logical form; $\exists x, \text{BOOK}(x)$)).
- (b) The evolving syntactic parsing for a sentence such as (1) tags all its lexical constituents and its linguistically motivated gaps—viz., the gaps for syntactic positions that may be optionally filled-in lexically. As for (1), the gap is potentially in the VP, as in $[\text{VP } [\text{V}^0 \text{ began } [\text{V}^0 e [\text{OBJ NP}]]]]$.
- (c) The concepts that are accessed (mapped onto) by each lexical item are premises for *synthetic* inferences whose consequents are experience-based relations yielding between predicates (thus, a possible inference would be $\forall x \text{ BOOK}(x) \rightarrow [\text{READ}[\text{ABLE}]](x)$)).
- (d) The meaning of a sentence is obtained by combining the token concepts—the translations of morphemes—into the evolving logical form, such as $\exists x(=\text{MAN}), \exists y(=\text{BOOK}) (\text{BEGIN } (x, y))$ (or, alternatively, $\exists w (\text{BEGIN } (x, y, w))$); that is the shallow, *unenriched* interpretation of (1).
- (e) Many processes of enrichment ensue; among them are the processes of filling the gaps identified during syntactic structuring with the concepts that were part of the postulates triggered by (i) the utterance context, and (ii) the co-text.

We can only make brief observations about (8)—but we trust that the contrast with (2) is quite clear. First, notice that the *meaning* of *book* is not a *sense*; and, according to our proposal, there are no senses stored *with* the meanings of words. We do not deny that there are *uses*, but *uses* are obtained pragmatically (they are *synthetic*; see below), within the inferences that run after conceptual tokening (as in 8a) and conceptual composition. Also, as suggested in (8b) there are linguistic arguments for holding a syntactic gap within the VP of sentences such as (1) without appealing to effects of “coercion”.⁹ And we hold that the coercion effects shown in most experimental studies could be effects of this gap as they can also be effects of inferences that the indeterminate sentence triggers.

The advantage of a proposal such as the one sketched in (8), in summary, is that it does away with analyticity. For any of the proposals appealing to analytic properties, the burden is to determine the criterion for separating analytic from synthetic

⁹Several linguistic arguments for the VP gap hypothesis appear in de Almeida and Dwivedi (2008) and in de Almeida and Riven (2012). Also, see arguments against coercion alternatives in de Almeida and Lepore (2018) and in de Almeida et al. (2016), which we cannot begin to discuss here.

properties. We do not appeal to such properties because to us concepts are atomic, but we see a role for such properties in the *inferences* that ensue upon conceptual tokening and semantic composition.

3.3 Conclusion: Atomic Concepts and Inferences

We conclude by stressing a few points about our proposal. First, in the sense we take in the present proposal, the inferences about lexical-conceptual properties are mostly (if not all) synthetic, not analytic, as mentioned above. Thus, one can know what a dog is without knowing what an animal is or what a pet is, for that matter. Crucial to this approach is the idea that all such relations, commonly known as constituent *features*, are synthetic and thus the inferences that run over them are not *necessary* for content attainment. In fact, only the content that each individual symbol instantiates suffices, independent of the inferences it generates. If inferences are synthetic, they cannot be part of the meaning of a token item. And if they are not part of meaning, we can dispense with a semantics that attempts to legislate on experience and world knowledge.

Second, we assume that many of the inferences that run as a consequence of a concept being triggered are common to many inhabitants of the same community, those sharing similar kinds of experiences. We cannot be precise on this idea because it points to something whose variables are virtually infinite. Crucial to our approach, in fact, is the idea that these commonalities cannot be legislated on. We also suggest that many, perhaps most effects found in the literature—from priming to prototypicality—are manifestations of these inferences; they are effects of the causal connectedness established between concepts as a function of use and experience. And we even acknowledge that it may be difficult to dissociate—empirically—between inferences computed upon tokening concepts and effects of “activation” of properties. However, we have presented some clear signs from the literature that point against decomposition.

We do hold that there is a crucial distinction, upon which a theoretical advantage stands: by not taking properties to be analytic, there is no commitment to building a semantic theory whose foundations are faulty. The crucial distinction between atomism and molecularism is that the former, but not the latter does not require semantic analysis based on features or synonymy and, because of that, there is no analysis of content other than assuming that concepts (and their lexical labels) are largely referential, symbols that point to things, events, ideas, and so forth. Reference does not entail being in the presence of the object or event: it entails bringing to fore the relation between the symbol and the thing/event/idea it designates.¹⁰

If semantics appeals to features, without an analytic/synthetic distinction, it turns to holism, which is the antithesis of semantics—at least of a semantics committed to

¹⁰This point was made by Russell (1913, Chap. 3) and, more recently, by Fodor and Pylyshyn (2015, Chap. 5) regarding reference “beyond the perceptual circle”.

compositionality and productivity. If semantics appeals to properties of the world to fix properties of mental representations, it may fall into the intentional fallacy trap. The way semantics can avoid all this trouble is to turn to atomism cum inferences.

Acknowledgements We are grateful to two anonymous reviewers for very useful comments. We are aware we did not address all their concerns, but we hope to have further clarified our position. We also thank the National Sciences and Engineering Research Council of Canada (NSERC) for financial support.

References

- Carnap, R. (1956). *Meaning and necessity* (2nd ed.). Chicago University Press.
- Croft, W. (2012). *Verbs: Aspect and causal structure*. Oxford, UK: Oxford University Press.
- Culicover, P. (1970). One more can of beer. *Linguistic Inquiry*, 1(3), 366–369.
- Davidson, D. (1978). What metaphors mean. *Critical Inquiry*, 5(1), 31–47.
- de Almeida, R. G. (1999a). *The representation of lexical concepts: A psycholinguistic inquiry*. Unpublished doctoral dissertation, Rutgers University, New Brunswick, NJ.
- de Almeida, R. G. (1999b). What do category-specific semantic deficits tell us about the representation of lexical concepts. *Brain and Language*, 68, 241–248.
- de Almeida, R. G. (2018). Composing meaning and thinking. In G. Preyer (Ed.), *Beyond semantics and pragmatics: New essays* (pp. 201–229). Oxford, UK: Oxford University Press.
- 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., & Lepore, E. (2018). Semantics for a module. In R. G. de Almeida & L. Gleitman (Eds.), *On modules, concepts, and language: Cognitive science at its core*. Oxford, UK: 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, NY: Springer.
- de Almeida, R. G., Riven, L., Manouilidou, C., Lungu, O., Dwivedi, V. D., Jarema, G., et al. (2016). The neuronal correlated of indeterminate sentence in interpretation: An fMRI study. *Frontiers in Human Neuroscience*, 10, 614. <https://doi.org/10.3389/fnhum.2016.00614>.
- de Almeida, R. G., Mobayyen, F., Kehayia, E., Antal, C., Nair, V., & Schwartz, G. (2021). Category-specific semantic deficits for verbs in Alzheimer's disease: Evidence from static and dynamic action-naming tasks.
- Engelberg, S. (2011a). Frameworks of lexical decomposition of verbs. In C. Maienborn, K. von Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. 1, pp. 358–399). Amsterdam: De Gruyter Mouton.
- Engelberg, S. (2011b). Lexical decomposition: Foundational issues. In C. Maienborn, K. von Heusinger, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. 1, pp. 124–144). Amsterdam: De Gruyter Mouton.
- Fillmore, C. J., & Baker, C. (2009). A Frames approach to semantic analysis. In B. Heine & H. Narrog (Eds.), *The oxford handbook of linguistic analysis*. Oxford, UK: Oxford University Press.
- Fodor, J. A. (1970). Three reasons for not deriving 'kill' from 'cause to die'. *Linguistic Inquiry*, 1, 429–438.
- Fodor, J. A. (1998). *Concepts: Where cognitive science went wrong*. Oxford, UK: Oxford University Press.
- Fodor, J. A., Garrett, M. F., Walker, E. C., & Parkes, C. H. (1980). Against definitions. *Cognition*, 8(3), 263–367.

- Fodor, J. A., & Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28(1), 3–71.
- Fodor, J. A., & Pylyshyn, Z. W. (2015). *Minds without meanings: An essay on the content of concepts*. Cambridge, MA: MIT Press.
- Fodor, J. D., Fodor, J. A., & Garrett, M. F. (1975). The psychological unreality of semantic representations. *Linguistic Inquiry*, 6, 515–531.
- Gentner, D. (1981). Verb semantic structures in memory for sentences: Evidence for componential representation. *Cognitive Psychology*, 13, 56–83.
- Jackendoff, R. (1983). *Semantics and cognition*. Cambridge, MA: Cambridge University Press.
- Jackendoff, R. (1990). *Semantic structures*. Cambridge, MA: MIT Press.
- Jackendoff, R. (2002). *Foundations of language*. Cambridge, MA: MIT Press.
- Kintsch, W. (1974). *The representation of meaning in memory*. Hillsdale: Lawrence Erlbaum Associates.
- Levin, B., & Hovav, M. R. (1995). *Unaccusativity: At the syntax-lexical semantics interface*. Cambridge, MA: MIT press.
- Levin, B., & Rappaport-Hovav, M. (2005). *Argument realization*. Cambridge, MA: Cambridge University Press.
- Manouilidou, C., & de Almeida, R. G. (2013). Processing correlates of verb typology: Investigating internal structure and argument realization. *Linguistics*, 51, 767–792.
- 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(2), 167–186.
- McKoon, G., & McFarland, T. (2000). Externally and internally caused change of state verbs. *Language*, 76, 833–858.
- McKoon, G., & Ratcliff, R. (1992). Inference during reading. *Psychological Review*, 99, 440–466.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
- Quine, W. V. O. (1953). Two dogmas of empiricism. In *From a logical point of view* (2nd ed.). Cambridge, MA: Harvard University Press.
- Rayner, K., & Duffy, S. (1986). Lexical complexity and fixation times in reading: effects of word frequency, verb complexity, and lexical ambiguity. *Memory and Cognition*, 14(3), 191–201.
- Russell, B. (1913/1984). *Theory of knowledge: The 1913 manuscript*. London: Routledge.
- Titchener, E. B. (1909). *Lectures on the experimental psychology of thought-processes*. New York, NY: Macmillan.
- 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.
- Thorndyke, P. W. (1975). Conceptual complexity and imagery in comprehension and memory. *Journal of Verbal Learning and Verbal Behavior*, 14, 359–369.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

