Truth-conditions for cities, countries and schools

Agustín Vicente

Ikerbasque/University of the Basque Country (UPV/EHU)
Dpt. Linguistics and Basque Studies, Vitoria
agustin.vicente@ehu.eus

Abstract. In this paper will try to show that semantics can explain word-to-word relations and that sentences 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, forth.). Against this view, I will hold that, typically, semantic knowledge can account for the denotational uses of words of individual speakers.

Keywords. Semantics, co-predication, Chomsky, truth conditions.

1. Introduction

The idea that the meaning of a sentence of a natural language determines its truth-conditions has a long story. 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 it accurate. If the perceptual representation is about a black flying object, then the representation is accurate only if it is tokened when there is a 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. Thus, if a sentence of a natural language is representational, its content has to determine its truth conditions.

The venerable 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 of 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 exact the accuracy or appropriateness conditions of the use of a certain word. Examples abound, but probably the most striking 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 H_2O only in the scientific "language game", so to speak[1]. Otherwise, we do not call *water* a cup of tea, even if it may have a higher percentage of H_2O molecules than the liquid that comes from a well (see Malt, 1994, Pietroski, forth.) The accuracy conditions of uses of *London* or *water*, thus, are not univocal or even consistent. The denotation that fixes such conditions, so the reasoning goes, has to be some entity in the world, and there is nothing in the world that is a place, a political institution, a way of life, etc., all at the same time. Also, there is nothing in the world that covers all the uses we make of the term *water*, which are not based on how much H_2O a certain liquid has.

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 been painted (see Vicente, 2012, 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, forth., Yalcin, 2014). This is the thesis I will criticize in this paper. I will try to show that semantics can explain word-to-world relations and that sentences have meanings that determine truth-conditions. 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). Against this view, I will hold that, typically, semantic knowledge can account for the denotational uses of words of individual speakers. I will summarize my position at the

[1] 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).

end of this introductory section. Before that, I briefly survey different ways in which one can react to the critics' attack on truth-conditional semantics.

There are different ways to react to the 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 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 (Recanati, 2010). However, one can also think that formal truth-conditional semantics has the resources to explain meaning shifts and provide "intitutive" truth conditions without invoking pragmatic operations. A number of rules have been proposed to account for regular polysemies (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 selection 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.

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, etc. According to this view, the ontology of the world includes "mongrel" entities formed by simpler entities: a statue is a mongrel 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 mongrels 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 mongrel entities, but we can also use the words that denote these mongrels 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. The selection of these particular aspects is mediated by some semantic mechanism: *huge* is a predicate that requires a "place" argument, whereas *unfriendly* requires an "animate entity" argument. Pustejovsky (1995) introduced a new type in type theory, which he called 'dot objects', a type formed by two or more different types, which are then called the 'aspects' of the dot object. The idea that has been pointed at in the previous paragraph is to consider that our word-types may denote, or

be about, dot-objects, dot-objects being part of the world we talk about. Some formal ontologists have indeed pursued this line (Arapinis, 2013, Arapinis and Vieu, 2015).

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 word, but stand for conceptual structures that have a representational content and 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 and Martínez-Manrique, 2016). This variety in 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 denotational possibilities, related not to different ways of categorizing certain entities or substances, but to different pieces of information stored in the 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 the 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 denotational possibilities (overlapping vs. not overlapping) can be accounted for 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 mongrel 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 more simple sentences. The general upshot of the approach 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 words 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. Sentences are connected to

wordly entities by means of their meanings and individual acts of denotations are based on the semantic knowledge of speakers.

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 to a word is difficult to handle in terms of complex objects[2]. Think for instance in *school*: *school* can stand, at least, for: a building; an educational process (*she went to school*); a daily event (*school starts at 9:00*); the institution that is associated with the building (*the school has prohibited wearing hats in the classroom*); the global institution (*the Government wants to reform the school again*); the people who run the local institution (*I have talked to the school about it already*); the kids that attend a particular local institution (*Today the school went for a visit to the Cathedral*), etc. 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 kids, but it is unclear that this unity may be captured in terms of a complex object that is a compound of the different denotations. Besides, this kind of account seems unable to explain the *water* or the Travis' cases: water is not a dot object formed by H_2O and the rest of the liquids that resemble H_2O to a certain extent; and a green leaf is not a complex formed by intrinsically green leaves and leaves painted in green.

Here I will attempt at explaining denotational variability in terms of our conceptualization of the world. A concept, according to the standard view in psychology, is a body of knowledge stored in long-term memory that is involved in higher order 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.

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 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

[2] 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.

building blocks of thoughts will be parts or aspects of the concepts psychologists-sense.

The hypothesis that I want to pursue is that the meaning of a lexical item is a concept. Concepts are representational entities: they are bodies of information about a certain category. However, concepts offer different denotational possibilities, 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. 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 the prototypical information only; accordingly, her thought "there is a car approaching to a gentle grandmother" will have a stereotypical 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.

Qualia structures and beyond

Puestejovsky (1995) represented the conceptual information stored in the lexical entries of nominals in terms of a *qualia* structure inspired in Aristotle's four types of causes (Moravcsik, 1998). This *qualia* structure stores information about the origin, the function or *telos*, the constitution, and the classification in the ordinate-superordinate hierarchy of the category the nominal picks up. Recently, Del Pinal (2015, ms.) has put to use a slightly revised version of Pustejovsky's account in order to explain the behavior of a number of adj+noun and noun+noun constructions. Del Pinal's structures include a perceptual quale that stores information about the stereotypical appearance of the category. According to his view, modifications can apply in principle to any of the *qualia* roles. Thus, *green N* can be three ways ambiguous: when *green* modifies ORIGIN we get the reading "naturally green"; when *green* modifies PERCEPTUAL 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[3].

So, the different *qualia* roles can be on or off. Now, there is no reason why the list of *qualia* has to be limited to four or five. Also, there is no reason why a *quale* has to await for modification to be turned on or off. Essentially, *qualia* provide ways to cate-

[3] Saying that an (e.g. color) adjective applies to the constitution *quale* associated to a noun, however, may not be enough to know which parts of the object referred to by the noun have to display the property referred to by the adjective in order for the object to be a "propert-ied object". It is usually noted that *red apple* is red-skinned apple, while *red watermelon* is red-interior watermelon. So something else has to be said about the way adjectives select parts offered by the "constitutive" *quale*. I'm not concerned with composition here –these applications of the *qualia apparatus* are only here for the purpose of illustration-; McNally and Boleda (forth.) provide an interesting model of more detailed mechanisms of composition, which complement the idea sketched in this brief description.

gorize: by origin, by appearance, by function, by constitution, etc. They can be used to categorize green leafs in different ways, but they can also be used to categorize birds or cats or water in different ways. If we turn on the BIRD, the CAT, or the WATER perceptual *quale* while turning off their respective origin and, in the case of water, constitutive *qualia*, we get categorization by prototype. If we do the opposite, we get categorization by theory. This explains why *water* (or bird or cat) can denote different things. *Water* can denote H_2O, thus excluding the liquid from the well, or it can stand for a liquid that looks enough like the stuff in rivers and lakes. This kind of structure also explains cases where adjectives and nominals can take a prototypical denotation, such as Your friend is very *German* or *The platypus is more a mammal than a bird* (Sassoon, forth.).

As I say, meanings can have more information than that stored in a Pustejovskyan *qualia* structure. The *qualia* structure is particularly relevant to explaining what content nominals offer for the composition of meanings. However, such a structure does not explain some "simple" variations in the semantic contribution of a nominal. In particular, a *qualia* structure cannot account for many regular polysemies such as those exemplified by the *London*, the *book* or the *school* cases. This is where dot objects enter the picture. Pustejovsky holds that dot objects are complex types formed by two contradictory types, say, physical object and information content (*book*). The way he deals with inherent polysemy is by supplementing the *qualia* structure with the information that the entity that the qualia structure is about is a complex type, say physical_object•informational_content. However, this way of integrating the regular polysemy in the information stored in the lexical entry is somewhat misleading, as it suggests that the denotation of a word such as *book* is a mongrel entity.

However, if we think about concepts as, roughly, structured encyclopedic information, we can think about the aspects that constitute dot objects as specially relevant and accessible *qualia*-like features of a concept, not as conceptualizations of a complex entity. Let me illustrate this idea by looking into a possible representation of the SCHOOL concept.

The concept of school can be thought as a structure that encodes the following information:

SCHOOL:
Kind: Social institution
-*Telos*: for learning, socialization, and enculturation of young people.
Associated prototypical knowledge about how the institution fulfils its role + Ideals, exemplars…
- Temporal realization [associated to the *telos* of the institution]: a process of x years of learning, socialization and enculturation.
Protoypical information about how the process is realized + Ideals, exemplars…

- Occupants [associated to the *telos* of the institution]: kids
  Prototypical and stereotypical information about the occupants + Ideals, exemplars…

-*Physical realization* [institutions have to be physically realized]: building.
  Qualia-like information about the building, including the associated stereotype: how schools look like + Ideals, exemplars…

-*Social realization* [institutions have to be organized in some way]: director and staff.
  Qualia-like information about directors and staff of a school + Ideals, exemplars…

-*Representation* [institutions are represented in society at large in several ways]: football team, basketball team, head of the institution, an elected representation of the kids…

This of course just a rough sketch of what the lexical meaning of *school* can be like and how conceptual knowledge can be represented. Let me now explain a bit this sketchy representation. The concept of school is labeled, or tagged, as *social institution*. A good part of what we find into the concept/lexical meaning of *school* derives from such a labeling or tagging and an inheritance system (Pustejovsky, 1995) or semantic network that puts concepts in relation with its superordinates such that subordinate concepts non-monotonically inherit information from their superordinates. If something is a social institution then, unless some information stored in the concept contradicts it, it will have a function (a *telos*), a physical realization, a social realization, and ways of being represented.

Given that school is a social institution, its *telos* specifies its essence. A way of thinking about the whole concept/lexical meaning of *school* is as a theory-like structure whose features are causally related to its essence (see Weiskopf , 2011, for an excellent introduction to theory-like concepts). From the essence attributed to school follows the idiosyncratic knowledge that schooling takes some time, which can be thought of as a process (relevant for understanding expressions such as *I went to school;* meaning: I completed the process of schooling: contrast with *I went to church*), and that schools have kids as occupants, as well as the particular physical realization, social realization and the ways in which the school can be represented. These features are explained by the essence of school, but are not necessitated by it, so they do not provide, all together, a definition of the SCHOOL concept.

Each node or feature in the theory-like structure is associated to representations about prototypes, ideals, and exemplars of the relevant features. The “process” feature represents default knowledge such as that schooling consists in attending classes at the school according to a certain timetable (present in *school starts at 9*, but absent in *home schooling*). The occupant feature gives access to the ideal of a school kid as a

well-behaved one (ideal in play in *you are such a school boy!*), etc. It is also plausible to think that the features PHYSICAL REALIZATION: BUILDING and SOCIAL REALIZATION: DIRECTOR AND STAFF have their own associated *qualia*-like structure that can store information about constitution, an eventual *telos*, origin and prototypes, stereotypes, ideals, exemplars, etc.

Perhaps it is somewhat artificial to separate theory-like structures and prototypes, ideals and exemplars in this way. Probably theories and prototypes are more integrated, to begin with (Hampton, et al. 2009, Rice, 2014, Bloch-Mullins, 2017). However, this is a discussion that exceeds the purposes of this paper. For the time being, we can do with this, somewhat complex, toy model. For what interests us here is not so much the detailed representation of the knowledge structure as the general idea about how this kind of structure explains our different uses of a certain word. The hypothesis is that speakers target or select particular pieces of information stored in a concept or knowledge structure. These can be thought as parts or aspects of the whole concept. We have seen that speakers can target the prototype associated to *grandmother* instead of its theory-like component (alternatively, we can say that they categorize by prototype instead of by theory). The suggestion is that they can also target the information about the physical realization of *school*, its representation in society, the prototypes or ideals associated to the different subentries (as in *you are such a school boy!* above) etc. They are all salient features of the concept, given our experience with schools in particular, and with this kind of social institutions in general. So they are all active when we think about schools in any of its different senses: this is the main reason why a speaker would use the word *school* to refer to any of these parts and why the hearer will easily understand her. (For empirical results, see Frisson, 2009, 2015, Klepousniotou et al. 2008, 2012).

When we categorize we do not use whole concepts. Rather, we typically categorize something as an *x* (say, a school) if it has some relevant properties associated to our knowledge about *x*, determined by the context. This is what authors working on contextual categorization, such as Barsalou (1999), Prinz (2002), and Casasanto and Lupyan (2015) have shown[4]. We draw from our knowledge in long term memory and retrieve the relevant pieces of information for use in working memory. This means that what we denote, or refer to, either in thought or in language, is not the denotation of a concept, but the denotation of that part of the concept that we have retrieved in working memory (see Vicente and Martínez-Manrique, 2016, for development). In turn, this means that concepts are not the building blocks of thoughts: selected parts of concepts constitute the building blocks of thoughts. This claim can be disputed on terminological grounds: some would prefer to call *concepts* to whatever forms part of a thought. However, if we adopt the psychological terminology, and conceive of concepts as bodies of information stored in long term memory, thoughts are not composed of concepts.

[4] Although Casasanto and Lupyan do not believe that there is anything that corresponds to a stored concept.

In the view defended here, concepts provide possibilities of categorization, i.e., of denotation. If a speaker uses the prototypical structure associated to a concept or to 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 denotational possibilities, i.e., a variety of possible denotations from which the speaker has to select. However, these possible denotations are stored in the lexical entry. 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[5].

2. The puzzle of co-predication

There are sentences, though, that seem to create problems for this kind of account. The account has it that word meanings are conceptual structures that offer denotational possibilities by having aspects or parts –which are categorizational/referential units- that can be selected. However, what can we say about sentences such as

(1) Brazil is a large Portuguese-speaking republic that has deep problems associated with inequality and won five Word Championships?

It seems that co-predication generates a problem, since, prima facie, *Brazil* in sentences like this "intends" to stand for many aspects of the concept simultaneously. As Pietroski, Chomsky, et al. put it, there is nothing in the world that corresponds to this *Brazil*. However, note that this kind of sentences is also a problem for those who hold that referring is something that individuals, not language, do. A particular's utterance

[5] Usually, the number of contents that a sentence can have will appear to be smaller than the number of denotational possibilities associated to 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. Similarly, in *he is very German*, the intensifier selects the gradable, stereotypical, sense of *German* and discards the definitional sense. However, this point is not obvious. Some authors prefer to talk about selectional preferences instead of selectional restrictions (Zarcone, 2014), given that contexts can be concocted where the alleged selectional restrictions are violated.

of (6) is not about anything, since there is nothing in the world that can have the various properties we predicate of *Brazil*.

One way to dissolve the problem created by co-predicative structures like (1) is to hold that co-predicational structures are shorthands of larger, conjunctive, structures. Thus, (1) can be taken to be shorthand of:

(1a) Brazil [place] is a large piece of land & Brazil [GENpeople] is Portuguese-speaking & Brazil [State] is a republic & Brazil [economic system] has deep problems associated with inequality & Brazil [football team] has won four Word Championships.

The idea may not be appealing because paraphrasing is not appealing at least since syntax abandoned deep structures (see Partee, 2015). However, I think this is the best option for anybody who wants to hold at least that linguistic utterances have representational contents. As explained 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, forth.) 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 (1) fails to refer or that she is ultimately expressing something like the paraphrase (1a) by some sort of previous compilation of the different senses. That is, a speaker seems to intend to refer to the different aspects that form the BRAZIL concept, since she predicates properties that correspond to those different aspects. However, the way she refers to those different aspects 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 all the different aspects (which is typical in regular polysemy: see Frisson, 2015). Thus, it is easy for the hearer to establish a correspondence between the predicates and the entities such predicates ascribe properties to. But then note that this kind of response is also available to someone who wants to hold that word meanings provide denotational possibilities and that speakers select among them.

Unfortunately, this cannot be the whole story about co-predication. An account of structures such as (1) has to also provide some explanation as to why some aspects or parts of concepts admit co-predication and anaphoric binding and others do not. This is well known problem in the study of polysemy, and is the reason why dot objects were posited in first place. It was noticed that only some regular polysemies passed co-predication and anaphoric binding tests. For instance, the pattern content-for-container (*beer* to refer to a glass of beer) accepts co-predication more naturally than the container-for-content pattern (*bottle* to refer to the water in the bottle): *Peter put down the beer and drank it a few minutes later* seems more natural than *Tim drank the bottle and dropped it* (Schumacher, 2013). And Schumacher (2013) reports that content-for-container switches involve a less costly processing than the inverse switches. Using the name of a container to refer to its content might involve some sort of coercion (or some other operation), where the container sense is accessed first and then adjusted to the content sense to fit the demands of compositionality (see also Asher, 2011). Typically, if there is coercion, the primary sense (i.e., the one that plays the

role of input of the operation) is no longer available for anaphoric reference or co-predication.

To explain how co-predication was possible at all, semanticists like Pustejovsky (1995) and Asher (2011) decided to introduce a new type: dot objects. The idea then was: co-predication is possible if and only if the different senses are aspects of a dot object. The different behavior of some regular polysemies and others, thus, is that the regular polysemies that do not allow for co-predication and anaphoric binding involve linguistic operations such as coercion, whereas the regular polysemies that pass the tests provide the processor with (at least) two aspects/senses simultaneously. On second thoughts, however, the idea of distinguishing some polysemies from others in terms of dot objects is not very illuminating and lacks predictive power: actually, Asher (2011) characterizes "logical" (dot-object) polysemy as that polysemy which passes the co-predication test. That is, we do not know in advance whether a certain polysemy pattern will allow speakers to resort to sense-compilation. However, the work of formal ontologists such as Arapinis (2013) and Arapinis and Vieu (2015) may provide a way to complement the account in an explanatory way: their basic idea is that senses form dot objects (and thus allow co-predication) when the denotations of such senses hold certain metaphysical dependency relations and there is spatio-temporal coincidence between said denotations. This seems to leave co-predications involving representations unexplained, even though it is a well-established fact that nominals can easily switch between represented things and their representations (Jackendoff, 1992), but maybe the representation case has to be dealt with separately.

Leaving at one side the notion of dot object, which, as mentioned, may give rise to ontological problems, a response to the question as to when sense compilation is possible can resort to this kind of proposals. For it is certainly possible that some senses receive higher activation than others, such that they are always available both for co-predication and for anaphoric binding, it being the case that such senses are typically linked either by dependency and coincidence relations or by a representation relation. That is, it is plausible that parts of concepts whose denotations have some especially intimate relationship, such as spatio-temporal coincidence, tend to receive a stronger synchronous activation. This would explain why they are always "there", all active as long as one of them receives activation, such that one can easily target each of them for predication and anaphora.

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, the semantic properties of human language cannot have evolved from animal signing –or from any other animal cognitive capacities-. In the view presented here, the semantic properties of language draw from our conceptual structures. Perhaps there is nothing like these conceptual struc-

tures 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 or that the meanings of the words of our language are mysterious in any sense. 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 to a word-type is a concept that offers different denotational/representational possibilities. The meaning of a sentence depends on the denotational possibilities 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 denotational 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.

References

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.

Barsalou, L. (1998). Perceptual Symbol Systems. *Behavioral and Brain Sciences* 22, 577–660

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

Casasanto, D., & Lupyan, G. (2015). All concepts are ad hoc concepts. In E. Margolis & S. Laurence (Eds.), *The conceptual mind: New directions in the study of concepts* (pp. 543–566). Cambridge, MA: MIT Press.

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.

Copestake, 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: 1-53.

Del Pinal, G. (ms) A multidimensional Putnam-style semantics for Truth-conditional Pragmatics

Dölling, J. (forth.). Systematic polysemy. In L. Matthewson, C. Meier, H. Rullmann, & T. Ede (Eds.), *The Blackwell Companion to Semantics*.

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.Geeraerts, D. (1993). Vagueness's puzzles, polysemy's vagaries. *Cognitive Linguistics*, *4*(3), 223-272.

Hampton, A. James, Storms, G., Simmons, Claire L., and Daniel Heussen (2009). Feature integration in natural language concepts. *Memory and Cognition*, 37(8): 1150-1163.

Jackendoff, R. (1992). Mme. Tussaud Meets the Binding Theory. 10:1-32.

Klepousniotou, E., Titone, D., & Romero, C. (2008). Making sense of word senses: the comprehension of polysemy depends on sense overlap. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*(6), 1534-1543.

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.

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 H2O. *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. (1998). *Meaning, creativity, and the partial inscrutibility of the human mind*. California: CSLI Publications.

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

Partee, B. (2015) The Garden of Eden period for deep structure and semantics. In *50 Years Later: Reflections on Chomsky's Aspects*, eds. Ángel J. Gallego and

Dennis Ott, pp. 187-198. Cambridge, MA: MIT Working Papers in Linguistics.
Pietroski, P. (2005) Meaning Before Truth. In Contextualism in Philosophy (G. Preyer and G. Peters, eds.), 253-300, Oxford University Press.
Pietroski, P. (forth.). Semantic internalism. In J. McGilvray (Ed.), *The Cambridge Companion to Chomsky*. Cambridge: Camrbidge University Press.
Prinz, J. (2002) *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT Press
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. (forth.) Comparisons of nominal degrees. *Language*
Schumacher, P. (2013). When Combinatorial Processing Results in Reconceptualization: Towards a New Approach to Compositionaliy. *Frontiers in Psychology*, 4, 677.
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: 59-88.
Weiskopf, D. (2009) The plurality of concepts. *Synthese* 169:145-173.
Weiskopf, D. (2011) The Theory theory of concepts. Internet Encyclopedia of Philosophy (on.line, April, 29, 2011)
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.
Zarcone, A. (2014). *Event Knowledge and Models of Logical Metonymy Interpretation*. Universität Stuttgart. PhD Dissertation.