

**Chomskyan arguments against truth-conditional semantics based on variability  
and co-predication**

Agustín Vicente

Ikerbasque: Basque Foundation for Science & University of the Basque Country  
(UPV/EHU)

agustin.vicente@ehu.eus

*Abstract*

In this paper I try to show that semantics can explain word-to-world relations and that sentences can have meanings that determine truth-conditions. Critics like Chomsky typically maintain that only speakers denote, i.e., only speakers, by using words in one way or another, represent entities or events in the world. However, according to their view, individual acts of denotations are not explained just by virtue of speakers' semantic knowledge (since, according to them, semantic knowledge is very scarce: see Pietroski, 2018). Against this view, I will hold that, in the typical cases considered, semantic knowledge can account for the denotational uses of words of individual speakers.

*Introduction*

The idea that the meaning of a sentence of a natural language determines its truth-conditions has a long history. Minimally, the idea comes to the view that, given that a sentence is a representation that has a predicative structure, its representational content determines the conditions under which the representation is true. In general, the content

of a representation, or what the representation is about, determines its *accuracy* conditions. For instance, the content of a perceptual representation determines the conditions under which tokening that representation would make the tokening accurate. If the perceptual representation is about a black flying object, then the tokening of the representation is accurate only if it is tokened when there is a black flying object in the environment. All representations have accuracy conditions determined by what they are about. The accuracy conditions of representations that have a predicative structure are truth-conditions: the representation is accurate if it is true<sup>1</sup>. Thus, if a sentence of a natural language is representational, its content has to determine its truth conditions.

The tradition of thinking that the meaning of a sentence can be expressed in terms of the truth conditions that its content determines is under heavy attack these days. The more we know about language, the more clearly we see that it is impossible to pair sentences with truth conditions one to one, even if we remove from our vocabulary all those expressions that require some contextual parameter to get a denotation (from demonstratives to gradable adjectives). This is because the rest of the lexical items also fail to have a definite denotation or content. No single word-type, it seems, has a representational content in the minimal sense that was introduced above, such that it seems impossible to provide the accuracy or appropriateness conditions of the use of a certain word. Examples abound, but probably the most striking ones concern proper and natural kind terms. As Chomsky (2000, 2016) has long argued, a proper name such as *London* can be used to refer to a place, to a set of people, to a political institution, to an economic centre, and even to a way of life, while the natural kind term *water* refers to

---

<sup>1</sup> I do not mean to say that only sentences have predicative structure. The claim is that if a representation has a predicative structure, then its accuracy conditions are truth-conditions.

H<sub>2</sub>O only in the scientific “language game”, so to speak<sup>2</sup>. Otherwise, we do not call *water* a cup of tea, even if it may have a higher percentage of H<sub>2</sub>O molecules than the liquid that comes from a well (see Malt, 1994, Pietroski, 2018) The accuracy conditions of uses of *London* or *water*, thus, are not univocal or even consistent.

In the case of perception, it has proven to be notoriously difficult to specify what a certain perceptual representation represents. There is a long debate concerning what kind of entity a frog may be representing when she snaps her tongue at a fly. However, we can say that a certain perceptual representation produced by the frog’s visual system represents, or is about, some particular content –we only do not know what that content is, and thus, we are not in a position to state its accuracy conditions-. A word-type, in contrast, seems to lack a particular denotation. That is, the issue is not that we do not know what content a word has; the issue is that a word does not have a content, i.e. an entity or event that it is about. As Chomsky and Pietroski insist, word-types fail to have denotations. A denotation, so the reasoning goes, is some entity in the world (or in the world as we believe it to be), and there is nothing in the world that is a place, a political institution, a way of life, etc., all at the same time<sup>3</sup>. Also, there is nothing in the world that covers all the uses we make of the term *water*, uses which are not based on how much H<sub>2</sub>O a certain liquid has.

---

<sup>2</sup> Here I will be concerned with cases similar to the *London* and the *water* case. For other examples and responses, see Kennedy and Stanley (2009), Forbes (2012), and Segal (2012).

<sup>3</sup> See Chomsky (2000; 37): “London is not a fiction, but considering it as London – that is, through the perspective of a city name, a particular type of linguistic expression – we accord it curious properties: as noted earlier, we allow that under some circumstances, it could be completely destroyed and rebuilt somewhere else, years or even millennia later, still being London, that same city. [...] We can regard London with or without regard to its population: from one point of view, it is the same city if its people desert it; from another, we can say that London came to have a harsher feel to it through the Thatcher years, a comment on how people act and live. Referring to London, we can be talking about a location or area, people who sometimes live there, the air above it (but not too high), buildings, institutions, etc., in various combinations (as in *London is so unhappy, ugly, and polluted that it should be destroyed and rebuilt 100 miles away*, still being the same city). Such terms as *London* are used to talk about the actual world, but there neither are nor are believed to be things-in-the world with the properties of the intricate modes of reference that a city name encapsulates”.

Travis' famous cases point at the same phenomenon (Travis, 1996, 2008): in *There is milk in the fridge*, *milk* may stand for some portion of drinkable milk, or for some portion of milk that has been spilled inside the fridge. In *The leaves are green*, the leaves we speak about can be the leaves as they are intrinsically, or the leaves as they look, after having been painted (see Vicente, 2015), etc. So, declarative sentences do not have a particular truth conditional content, because their constituents do not have a denotational semantics, i.e. their constituents are not representations in the sense explained above: they are not about particular entities or events in the world. From here critics conclude that semantics is not in the business of explaining word-to-world relations (Chomsky, 2000, Pietroski, 2005, 2018, Yalcin, 2014, Collins, 2017).

This latter is the thesis I will scrutinize in this paper. I will try to argue that the arguments that critics present do not show that semantics cannot explain word-to-world relations because sentences do not have meanings that determine truth-conditions<sup>4,5</sup>.

Critics typically maintain that only speakers denote, i.e., only speakers, by using words in one way or another, represent entities or events in the world. However, according to their view, individual acts of denotations are not explained just by virtue of speakers' semantic knowledge (since, according to them, semantic knowledge is very scarce: see Pietroski, 2018). Against this view, I will hold that, at least in the cases critics seem to find more problematic, semantic knowledge can account for the denotational uses of

---

<sup>4</sup> That sentences have meanings that determine truth-conditions does not imply that sentences have *univocal* truth-conditions. It can be that sentences have several meanings, and that each of these meanings determines a set of truth-conditions. That is, in the view, as it is stated, sentential meanings cannot be captured in a simple T-schema.

<sup>5</sup> Also, I have to note that I will be less concerned with sentences and truth-conditions than with (certain) words and denotations. That words have denotations is a pre-requisite that has to be satisfied so that sentences can have truth-conditions. But then it has to be explained how composition rules work: this is the part that I will be less concerned with in this paper. Also, the paper will be focused on nominals and what nominals denote because the arguments based on variability and co-predication affect nouns (proper or common). As I will argue, this is also the case with respect to Travis cases. The thesis that I defend may not be valid for other classes of words (though see below). But the dialectics of the paper is: if the biggest problem that truth-conditional semantics has to face is the one that the critics signal, truth-conditional semantics is not in such big trouble.

words of individual speakers. I will summarize my position at the end of this introductory section. Before that, I briefly review different ways in which one can react to this kind of attack on truth-conditional semantics.

There are different ways to react to this attack on truth-conditional semantics. One is to single out a particular denotation among the different denotations that a word can have and explain the rest of possible denotations in terms of semantic or pragmatic mechanisms. For instance, it can be held that the literal meaning (i.e., the “real” denotation) of *London* is a certain spatial location, the rest of the possible denotations being generated by means of metonymical operations on that literal meaning. If the explanation provided for these denotation shifts is pragmatic, the result would be that although word-types have denotations and sentences contents that determine truth-conditions, the truth-conditions of particular utterances of those sentences will typically not coincide with the truth-conditions of the sentence-type. This gives rise to the debate between minimalists and contextualists over what type of content theoreticians should attempt to explain: literal truth-conditions, which are usually not entertained in linguistic processing, or “intuitive” truth-conditions, which incorporate already shifted denotations (Recanati, 2010). This kind of “go pragmatic” reaction to arguments from variability is now losing traction. On the one hand, as some authors have noted (e.g., Asher, 2011, Del Pinal, 2018), it is certainly difficult to specify the when and why of the pragmatic mechanisms that are said to operate on literal meanings. An account such as Recanati’s, which allows for free modulation of literal meanings according to contextual information, seems to lack the resources to explain why *John finished the apple* can easily be read as “John finished eating the apple” while *John stopped the apple* will never be read in such a way (Asher, 2011). Similarly, it is easy to understand *Mary began the book* as meaning that Mary began reading the book, but an utterance of

*Mary began the elevator* will not be understood as meaning that Mary began going up in the elevator. On the other hand, the idea that lexical items have literal meanings that are simple and stable denotations is an idealization that has at least two problems. The first problem is that the idealization is insensitive to a lot of systematic ways in which words combine (see below on Travis cases) and to the richness of such combinations: for instance, a privative adjective such as *fake* does not simply affect the denotation of the noun it modifies, but it also affects in a systematic way several other dimensions of its meaning (function, origin, sortal: Del Pinal, 2015). The second problem is that the idealization is simply untenable on the face of how entrenched polysemy is and how senses of polysemous expressions are typically stored and processed. A literalist hypothesis would predict that a particular sense of each word is accessed easier in all conditions. However, by and large, this is not what is observed, especially in the case of polysemies that pass copredication tests (Frisson, 2015, Schumacher, 2013, Ortega-Andrés & Vicente, 2019).

On the other hand, one can also think that truth-conditional semantics has the resources to explain meaning shifts and provide “intuitive” truth conditions without invoking pragmatic operations. A number of rules have been proposed to account for regular polysemies under names such as “meat grinder” or “universal grinder” (Copestake and Briscoe, 1995), and some authors make massive use of coercion to explain how a literal denotation can be coerced into a different meaning if the type of the literal denotation clashes with the selectional restrictions of the surrounding linguistic material (Asher, 2015). This kind of reaction to the variability argument tries to show that the denotational semantic apparatus is not as meager as the skeptics, as well as the pragmaticians, think. The approach, however, shares some problems with the pragmatic approach, since there is evidence against any kind of general literalist approach. In

many cases (and especially in the cases of copredication), switching from one sense to another of a polysemous expression is smooth, there being no trace of an operation such as coercion (Frisson, 2015, Schumacher, 2013, Ortega-Andrés & Vicente, 2019).

Another way of responding to the skeptics challenge is to hold that, contrary to what they claim, there are indeed entities in the world with the adequate profile to be denotations of word-types. Thus, *London* can be said to denote a complex entity formed by *aspects* or parts that specify a place, a political institution, a set of inhabitants, etc; *school* can be said to denote a complex entity formed by aspects that relate to a certain building, an institution, the people who run the institution, the kids that go to a certain building each day to attend classes, etc. According to this view, the ontology of the world includes complex entities formed by simpler entities: a statue is a complex formed by a piece of matter and some structure or organization; a human being is – maybe- a mongrel formed by a body and a person, etc. Such complexes have distinct parts, each with its own persistence conditions; however, when they are together, they can be considered a single entity. Word-types denote these complex entities, but we can also use the words that denote these complexes to denote only parts of them. In a sense, the literal meaning of a word for, e.g., a city, is a complex entity. However, we can also single out a part or an aspect of that complex entity and refer only to that particular aspect, as when we say *London is huge* or *London is unfriendly*. The selection of these particular aspects is mediated by some semantic mechanism<sup>6</sup>: *huge* is a predicate that requires a “place” argument, whereas *unfriendly* requires an “animate entity” argument. It is relatively easy for us to believe that, say, John is a person *and* a body. When we say *I met John*, we say that we met the whole thing. But when we say *John is fat* we refer to his body, while when we say *John is kind* we refer to the person. Similarly, *book*

---

<sup>6</sup> Although explaining how this semantic mechanism works is no easy task (see Asher, 2011, Gotham, 2016).

denotes an informational content and a physical object. When we say *I have a new book in the market*, we refer to the whole complex; when we say *the book is heavy but interesting* the book we talk about is still the text and tome complex, which offers its two aspects for separate predication. But in *the book is interesting* we only denote the text aspect, and in *the book is heavy* we refer to the tome aspect.

Pustejovsky (1995) introduced a new type in type theory, “dot types”, which refer to complex entities: dot objects. Dot objects are formed by two or more different objects that belong to different types, which are then called the ‘aspects’ of the dot object. The denotation of *book* is *text•tome*, the denotation of *lunch* is *food•event of eating the food*, etc. The idea that has been pointed at in the previous paragraph is to consider that our word-types may denote, or may be about, dot-objects, dot-objects being part of the world we talk about. Some formal ontologists (Arapinis, 2013) as well as formal semanticists (Gotham, 2016) have indeed pursued this line. I will criticize this view in the next section.

However, it is also possible to think about these dot objects (and other similar posits) not as things in the world, but as descriptions or representations of conceptual structures in our minds. In this view, which is a yet third way to react to the attack of the skeptics, word-types do not denote entities in the world, but stand for conceptual structures that offer different possibilities of denotation. This is the line I am going to pursue here. The picture I will argue for is the following: the meaning of a word-type is a concept, a concept being a body of knowledge about a certain category stored in long-term memory. Concepts understood in this sense are structured mental entities that support different ways of categorizing the categories they are about and of supporting inferences. Categorization and inference can be based (at least) on theory-like or prototypical knowledge, as well as on stored exemplars or on idealizations (Machery,

2009, Weiskopf, 2009, Murphy, 2002, 2016, Rice, 2014, Vicente & Martínez-Manrique, 2016). This variety of ways of categorizing explains why words can be used to denote only partially overlapping categories (say, water perhaps including tea but excluding water from a well vs. water including water from a well but excluding tea). However, concepts also provide a different kind of possibilities of denotation, related not to different ways of categorizing certain entities or substances, but to different pieces of information stored in the concept or knowledge structure. These are the cases of *London*, *book*, or *school*. In these cases, word tokens have denotations that do not overlap; rather, they refer to different kinds of entities in the world. But such denotations are represented by highly salient parts or aspects of a concept which, given the activation they receive, are easily targeted by the use of the word the concept is related to.

I will try to explain how these two different kinds of possibilities of denotation (overlapping vs. not overlapping) can be accounted for. My focus, however, will be the non-overlapping cases, which I will examine by looking at the *school* example. I will provide a sketch of a conceptual structure that can be thought as the lexical meaning of *school*, which accounts for its different word-token denotations. This view should be able to provide a plausible account of how lexical meaning connects to (word-token) denotations. The most difficult issue to tackle in such a picture is co-predication, i.e. those cases where, apparently, we refer to the complex entities mentioned above. Co-predication is a problem also for the authors that take a more Strawsonian view on denotation (i.e., language is not representational, but individuals do refer to worldly entities by using language). I will suggest that co-predicational sentences are compilations of complex semantic structures. It still has to be explained when these compilations are available: I will suggest that, typically, it is a matter of how certain

aspects of a concept enter into co-activation patterns. The general upshot of the approach, in any case, will be the following: whereas lexical meaning is not denotational in a simple, direct, way, it does determine a set of possible denotations. The critics are right about the claim that word-types do not have denotations, but not about the argument they mount on the basis of that claim. That word-types do not have denotations does not mean that lexical meanings do not contribute to determining truth-conditions or that semantics is not in the business of explaining word-to-world relations. I will restrict my discussion to the case of nouns. It is possible that other classes of words (e.g., verbs) do not behave in the way that nouns behave. It is credible that verbs have the kind of scarce or thin meaning that Pietroski says all words have, encoding only general or schematic information of the kind that Levin and Rappaport Hovav have been pointing at for a long time (Levin & Rappaport Hovav, 2013; see also Allott & Textor, 2017). This possibility suggests that semantics can be more complicated than expected, such that the truth-conditions of a given sentence arise from an interaction between the rich meanings of nouns and the schematic meaning of verbs and other classes of words. For the most part, I will not enter into this issue here, though, as a reviewer reminds me, cases analogous to co-predication can also be found in the verbal domain: e.g., *Elena writes movingly but illegibly*. In any case, my claims will concern the semantic knowledge associated with nouns, which is, after all, where the critics have lately placed the battleground<sup>7</sup>.

---

<sup>7</sup> What can be called a “rich meanings” or multidimensional account of verbs can be found in Zeevat et al. (2017), as applied to the polysemy of *fall*, and especially in Loebner’s (ms) cascade account of verbs meanings. According to Loebner, verb meanings can be explained in terms of levels of what I will here call “realization”. A case in point is the example above of *writing*, which can denote different actions that are in realization relations. Loebner makes use of frame semantics to flesh out his proposal. Frame semantics are also close in spirit to the proposals here presented, though I prefer not to commit to any particular account in linguistic semantics (not even to Pustejovsky’s, which I will take as a starting point later on). Thanks to a reviewer for directing me to Loebner’s work.

*Dot objects and concepts as routes to denotations*

It is certainly weird to believe that there is an entity that is a place, a political entity, an economic centre, a way of living, and maybe a football team, all at the same time. It may make sense to believe that there is an entity that is a person and a body, or a text and a tome, but the denotational variability associated with a word is difficult to handle in terms of complex objects<sup>8</sup>. Think for instance about *school*: *school* can stand, at least, for:

- (1) a building: *the school has been painted in red*;
- (2) an educational process: *she went to school*;
- (3) a daily event: *school starts at 9:00*;
- (4) the rules of the institution that is associated with the building: *the school has prohibited wearing hats in the classroom*;
- (5) the people who run the local institution: *I have talked to the school about it already*;
- (6) the kids that attend a particular local institution: *Today the school went for a visit to the Cathedral*.

There is a unity to all these different denotations, for all of them relate to a certain social artifact that is there for the purpose of educating people, but it is unclear that this unity can be captured in terms of a complex object that is a compound of the different denotations.

The literature on dot objects discusses whether the alleged dot objects can be counted as such. Asher (2011), for instance, holds that we can only count informational books or

---

<sup>8</sup> The *book* case is not so easy either: in *yes, the book is beautiful but not credible* we seem to denote two different senses associates to the “text” meaning: say, the writing and the plot. Thanks to my student Marina Ortega-Andrés for pointing this to me.

physical books, not pairs formed by physical objects (tomes) and contents (texts). From this he infers that dot objects, understood as complex entities, do not exist (if you cannot count them, they don't exist). Gotham (2016), in contrast, claims that in *There are three red interesting books* we count physical objects paired with the contents they realize. Arapinis and Vieu (2015), on the other hand, try to establish conditions that the physical objects and the contents have to fulfill in order to be the objects of counting as complex entities. I think that, pressing as this counting objection is, there is an even more pressing objection against the existence of the complex entities that allegedly figure as denotations of polysemous terms such as *book* or *school*. This objection against the “ontologization” of dot objects revolves around the idea that the persistence conditions of the wholes that inherently polysemous terms allegedly denote are unclear. If we say (Chomsky, 2000): *London is so unhappy, ugly, and polluted that it should be destroyed and rebuilt 100 miles away*, it seems that we are claiming that the alleged whole would survive even if only one of its parts (its reconstructed buildings and streets) would survive (nowhere is it said that inhabitants and councils are moved 100 miles away). However, it seems that the alleged whole should not exist when only one of its constitutive parts persist: if we say that a human being is constituted by body and mind, then the human being goes out of existence if either part does. In contrast, if, in the Chomskyan situation, all Londoners and the London institutions go on exile and decide not to move to the new location, we can also say that London is wherever the Londoners and the institutions are. So, if London could persist with the persistence of any of its parts, and parts can persist independently from each other, then we have too many Londons<sup>9</sup>. This, I take it, is a serious problem for defenders of the “complex entities” approach.

---

<sup>9</sup> Riaño, a village in the North of Spain, was flooded after a dam was constructed. The Government built a

I will take a very different route: I will attempt at explaining facts about co-predication in terms of our conceptualization of the world, a point of view that can also explain the *water* and the Travis's cases (although I will not say much about these cases here). This means, to begin with, that the proposal has more explanatory power than any other proposal that tackles each of these phenomena with different tools. As said above, a concept, according to the standard view in psychology, is a body of knowledge stored in long-term memory that is involved in higher level cognitive tasks, especially in categorization and inference. Concepts have some internal structure and are not isolated from each other. There are different accounts about how concepts are structured: in terms of exemplars and a similarity metric, in terms of prototypes that abstract statistical information from exemplars, in terms of theory-like or causal structures, or as hybrids of all these different structures –and maybe more, such as ideals (Murphy, 2002, Weiskopf, 2009, Machery, 2009, Rice, 2014, Bloch-Mullins, 2017)-. Here I will assume that a concept typically includes prototypical and theory-like information, maybe meshed together (Hampton et al. 2009, Rice, 2014, Bloch-Mullins, 2017), and can also recruit information from stored exemplars and idealizations (see Vicente & Martínez-Manrique, 2016).

The notion of concept that (some) psychologists use is said to be different from the notion used in philosophy (Machery, 2009). The following difference will be relevant to

---

new village by the side of the lake that emerged, but many people refused to go to the new Riaño. Suppose that all the inhabitants of the old Riaño, including the major and the town council, refused to move to the new location. In that case, the following three assertions would have been true:

(1) Riaño is now in the middle of the lake.

(2) Riaño is an ugly village.

(3) Riaño refused to move to the new location and ended up settling down in a different place.

Nobody would utter (1-3) in a row, but each of (1-3) says something that is true. One cannot react, to any utterance of (1-3), by saying: 'well, strictly speaking, Riaño does not exist anymore'. My point, thus, is that, contrary to intuitive ontological principles, the defender of dot objects has to commit to the view that a particular dot object would persist even when its constitutive parts are pulled apart, with the consequence that the dot object entity is then able to persist in several different entities.

our concerns: whereas psychologists identify concepts as bodies of knowledge that have certain roles in higher order cognition, philosophers have it that concepts are constituents, or building blocks, of thoughts. The view that will be defended here is that concepts in the psychologist's sense do not form part of thoughts; rather, the building blocks of thoughts will be parts or aspects of the concepts psychologists-sense.

The hypothesis that I want to pursue is that nouns stand for concepts. In other words, the meaning of a noun is a concept. Concepts are representational entities: they are bodies of information about a certain category. However, concepts also offer different possibilities for denotation, depending on what kind of information the thinker focuses on and is therefore brought to working memory. To use an example borrowed from Machery and Seppälä (2010): there is a GRANDMOTHER concept, which is about grandmothers. This concept includes theory-like information about what grandmothers are -how they come into being-; but it also has prototypical information about grandmothers. Upon seeing an old lady, one can categorize her as a grandmother based on prototypical information only; accordingly, her thought "there is a car approaching a gentle grandmother" will have a prototypical representation of grandmothers as a constituent. The denotation of that representation in thought, and the denotation of the word-token *grandmother* used in giving voice to that thought, is not the set of actual grandmothers, but the set of old people that look like typical grandmothers.

This kind of view can easily explain the *water* case: *water* can denote H<sub>2</sub>O, thus excluding the liquid from the well, when categorization is guided by the theory-like, essentialist, structure. But it can also stand for a liquid that looks enough like the stuff in rivers and lakes when categorization uses the prototype structure<sup>10</sup>. An adequate

---

<sup>10</sup> It also explains cases like these, where the words in italics take a prototypical denotation:  
(1) Your friend is very *German*

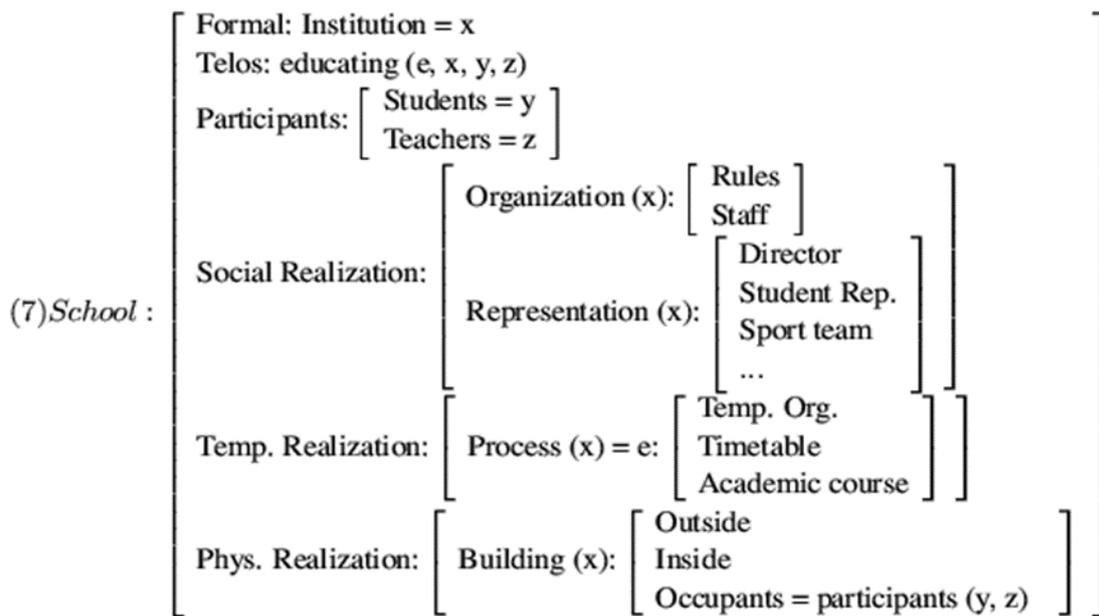
explanation of Travis' cases requires a paper on their own, but proposals such as Vicente's (2012, 2015) and Del Pinal's (2017) can serve as illustrations of how conceptual structures can account for the "green leaves" case. Del Pinal (2017), for instance, uses a slightly revised version of Pustejovsky's *qualia* structures (Pustejovsky, 1995) in order to explain the behavior of a number of adj+noun and noun+noun constructions. To put it very briefly: a *qualia* structure (Del Pinal's version) is a conceptual structure associated with a noun that includes information about the kind, the origin, the constitution, the stereotypical appearance, and the function, of the entity the noun denotes. According to Del Pinal's view, modifications can apply in principle to any of the *qualia*. Thus, *green N* can be three ways ambiguous: greenness can be predicated of the origin of the denotation of the nominal, of its associated stereotypical appearance, and of its constitutive structure. When green modifies ORIGIN we get the reading "naturally green"; when green modifies APPEARANCE the reading is "looks green"; and modification of CONSTITUTION in principle accounts for the variability observed in our judgments about how much of a certain object has to be green to be considered green. If an object has parts (that's what CONSTITUTION refers to), then it may be enough that some parts of the object are green for it to be considered a green object (e.g. a green apple has to have a green skin, a green tree is not expected to have a green trunk, etc) .

### *The school case*

---

(2) The platypus is more a *mammal* than a *bird* (Sassoon, 2017). Cases like (1) have been typically treated as cases of coercion, where the denotation of the predicate, which is not gradable, is turned into a gradable property: the prototype/stereotype associated to Germans. However, both in (1) and in (2), cases that Sassoon (2017) uses to assign a prototypical meaning to common names, can be seen as exhibiting that kind of polysemy that results from having different conceptual structures, or ways of categorizing, associated with a particular word.

In this section, I will present a somewhat simplified conceptual structure that represents the knowledge associated with the multiply polysemous word *school*. The structure attempts at capturing a relevant part of the knowledge that we have about schools, as well as how the different aspects of the informational structure relate to each other. We can begin by noting that a school is, categorically, an institution, and that its function is educating. Now, if that is the case, a typical school, being an institution, will need to have a physical realization (i.e., a physical object that hosts it and, in that sense, makes it real), and a certain social implementation or organization, including certain ways of being represented in society at large. Given the particular function that schools have, it is also expected (i) that schooling takes some time, which can be thought of as a process; (ii) that this process is in turn temporally organized, and (iii) that schools have students and teachers as participants or “inhabitants”. All these features, aspects or pieces of knowledge, are linked to, and to some extent depend on, our conceptualization of school as a social institution whose purpose is that students learn and socialize. A way of seeing how these pieces of information are all kept together or conform a coherent, robust, structure, is to think about all of them as specifying different realizers or implementers of an abstract entity (i.e. as things that the institution requires to be actualized). So, the knowledge structure is organized on the basis of our understanding of what a school is and how it is actualized in the world. A (Pustejovsky-inspired) way to represent the informational structure that captures the lexical meaning of *school* is the following:



There is a principled reason to include some information in the structure and leave some other information outside: the information that is included is all related to the realization (the making real or actual) of the institution. It captures what is *typically* required to have an institution with that function running. In order to have an institution for educating young people one needs some kind of organization, rules, staff, students, a location that is usually a building; also, one also needs to structure a schedule for teaching, organized in terms of hours, days, weeks, years, etc. Besides, any institution requires representatives that represent the institution in different social forums. The concept is structured and held together via explanatory relations: the fact that a school is the kind of institution that it is explains that there is a process, that the participants are teachers and young people, that there is some physical realization, typically a building, hosting the institution, etc. However, it is usually the case that entities and events can be realized in different ways. So there may be various ways in which the telos of the institution can be realized. While it is certainly typical to locate schools-institution in buildings, one can find schools-institutions in barracks, for instance. In such a case, the

word ‘school’ will apply to the barracks. So, the proposal is that a word that designates an institution (and, in general, any kind that requires realizers) will also denote its different realizers or implementers. The structure in (7) represents the realization structure of the school concept and assigns default values to each item in the structure<sup>11</sup>. For those familiar with the concepts literature, it represents some kind of integrated theory-prototype conceptual structure, since it captures explanatory relations and prototypical knowledge at the same time.

This kind of account, besides explaining polysemy and co-predication (see below), is able to explain the linguistic facts that the Pustejovskyan *qualia* structures are supposed to explain, without resorting to the dubious dot objects or types. Pustejovsky (1995) introduced the *qualia* structures to account for the different ways in which predicates can modify nouns: a fast car is a car that drives fast, while a plastic car is a car made of plastic. In the first case, the modifier regards to the telos or function of the car, while in the second case it regards to its constitution (see above on Del Pinal’s treatment of Travis cases). In a parallel way, the different kinds of predicates that ‘school’ combines with are traceable to the knowledge structure in (7): the school (building) can be in ruins, it (process) can be expensive, it (rules) can be strict, etc. The structure in (7) implies a modification to the Pustejovskyan schema, where only four *qualia* are always considered (without much more motivation than reflecting the four causes distinguished by Aristotle: see Moravcsik, 1975). Although this is an issue that would take us too far, it is probably sensible to assign different kinds of knowledge structures/concepts to different classes of entities. Like in the original *qualia* structures, the aspects to be taken into consideration are the aspects or features that characterize a certain kind. However,

---

<sup>11</sup> Again, it is interesting to compare the approach to Loebner’s study case of *writing*. The high-level action of writing (as when one writes movingly) requires realizers. The typical realizer of that action used to be moving a pen in certain ways on a sheet of paper. Perhaps that is still the prototypical realizer of writing-high-level, even though it is now more usual to realize that action by typing.

not all kinds of entities are characterisable in the same way. In the kinds that concern us now, characteristic features include features that typically realize, make actual or implement the kind (plus: they are typically available for being the objects of predication). This holds for institutions and their different ways of being part of the world (as social organizations, as part of society, as physical entities, etc.), as well as for countries, informational contents, drinkables, and so on. Drinkables, for instance, require containers to be “actualized” as such drinkables: this is the reason why it would make sense to include a “container” aspect in the knowledge structure corresponding to ‘beer’, for instance. Informational contents, in turn, require physical realizations, which explains that reference to a physical object will be in the knowledge structure associated to ‘book’, ‘letter’, etc. In sum: it seems that more features than the four Aristotelian causes are relevant (i) to characterize a kind, and (ii) to account for predication. Some of these features or aspects relate to ways of making certain (typically abstract) kinds real. Thus, the proposal in (7) is not simply an after-the-fact structure that accounts for the different senses the word *school*. However, a structure like this does account for the typical senses of *school*, namely:

(8) The school [building] is on fire.

(9) She went to school [process].

(10) School [timetable] starts at 9:00.

(11) The school [rules] has prohibited wearing hats in the classroom.

(12) I have talked to the school [director, staff] about it already.

(13) Today the school [participants] went for a visit to the Cathedral.

(14) The school [sport team representation] has won the championship again.

A structure like this locates all these different aspects in a single representation. How are these different aspects of a knowledge structure/concept turned into senses of a polysemous expression? The response depends on the general answer we give to why polysemy is so widespread. A fair bet is that polysemy (using the same word to mean different but related things) makes communication and learnability easier and more efficient (Falkum 2011; 2015; Xu et al. 2017). Polysemy is a way to comply with the maxim of Relevance that guides communicative exchanges in general: a way to maximize cognitive effects while minimizing processing cost (Falkum 2011). Applying this general view to the polysemy of *school*, we can say that communicative exchanges about the building where students go, the process they engage in, etc., are facilitated by the use of the same word, *school*. These different contents are already part of a coherent knowledge structure, which means that they are easily accessible to both speaker and hearer. The use reinforces the coherence of the structure and the accessibility of the different senses, but it is reasonable to think that the meanings are easily accessible from the start. We can say, thus, that the information structure provides different possibilities for denotation for the word *school*, and that speakers and hearers are aware of this.

The consensus in psycholinguistics seems to be that the polysemous senses of a word are stored in one single lexical entry. At least, the thesis seems to hold for closely related senses, given that closely related senses co-activate each other instead of competing for activation. This co-activation, on the other hand, may last for 750 ms. or more (McGregor et al. 2015). Now, the senses of a polysemous expression such as *school* are particularly closely related. There may be some dispute about the extent to which the 'one representation' thesis applies to all kinds of polysemous words, but it surely applies to *school* and similar cases. The next step is tentative, but seems to have

some empirical support: whereas many different polysemy patterns show sense co-activation, the more related the aspects are, the stronger the co-activation pattern they form (see Schumacher, 2013, forth.). That is, if two, or three, or  $n$  senses are very closely related, it will be usually the case that the activation of one of them results in a high and enduring activation of the others.

In what follows, I will try to explain why senses that belong to structures like (7) may form especially coherent activation packages. The idea will be that these senses are more closely related than any other set of senses because we picture the entities that they denote very closely related by means of realization or actualization relations.

### *Co-predication*

The notion of activation packages is the starting point for an account of co-predication patterns. Co-predication occurs when the same polysemous nominal expression has simultaneous predications selecting for two different meanings or senses. Consider:

(15) The school that caught fire was celebrating 4th of July when the fire started.

(16) Brazil is a large Portuguese-speaking republic that is very high in inequality rankings but always first in the FIFA ranking.

(17) The city has 500,000 inhabitants and outlawed smoking in bars last year  
(Asher 2011)

(18) The best university of the country has caught fire.

Co-predication generates some puzzles that must be solved, apart from the puzzle that it generates to the truth-conditional semantics enterprise. Here I want to focus on two

issues, also relevant to the issue about truth-conditions: (i) why there are some senses that co-predicate and others that do not, and (ii) how we interpret co-predicative sentences.

Let's take (15): when the reader reads (15), the word *school* activates the whole informational structure (7), and, with it, all the different aspects. The predicate *caught fire* selects the building aspect of *school*. However, this does not mean that the other senses decay –that is, they do not get deactivated-. They are all still active when the reader encounters the next predicate: *was celebrating 4th of July*. Thus, she finds no problem in selecting another sense that would comply with the selectional preferences of the predicate, namely, the sense “participants”, which targets teachers and students. There is, in principle, no difference between co-predication and anaphoric binding: in both cases what is required is that the reader or hearer can select the sense that she needs in order to comply with the selectional preferences of the surrounding linguistic material. The initial answer to the question (i): what makes co-predication easy for some polysemes and not for others, is that co-predication will be easy and natural if the senses form an activation package, which, in turn, will occur when the senses are particularly closely related (as in the case of *school*). In polysemies that do not allow for co-predication (e.g., ‘the ham sandwich is impatient [sense: person] and cold [sense: sandwich]’), we see that the activation of one sense involves the de-activation of the other, and so that they do not form a package (see Schumacher, 2013 for the experimental evidence). More controversially, the same seems to happen with the marginally good ‘Tim drank the bottle [content] and broke it [container]’. In such cases, where we see some inhibitory effects, we also observe that the entities that the senses denote are not related by realization relations (e.g., a bottle (container) is not made actual by a content).

This proposal has some similarities with what Arapinis (2013) and Arapinis and Vieu (2015) propose, except that they want to argue for the existence of complex entities. Arapinis and Vieu (2015) take a mereological approach towards dot objects and address the question of what glues together the aspects of a dot object (e.g., school-as-institution and school-as-building) so that aspects can actually be considered parts or constituents of a complex entity. Their response, putting it briefly, is that we have dot objects when (a) the aspects constitute the entity; (b) the aspects are in some coincidence relation; and (c) the aspects are linked by some constant dependency relation that ensures and explains that they are typically in coincidence. The present approach to co-predication can be seen as a way of psychologizing Arapinis and Vieu's ontological proposal. In the current view, co-predication and anaphoric binding are particularly facilitated because senses that denote entities that are linked by realization relations plausibly form an activation package. If denotations of senses are related in this way, then thinking about one of them makes thinking about the others easy. This is, at the end of the day, what explains that these senses form activation packages: they form activation packages because we know that the entities they denote are related in a special way, and in such situations, representations forge strong association links. In a nutshell: there is a strong link between the different denotations that facilitates co-predication via sustained activation of different senses.

Note, however, that the ultimate explanation of what makes co-predication easy for some polysemes and not for others, is associationist: the activation of one sense activates another provided that their denotations are in a relation such that thinking about one facilitates thinking about the other. For this reason, typically, representations of denotations that are in realization relations will activate each other. Yet, there may be more requirements that two or more senses have to meet in order to actually persist into

a mutual sustained activation dynamic. For instance, it seems that the representation of the events (i.e. whole situations) in which the denotations participate need to fulfill some requirements. For instance, examples (19a-b) and (20a-b) suggest that the events have to be spatially or temporally coincident for co-predications to be acceptable.

(19) ?? a. On Monday morning, the school convened to set new rules against plagiarism. Later that day, it was re-painted<sup>12</sup>.

b. The school convened to set new rules while it was re-painted.

(20) ?? a. The school caught fire when it was on excursion.

b. The school caught fire when it was celebrating 4<sup>th</sup> of July.

What we see in (19a) and (20a) is that temporal or spatial displacement of the events involved makes interpretation difficult. Why could this be so? Phenomenologically, it seems that (19b) and (20b) are good because the reader can easily access a (perhaps imagistic)<sup>13</sup> representation of the second event and its participants, facilitated by the representation of the first event and participants. In contrast, in (19a) and (20a) the reader experiences problems about locating the participant denoted by the intended sense of *school* in the second event. It looks like the activation of the different senses holds at least as long as they refer to participants in events that are temporally or spatially coincident or very proximal. The first event sets the spatio-temporal frame, and co-predication will work if the second event describes a situation easily accessible from that frame. Otherwise, the activation of the other senses decays<sup>14</sup>. So a possibility is that all the senses are activated, and so easily retrieved while an event or situation is being

---

<sup>12</sup> Thanks to a reviewer for the example and for pressing the point.

<sup>13</sup> This is the account that Löhr (manuscript) is exploring.

<sup>14</sup> It would be good to be able to support these claims with empirical results, but, unfortunately, there is very little work done on co-predication. The interest that co-predication is receiving lately will hopefully change the situation soon.

described, but then sustainment or decay of the activation depends on there being strong association links between the representation of that event or situation and the representation of a second event or situation. I have to note, though, that there may be problems for such a view. Compare (20a), now (21a), with (21b):

(21) ?? a. The school caught fire when it was on excursion.

b. The school in NYC caught fire when it was celebrating 4th of July in Chicago.<sup>15</sup>

Why is (21b) better than (21a), given that the two events involved are explicitly spatially located in two different places? These are issues that a full account of co-predication has to solve. In this moment, we are very far from having such an account – we are even far from having the stock of examples that we would require to start doing good theorizing-. We are only beginning to scratch the surface of the phenomenon. So, for now, I want to restrict my commitments at this stage to the claim that all the senses in a structure such as the “school” structure above receive a synchronous initial activation and that the activation of one of the senses strongly activates the rest, forming thus an activation package. Then there seem to be conditions that favor suppression of some of these senses that are initially activated, and conditions that do not conspire against sustainment of the activation of all the senses. Until now, empirical research has been focused on studying polysemy in general in the latter kind of conditions, since the objective has been to research on the differences (competition vs. co-activation) between homonymy and polysemy. Hopefully, the interest that co-predication is attracting lately will foster empirical research that can help us know when and why sustained co-activation fails.

---

<sup>15</sup> The example was provided by an anonymous referee of a different journal and of a different paper.

### *Denotations and truth-conditions*

In the view defended here, concepts provide different possibilities for denotation. If a speaker uses the prototypical structure associated with a concept or with a part of a concept, the denotation of her word will differ from the denotation the word would have if she instead uses a theory-like structure (this is the case of *water*). This, of course, has an impact on the truth-conditions of the utterance she makes. Likewise, if she retrieves the physical realization part of the SCHOOL concept, the denotation of *school* will be a building, not an institution or a time in her life, or whatever. Thus, if lexical meanings are concepts, lexical meanings do not have denotations, but only offer possibilities for denotation, i.e., a variety of possible denotations from which the speaker has to select. That is, the denotation potential of a word-type is not explained in terms other than the information stored in the meaning of such a word-type. In this view, thus, a word is associated with a number of denotations, and a sentence with a number of contents that determine different truth-conditions. Usually, the number of contents that a sentence can have will appear to be smaller than the number of denotational possibilities associated with a single word, given that much of the selection of denotations is supposed to be intra-linguistic: in *I have talked to the school*, as *talk* has selectional restrictions for animacy in both of its arguments, there are some denotations of *school* that are ruled out.

Is language a representational device on this approach? Remember that, according to some authors, the language system does not deliver contents with truth-conditions, even if it does deliver units that have a predicative structure, i.e., sentences. In that view, language cannot be considered a representational device, unlike perceptual systems, monkey calls or individual uses of language (Chomsky, 2016). However, according to the view here put forward, language should *prima facie* be regarded as a representational device: sentences do have contents that determine truth-conditions. Note, again, that we have only taken into account nominal expressions. What has been said about the lexical meaning of nouns cannot be directly transferred to verbs, adjectives and adverbs, much less to functional words<sup>16</sup>. Some of these classes of words may have thinner meanings. But it is interesting to think that sentential meanings may derive from the interplay between rich meanings (of e.g., nouns) and thin meanings (of e.g., verbs). If something along these lines is true, then sentential meanings would actually determine what contents, and thus what truth-conditions, an utterance of a sentence may have. It is true that only individual utterances have some determinate, specific, truth-conditions, but, typically<sup>17</sup>, they inherit their truth-conditional content from the sentence-type they are tokens of: they realize or select one of the possible contents that the sentence provides. Alternatively, it can be said that sentence meanings determine disjunctive truth-conditions. They specify a definite number of situations in which the sentence will be true. Speakers and hearers have to select among the different possibilities, so there is a role that pragmatics has to play, but the role that pragmatics plays is that of selecting among alternatives provided by semantic knowledge.

---

<sup>16</sup> Even some nouns can have thinner meanings. Pritchard (2018) analyzes relational nominals such as *target* and argues that they have an abstract schematic meaning.

<sup>17</sup> Typically, because words can also be meaningfully used in novel, ad hoc, ways (Carston, 2015).

Still, co-predication creates a puzzle for this kind of account. The account has it that word meanings are conceptual structures that offer possibilities for denotation by having aspects or parts that can be selected. However, what can we say about sentences such as

(16) Brazil is a large Portuguese-speaking republic that is very high in inequality rankings but always first in the FIFA ranking?

It seems that co-predication generates a problem, since, *prima facie*, *Brazil* in sentences like this “intends” to stand for many parts of the concept simultaneously. An option that I think should be taken seriously is to hold that co-predicational structures are shorthands of more complicated structures. Thus, (16) can be taken to be shorthand of:

(16a) Brazil [place] is a large piece of land & Brazil [people] is Portuguese-speaking & Brazil [State] is a republic & Brazil [economic system] is very high in inequality & Brazil [football team] is always first in the FIFA rankings.

Actually, this explication of truth-conditions is the best option for anybody who wants to hold at least that linguistic utterances have representational contents. As mentioned above, co-predication is not a specific problem for defenders of truth-conditional semantics. It is also a problem for those who, like Pietroski (2005, 2018), following Strawson (1950), want to maintain that only individual speakers refer, i.e. that reference is an individual’s act. In this case, we can say either that an individual using (16) fails to refer or that she is ultimately expressing something like the paraphrase (16a) by some sort of previous compilation of the different senses. That is, a speaker seems to intend to refer to many of the different aspects that form the BRAZIL concept, since she predicates properties that correspond to those different aspects or parts of the concept. However, the way she refers to those different parts is by using a single, compilatory, term that binds (in the psychological sense) them all. At the hearer’s end, the singular tem *Brazil*

activates the package formed by all the different senses or parts of the concept. Thus, it is easy for her to establish a correspondence between the predicates and the entities such predicates ascribe properties to: senses are there, active, for her to retrieve them as she goes on processing the sentence. But then note that this kind of response is also available to someone who wants to hold that word-type meanings provide possibilities for denotation and that speakers select among them.

Providing the truth-conditions of a co-predicative (or anaphoric) structure in this way may be a non-trivial matter. I assume some examples will be more complicated than others. The sentence above (16), is structurally simple, and can be “paraphrased” by the procedure of stacking conjunctions. Initially, it may seem more difficult to provide paraphrases for:

(22) There is no school painted in red that is good enough.

However, one possible explicitation of the truth-conditions of (23) is:

(22a) there is no school [building] that is painted in red such that that school [institution] is good enough.

The problem, of course, is how it can be that *that school* in the final clause refers to an institution, given that apparently it has to make reference to a building. But this is possible: the hearer or reader of (22) can easily retrieve the sense “institution”. The relative clause in (22) does not have the effect of determining that the reference of the subject of the clause is a building, precisely because *school* offers different possibilities for denotation, all of which are active at the moment of interpreting the relative clause. This means that when the speaker has to find a meaning for *school* in the relative clause she can select the institution sense, give *school* the semantic value: [that institution: institution located in building painted in red], and thus get the truth-conditions in (22a).

Actually, assigning these truth-conditions to (22) is no more problematic than assigning its corresponding truth-conditions to a sentence like:

(23) I talked to the hospital but they told me that I have to wait still another week.

In (23) the reader accesses the sense “staff” and consequently uses the plural pronoun *they* to refer to the staff in the second clause. The sense of *hospital* is the same in both clauses. However, research on this kind of polysemy has shown that switching from one sense to another of a polysemous expression is only a bit more costly than repeating senses (Frisson, 2015), which means that assigning truth-conditions (22a) to (22) should not be more complicated than assigning its corresponding truth-conditions to (23).

I am aware that this way of explaining the truth-conditions of co-predicative sentences may clash with some well-established ideas and practices in linguistics and philosophy of language. For instance, the proposal has to reject the ideas that each NP/DP should have just one denotation and that sentences that are syntactically alike have to be semantically alike as well. The proposal has to reject also that any sentence of the kind ‘a is F’ entails an existential ‘there is an x that is F’. From ‘the school round the corner is big but not what we like for our kids’, it does not follow that there is a something that is round the corner, big, and not what we like for our kids’. It is therefore understandable that semanticists have not taken the turn that I commend here. Yet, my concern here has been with whether the truth-conditions of co-predicative sentences can be said to depend on semantic knowledge. Following Pietroski (Pietroski et al., 2009), it is possible to link truth-conditions with “verification procedures”: the truth-conditions of a given expression are linked to the way we actually try to verify whether the sentence is true or not. In the case of co-predicative sentences, it seems reasonable that we verify whether they are true by picking apart the different senses and pairing them to

their respective predicates. The claim then is that to do this we do use semantic knowledge, i.e., that the verification procedure that we use to assign truth-conditions to co-predicative sentences resorts to our lexico-conceptual knowledge associated with nouns that co-predicate.

### *Closing Remarks*

Chomsky (2016) (see also, Berwick and Chomsky, 2016) holds that there are two mysteries associated to language: one is how recursion appeared, and the other is where the atoms of meaning, the units that recursion operates on, came from. His view is that, whereas animal signs are referential, the basic units of human languages are mind-dependent, perspectival, concepts. Thus, he insists that the semantic properties of human language cannot have evolved from animal signing –or from any other animal cognitive capacities-. One of his major arguments is based on copredication, since copredicative sentences have NPs that cannot be said to refer to anything in the world. Followers of Chomsky such as Pietroski, Glanzberg (2014) or Yalcin (2014) interpret the import of the phenomenon somewhat differently. In their view, co-predication and variability show that meanings have to be underdetermined, to the point that they could be mere indexes to concepts.

In the view presented here, the semantic properties of language could draw from our conceptual structures. Perhaps there is nothing like these conceptual structures in animals, though there is plausibly some continuity between animals' concepts and human concepts (bodies of knowledge stored in long-term memory). However, this does not imply that our language is not referential, that the meanings of the words of our language are mysterious in any sense, or that they are underdetermined, schematic, or

indexical in nature<sup>18</sup>. The way in which signs connect with references is more complicated than customarily assumed, but still, language can be said to be representational. A lexical meaning, that is, the meaning associated with a word-type is a concept that offers different possibilities for denotation. The meaning of a sentence depends on the possibilities for denotation offered by words, the selectional preferences of such words, and the proper structure of the sentence. In some cases, it may look like certain word-tokens cannot have a denotation: instead of selecting one of the possibilities offered by the lexical meaning of the word-type, the speaker seems to be intending to refer to several of these denotations at the same time. However, the world does not contain entities that can be formed by an operation that takes as input such denotations. In these cases, speakers form a compilation of senses with different denotations, a compilation that is made possible by an especially high degree of co-activation, but refer to each of the denotations that form part of the compilation. That is, senses are compiled, though referential intentions relate to each of the different senses in the compilation.

## **Acknowledgments**

Versions of this paper, or parts of it, have been presented at the conferences CONTEXT (2017) and Semantics and Philosophy in Europe (2017), also at the Workshop on Ambiguity and Context-Sensitivity in Berlin, and at a talk at UCL. Thanks to audiences

---

<sup>18</sup> An anonymous reviewer suggests that the view here presented may not differ from Chomsky's much after all. I think it is dissimilar in at least two ways: (i) it is not internalist, since I try to explain how words relate to denotations; and (ii) it is not "mysterian" in any way. Chomsky has it that meanings emerged at some point in evolution, unconnected to communication systems in animals, basically because he thinks lexical meanings are not referential. What I try to say is that lexical meanings are referential, although multiply so. That a word is polysemous does not mean that it lacks a reference, only that it does not have a single reference. On the other hand, many of those who endorse a Chomskyan (internalist) semantics program take it that co-predication and variation show that meanings are underspecified or underdetermined representations that connect to concepts. I do not think that this is Chomsky's own (elusive) view, though, again, the view here presented differs from this other view in that meanings are taken to be multidimensional concepts directly. Underdetermined representations are unnecessary intermediate stations and it is unclear what role they could play in semantic theorizing and interpretation.

in these events, especially to Tim Pritchard, Robyn Carston, Manuel García-Carpintero, Elliot Murphy, François Recanati and Emanuel Viebahn. Thanks also to John Collins, Elena Castroviejo, Guido Löhr and two anonymous referees from *Erkenntnis*. Very special thanks to Marina Ortega-Andrés. Discussing co-predication and related issues with her has shaped up my views significantly.

## References

- Allott, N. & Textor, M. (2017). Lexical modulation without concepts: Introducing the derivation proposal. *Dialectica*, 71: 399-424
- Apresjan, J. D. (1974). Regular polysemy. *Linguistics*, 14(2), 5-32.
- Arapinis, A. (2013). Referring to institutional entities: Semantic and ontological perspectives. *Applied Ontology*, 8, 31-57.
- Arapinis, A. & Vieu, L. (2015) A plea for complex categories in ontologies. *Applied Ontology* 10, 285-296.
- Asher, N. (2011). *Lexical Meaning in Context*. Cambridge: Cambridge University Press.
- Asher, N. (2015). Types, meanings and coercions in lexical semantics. *Lingua*, 157, 66-82.
- Berwick, R. C. & Chomsky, N. (2016). *Why Only Us: Language and Evolution*. Cambridge, MA: MIT Press.
- Bloch-Mullins, C. (2017) Bridging the gap between similarity and causality: An integrated approach to concepts, *The British Journal for the Philosophy of Science*, <https://doi.org/10.1093/bjps/axw039>
- Carston, R. (2015). Contextual adjustment of meaning. *The Routledge Handbook of Semantics*, 195-217

- Chomsky, N. (2000) *New Horizons in the Study of Language and Mind*, Cambridge, MA: Cambridge University Press
- Chomsky, N. (2016) *What Kind of Creatures are We?* New York: Columbia University Press.
- Collins, J. (2017). The copredication argument. *Inquiry*, 60: 675-702
- Copetake, A., & Briscoe, T. (1995). Semi-productive polysemy and sense extension. *Journal of Semantics*, 12(1), 15-67.
- Del Pinal, G. (2015) Dual content semantics, privative adjectives, and dynamic compositionality. *Semantics and Pragmatics*, 8(7): 1–53.
- Del Pinal, G. (2017). Meaning, Modulation, and Context: A Multidimensional Semantics for Truth-conditional Pragmatics, *Linguistics and Philosophy*, 41(2): 165-207.
- Falkum, I. L. (2011). The semantics and pragmatics of polysemy: A relevance-theoretic account. *Doctoral UCL Dissertation*. London
- Falkum, I. L. (2015). The how and why of polysemy: A pragmatic account. *Lingua*, 157, 83-99.
- Forbes G. (2012) On Some Examples of Chomsky's Prospects for Meaning. In Schantz, R. (Ed.) *Prospects for Meaning* (pp. 121-142) De Gruyter
- Frisson, S. (2009). Semantic underspecification in language processing. *Language and Linguistics Compass*, 3(1), 111-127.
- Frisson, S. (2015). About bound and scary books: The processing of book polysemies. *Lingua*, 157: 17-35.
- Glanzberg, M. (2014) Explanation and Partiality in Semantic Theory, in *New Essays in Metasemantics*, ed. A. Burgess and B. Sherman, Oxford University Press, pp. 259-292.

- Gotham, M. (2016) Composing criteria of individuation in copredication. *Journal of Semantics*, 34: 1-39.
- Hampton, J., Storms, G., Simmons, C. L., and Heussen, D. (2009). Feature integration in natural language concepts. *Memory and Cognition*, 37(8): 1150-1163.
- Kennedy, C. & Stanley, J. (2009). On average. *Mind* 118: 583-646.
- Klein, D. E., and Murphy, G. L. (2001). Paper has been my ruin: Conceptual relations of polysemous senses. *Journal of Memory and Language*, 47: 548-570.
- Klepousniotou, E., Pike, G. B., Steinhauer, K., & Gracco, V. (2012). Not all ambiguous words are created equal: An EEG investigation of homonymy and polysemy. *Brain and Language*, 123(1): 11-21.
- Levin, B. and M. Rappaport Hovav (2013). 'Lexicalized Meaning and Manner/Result Complementarity', in B. Arsenijević, B. Gehrke, and R. Marín, eds., *Subatomic Semantics of Event Predicates*, Springer, Dordrecht, 49-70.
- Loebner, S. (ms). Cascades. Goldman's level-generation, multilevel categorization of action, and multilevel verb semantics, retrieved from Researchgate
- MacGregor, L.J., Bouwsema, J., & Klepousniotou, E. (2015). Sustained meaning activation for polysemous but not homonymous words: Evidence from EEG. *Neuropsychologia*, 68: 126-138.
- Machery, E. (2009) *Doing Without Concepts*. Oxford University Press.
- Machery, E. and Seppälä, S. (2011). Against hybrid theories of concepts, *Anthropology and Philosophy*, 10: 99-126.
- Malt, B. (1994) Water is not H<sub>2</sub>O. *Cognitive Psychology* 27:41-70.
- McNally, L., G. Boleda. forth. Conceptual vs. Referential Affordance in Concept Composition. In Yoad Winter & James Hampton (eds.) *Concept Composition and Experimental Semantics/Pragmatics*, Springer.

- Moravcsik, J. M. 1975. Aitia as Generative Factor in Aristotle's Philosophy, *Dialogue*, 14:622-36.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge,MA: MIT Press.
- Murphy, G.L. (2016). Is there an exemplar theory of concepts? *Psychonomic Bulletin and Review* Volume 23, Issue 4, 1 August 2016, Pages 1035-1042.
- Ortega-Andrés, M. & Vicente, A. (2019) Polysemy and Co-predication, *Glossa: A Journal of General Linguistics*, 4(1)
- Pietroski, P. (2005) Meaning Before Truth. In *Contextualism in Philosophy* (G. Preyer and G. Peters, eds.), 253-300, Oxford University Press.
- Pietroski, P. (2018). *Conjoining Meanings: Semantics Without Truth Values*. Oxford University Press.
- Pietroski, P., Lidz, J., Hunter, T. & Halberda, J. (2009). The Meaning of 'Most': Semantics, Numerosity and Psychology. *Mind and Language*, 24, 554–585.
- Pritchard, T. (2018). Analogical Cognition: an Insight into Word Meaning. *Review of Philosophy and Psychology*, DOI: <https://doi.org/10.1007/s13164-018-0419-y>
- Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, Massachusetts: The MIT Press.
- Recanati, F. (2010) *Truth-Conditional Pragmatics*, Oxford University Press.
- Rice, C. (2014). Concepts as Pluralistic Hybrids. *Philosophy and Phenomenological Research* DOI: 10.1111/phpr.12128
- Sassoon, G.W. (2017). Comparisons of nominal degrees. *Language* 93: 153-188
- Schumacher, P. (2013). When Combinatorial Processing Results in Reconceptualization: Towards a New Approach to Compositionality. *Frontiers in Psychology*, 4, 677.

- Schumacher, P. forth. Metonymy. Oxford Handbook of Experimental Semantics and Pragmatics
- Segal, G. (2012) Five Flies in the Ointment: Some Challenges for Traditional Semantic Theory In Schantz, R. (Ed.) *Prospects for Meaning* (pp. 287-307). De Gruyter.
- Travis, C. (1996) "Meaning's Role in Truth", *Mind*, 105: 451-466.
- Travis, C. (2008). *Occasion-Sensitivity: Selected Essays*. Oxford: Oxford University Press.
- Vicente, A. (2012). On Travis Cases, *Linguistics and Philosophy* 35:3-19.
- Vicente, A. (2015). The green leaves and the expert: polysemy and truth-conditional variability. *Lingua*, 157, 54-65.
- Vicente, A. & Martínez Manrique, F. (2016) The big concepts papers: a defence of hybridism. *British Journal for the Philosophy of Science*, 67 (1), 59-88
- Weiskopf, D. (2009) The plurality of concepts. *Synthese* 169:145-173.
- Xu, Y., Malt, B.C., and Srinivasan, M. (2017). Evolution of word meanings through metaphorical mapping: Systematicity over the past millennium. *Cognitive Psychology*, 96: 41-53
- Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In A. Burgess & B. Sherman (Eds.), *New Essays in Metasemantics*. New York: Oxford University Press.
- Zeevat, H., S. Grimm, L. Hogeweg, S. Lestrade, and E. A. Smith (2017), Representing the lexicon: Identifying meaning in use via overspecification. In: Kata Balogh and Wiebke Petersen (eds), *Bridging formal and conceptual semantics*. Selected papers of BRIDGE-14, 153-186. Düsseldorf: Düsseldorf University Press