

Compositionality and Concepts—A Perspective from Formal Semantics and Philosophy of Language

Francis Jeffrey Pelletier

Abstract It's no secret that different of the subfields in cognitive science dispute what the correct solution is to various problems that they each investigate in their separate ways. Sometimes this is due to differing antecedent ideas about what is the appropriate way to investigate the phenomenon, other times it is due to differing antecedent ideas about what principles an adequate solution should embody, and still other times it is due to differing antecedent ideas concerning what the dispute is about... as for example when they use the same terminology for different phenomena. This paper is an investigation into these differing antecedent ideas in the realm of meaning and compositionality as they play out in linguistics, cognitive psychology, and philosophy of language. The focus is on the notions of *subjective meaning* and *objective meaning*, and a main conclusion is that there needs to be a “two-factor semantic theory” to accommodate the overall goals of both sides. Some steps are made towards that end, but various previous attempts are argued to have missed the point. In the end it will be shown that, clearly, more work needs to be done along a number of specified dimensions, especially on the subjective side of the dispute.

1 (Some) Deep Background

I am going to be making some remarks about “what philosophers of language believe” and “what cognitive psychologists believe”—as well as some equally general-sounding claims about “what linguists/formal semanticists/cognitive scientists/etc. believe”. Like all such general statements made about a group, mine should not be taken as universally quantified.¹ (It would be almost beyond belief were all

¹In fact, most researchers into the field of generics think that such “generic statements” can be true—depending on the particulars of the case—even if only a relatively small percentage of the relevant population manifest the property in question, especially when this subgroup is “salient” in some way. See the discussion in Krifka et al. (1995).

F.J. Pelletier (✉)

Philosophy Department, University of Alberta, Edmonton, AB T6G 2E7, Canada
e-mail: francisp@ualberta.ca

© The Author(s) 2017

J.A. Hampton and Y. Winter (eds.), *Compositionality and Concepts in Linguistics and Psychology*, Language, Cognition, and Mind 3,
DOI 10.1007/978-3-319-45977-6_3

cognitive psychologists to believe the same theoretical things; it would be an even further stretch to think that all philosophers of language believed the same theory; and so on).

I furthermore will be making claims about the notion(s) of compositionality, and will use my explanation of this/these notion(s) as a (rather meagre) web upon which to link what philosophers of language, cognitive psychologists, formal semanticists, cognitive linguists, etc., believe. In fact though, I think that the major reasons for the difference in approaches to these topics come from two other sources, and not compositionality considered simply in itself—which I see as just a consequence of choices within these other sources. The first source is an ontological commitment to one of (what I will call) Atomism and Wholism, and the second source is a semantic commitment to one of (what I will call) Subjectivism and Objectivism. I will try to emphasize these options as the paper develops. At the end, I hope you'll agree that there are numerous “disconnects”, and that on the other hand there are some rather surprising places where agreement is reached.

I start with compositionality.

2 (Some) Features of Compositionality

There are two conceptions of compositionality that make their appearance in the literature. One version, which I'll call *Ontological Compositionality*² describes the way a complex item is put together. The underlying idea is that the complex is “nothing more than” the collection of its parts (and the parts of these parts, etc.), put together in some perhaps new configuration. Underlying this sort of compositionality is the slogan:

Definition 1 (*Ontological Compositionality*) A whole is a compositional system built up only from materials in the parts, with due consideration about how these parts are arranged.

Consider these two theories (don't worry about their correctness as theories—we're only focussing on whether they express ontological compositionality or not).

Suppose assemblage-1 of neurons is active during task-1 and assemblage-2 is active during task-2. Now consider a “supertask” which intuitively involves doing task-1 and task-2 together. Then: is the assemblage of neurons involved in this supertask made up out of, and only of, those neurons involved in either task-1 or task-2? Or does the new supertask bring an entirely new assemblage of neurons to bear on the new task? Many theorists think that, at least after the new supertask becomes a learned, repetitive action, there is a new group of neurons responsible. And thus, motor neural accounts are *not* ontologically compositional.

²Or perhaps a better name would be *Building-Block Compositionality*.

One view of prototypes³ says that a prototype is a structure that has a number of attribute-value pairs. A question of interest in this area is: Given that prototype-1 has a structure [A: value-a; B: value-b; . . .] and that prototype-2 has the structure [C: value-c; D: value-d . . .], is the prototype of the “combined prototypes” (that is, the conceptual combination of the two prototypes) made up only from the material in the two component prototypes, that is, from the attribute-value pairs that are in one or the other of the prototypes? Many theorists answer ‘yes’ to this question, thereby claiming that this manifests the (ontologically) *compositional* view of prototypes. As an example, the prototype of PET FISH would be (ontologically) made up compositionally from the prototypes of PET and FISH because all the “relevant material” (the attributes) could be found in the these smaller, embedded prototypes, and merely needs to be “re-arranged”.

A name for the underlying theory that holders of Ontological Compositionality believe is *Atomism*—the view (in whatever area is under discussion) that there are smallest items (atoms—that is, whatever the area claims to be its minimal elements) and that all other ontological items (in the area) are larger and larger wholes that are ultimately made up of these atoms put together in some specific way. (Different ways will generate different objects out of the same parts.) The denier of Ontological Compositionality believes in *Wholism*—the view (in the area under discussion) that there are some items whose makeup at least partially consists in things that are not in its parts.⁴

Since Wholism takes it that some complex object (of whatever type) can contain things that are not in any of the parts, Atomists naturally find this quite puzzling. One reaction is:

Those who say that a whole is greater than the sum of its parts either don’t know what a whole is, or don’t know what a part is, or don’t know what being a sum is, or don’t know what being greater than consists in.

To take our example of assemblages of neurons, the Atomists would claim that the alleged complex object does *not* have the initial assemblages of neurons as parts. It was a completely different object that was created by learning or practice. It was not the initial assemblages, rearranged, plus something new added on—to make it into a whole. It was just a different object. So it neither obeys nor violates Ontological Compositionality: we just have two independent objects. Or at least, this is the reaction that holders of Ontological Compositionality would claim.

But it is never totally clear when Wholism applies. According to the Boodin quote in footnote 4, any whole that is generated in accordance with gestaltism violates Atomism. Is a heap of sand, made from putting a bunch of grains into close proximity to one another, a *new whole* over and apart from the grains? Or is it just the grains, rearranged by being placed into the relevant proximity? Is a sentence . . . for instance one which is made up of a noun phrase plus a verb phrase . . . such a whole? Did something new come into existence, a *sentence*? Or is the alleged whole actually *nothing more than* the noun phrase and the verb phrase *put together in the appropriate way*? How can we tell in every case whether there is an *entirely new, different*

³E.g., in Hampton (1979, 1982), Smith et al. (1988); the idea can be seen already in Rosch (1975).

⁴The OED cites Boodin (1939) as the originator of the word ‘wholism’: “Two conceptions . . . namely, creative synthesis (or emergence), and wholism (or gestaltism). . .”.

object coming into being, or just the parts that are arranged appropriately? Are the (computational) connectionists right in saying that *every* property of a connectionist system is an emergent property?⁵ We leave these questions aside at this point.

But compositionally of the Ontological variety is not the only notion in the neighbourhood. When discussing linguistic matters, it is far more common, at least in philosophy of language and formal semantics, to see a discussion that presumes there to be an already-given, well-defined system of the sort that Ontological Compositionality has characterized—namely the syntactic structures of language, treated as a complex object made from atoms (and not as a newly-created whole). The atoms are lexical items and the wholes that are made from them are syntactically-defined larger and larger linguistic units. But if we become interested in some *property* of this system—for instance, in the *meanings* of the items—the issue then becomes whether there is a way to define this property—the meaning—of *all* complex items solely in terms of the meanings of the syntactically-given parts of the complex item and their syntactic method of combination. A ‘yes’ answer to this for every complex member signals that the semantic theory assigns meanings compositionally. One might note that in this kind of compositional theory, there are *two* structures: the first, antecedently-given one (in the language case, the syntactic structure), and a second structure (in the language case, it would be the structure that “the meanings” manifest) which is a kind of mirror of the first structure. So another way of asking whether a semantic theory is compositional in this sense is to ask whether there is a homomorphic mapping from the syntactic structure to the semantic structure.

In the language case, this mapping is called “the meaning function”, which most authors writing about semantic compositionality symbolize by μ . So, $X = \mu(A)$ means that (i) A is some element of the first structure [i.e., A is some syntactic item], and (ii) that X is the item in the meaning structure that is paired up with A (i.e., X is the meaning of A). Then this second conception of compositionality asserts that there is a function f such that whenever A is composed (in the sense of Ontological Compositionality) of B, C, D ... by means of syntactic method R, then $\mu(A)$ is f applied to $\langle \mu(B), \mu(C), \mu(D), \dots \mu(R) \rangle$. That is: the system is compositional if and only if, there is a function f such that for every (syntactically) complex item A in the syntactic system, its meaning $\mu(A)$ is a function of, and only of, the meanings of A’s syntactic parts, together with the way they are combined. If NP_1 and VP_1 make up the entirety of S_1 and they are combined by rule-R to do so, then $\mu(S_1) = f(\mu(NP_1), \mu(VP_1), \mu(R))$. Underlying this type of compositionality is the slogan:

Definition 2 (*Functional Compositionality*) The μ of a whole is a function of the μ ’s of its parts and the ways those parts are combined.

⁵I am speaking generically here about connectionists and generically about their views of most properties.

If we restrict our attention to cases where μ yields meanings, then this abstract notion of Functional Compositionality becomes “semantic compositionality”.

A difference between the two notions of compositionality concerns whether some “whole” can contain things not in the parts. According to the ontological building-block view, no; but according to the functional version, yes. For, the first notion allows the whole to contain only what is in the parts, possibly re-arranged in some manner. But the second allows the thing associated with a whole (in the linguistic case: the meaning of a complex whole) to be a *function* of the things associated with the parts (in the linguistic case: a function of the meanings of the syntactic parts and syntactic mode of combination). And then there is nothing to stop such a function from introducing new material into the thing associated with the whole—that is, the function can make the meaning of a whole contain many new and radically different things from what are contained in the meanings of the parts. According to functional compositionality, all that is required is that this be a function—which merely means that it must introduce this same material and generate the same (type of) result every time it is faced with the same (type of) parts and manner of combination.

A simplified example to illustrate the difference between Ontological and Functional compositionality is this. Many theories have thought that adjective-noun combinations were to be analyzed semantically as a conjunction.⁶ Thus, *pink ribbon* is assigned a meaning like PINK(X) & RIBBON(X). Note, however, that we have introduced something into the meaning that is not in the initial phrase nor in the meanings of PINK or RIBBON: the meaning contains &, a conjunctive operator, which is not to be found in the syntactic representation of *pink ribbon*. So, the letter of the definition of Ontological Compositionality would not allow this to happen. On the other side, there is no problem for Functional Compositionality to generate such a representation by saying that the meaning of the syntactic operation ADJPHRASE \Rightarrow ADJECTIVE NOUN is to insert a & (plus add the variables).

However, finding some arbitrary function to do this job is easy—too easy, according to many theorists. For, it seems possible, without any further constraints, to gerrymander the function so as to arbitrarily (but in accord with the function’s definition) admit or omit anything at all as the value of any object in the function’s domain, because a function need not be “regular” or “systematic” or “uniform” or “orderly” in any intuitive sense. For example, Johnson (2014) remarks that while the following three French sentences in (1) have various parts in common, it is still possible to compositionally assign them the respective meanings in (2).

- (1) a. Le chien aboie.
The dog barks.
b. Le chat aboie.
The cat barks.
c. Le chat pue.
The cat stinks.

⁶Although current semantic theory does not hold to this account. I’m using it here merely as an example to illustrate a difference in the two conceptions of compositionality.

- (2) a. The dog barks.
 b. The cat dances.
 c. The skunk eats.

Johnson says that we can see that sharing a verb does not say anything about similarity of meanings, sharing a subject NP says nothing about a similarity of assigned meanings. “Nevertheless,” he remarks, “there exists a function that takes the syntax, and the meanings of the morphemes...and maps it to these meanings....In fact, any random, unsystematic assignment of meanings to sentences is compatible with the functional conception of compositionality....This is ‘dependence’ only in the weakest sense of that word.” In response to this general sort of worry, (Szabó 2012, p. 71) argues for the inclusion of a requirement that the function be one that makes the meanings of the parts *determine* the meaning of the syntactic whole. (Pelletier 2012, pp.160–1) says that the function has to correspond to the way that the meaning of the complexes *really depend* on the meanings of their parts... where ‘really depend’ is taken in some sort of ontologically-explanatory way. Pagin (2009, 2012) suggests that the function has to be a *computable function*, or perhaps an “elementarily computable” function—or even a “humanly-computable-by-anyone” function.

The position that opposes Functional Compositionality is—at least when the realm of meaning is being discussed—(*semantic*) *holism*.⁷ A semantic holist takes it that the meanings of some complex expressions cannot be accounted for by *any* function that “really depends” on the meanings of their parts. We discuss this notion later, in Sect. 6.4.

Let me emphasize that Ontological and Functional Compositionality—as well as their respective contraries, Wholism and Holism—can be applied to many different realms of inquiry. I focussed on the language-meaning case with respect to Functional Compositionality, but that was just for convenience, and because there is a large literature about this application. We return to the topic of compositionality and its various manifestations after introducing a bevy of *dramatis personae*.

3 (Some) Philosophers of Language

It would be very difficult to find any modern philosophers of language who hold that meaning is “in the mind”, or even just the weaker claim that the meaning of a lexical item is a mental concept.⁸ Locke (1690) had said

⁷Note: without the ‘w’. It’s a strange word: it ought to mean a belief in the ontological reality of holes, or perhaps a belief in the virtues of holiness. But it doesn’t. The OED attributes the first occurrence to Smuts (1926).

⁸I *did* say that I’m speaking in generic sentences! There are works such as Hinzen and Sheehan (2013) in which an individual’s concepts are viewed as having been generated pre-linguistically and then having some linguistic label attached to them. But then this label becomes subject to categories/conditions/features of language—both of “universal grammar” and the grammar of the

Besides [the ability to make] articulate sounds, therefore, it was further necessary that [Mankind] should be able to use these sounds as signs of internal conceptions; and to make them stand as marks for the ideas within his own mind, whereby they might be made known to others, and the thoughts of men's minds be conveyed from one to another. *Essay Concerning Human Understanding* iii.1.2

... a word is made arbitrarily the mark of such an idea. The use, then, of words, is to be sensible marks of ideas; and the ideas they stand for are their proper and immediate signification. *Essay* iii.2.1

This is nowadays cited in introductory philosophy of language textbooks as “the most naïve view of meaning”, and the remainder of the text is taken up with a brief rebuttal and then an exploration of alternative theories. Even the few avowedly mentalistic theories (particularly that of Grice 1989) eschew the notion of a mental *concept* in their theory, and employ instead the notion of *intention*, as in “intention to communicate something”, where this “something” is normally something “in or about the world”—such as the proposition that it is raining outside.

Philosophers of language and formal semanticists more generally hold the view that the meaning of a lexical item⁹ is some extra-mental item, which they (unfortunately?) call a concept. Mind you, different ones of these philosophers and semanticists hold different views concerning what such an extra-mental concept is—some hold they are “atomic” and cannot be analyzed. Others think that lexical decomposition is an observable phenomenon, so the corresponding concept must be decompos-

(Footnote 8 continued)

particular system becoming internalized. Likewise there are the various works of Fodor (1975, 1983, 1998, 2000, 2001; among many others) which have advocated a “language of thought” whose atomic members are somehow associated with natural-language lexical items. But the atomicity of the LOT items entails that the linguistic items also must be atomic and not subject to any definition or decomposition. Many theorists have found this implausible. Probably they will find the more recent Fodor and Pylyshyn (2015) even more implausible, where there is *no* intensionality at all in a concept, only extensionality considerations. (That is to say, only membership conditions for instantiating the concept.).

⁹Actually, we mean an individual *sense* of a lexical item. So we really wish to discuss separately, and not as if one word, such cases as: *Pitch: the degree of acuteness or gravity of a tone or sound* at the same time we discuss *Pitch: the central part of a cricket field* or *Pitch: the act or manner of pitching a baseball to a batsman* or *Pitch: a dark-coloured viscous substance* or *Pitch: a stroke with a lofted golf club* and so on and on. Or if that seems obviously to be a group of separate words—and hence obviously separate concepts—to you, consider the case of the noun *Lamb* whose senses include *a young sheep* and *the flesh of a young sheep used as food*. If those still seem as obviously being separate meanings, each to be a separated sense, consider *Collection: the activity of gathering things together* versus *Collection: A group of things gathered together*. The former is a sense indicating an activity; the latter is a sense indicating the result of the activity. Are these separate senses, even though closely related, or not? Some say yes, others say no. If they are separate, then they need to be treated as separate “lexical items”. But all this depends on how finely we would like to individuate senses. That having been said, we will not delve into that topic in this paper. For interesting thoughts on the matter, see Chapt. 8 of Atkins and Rundell (2008). (Both Atkins and Rundell are dictionary editors with many years of experience in issues concerning the individuation of senses.) We also wish only to countenance the “base form” of such a lexical item—a lexeme, as they are sometimes called.

able. Still, the main view in philosophy of language is that the meaning of a lexical item is a concept, in this extra-mental understanding of ‘concept’.

For (Frege 1891, 1892a, b), a concept is a particular type of function: one which, when given a proper name (or as we would say now, a singular term) as an argument, yields a truth-value as a result. In more modern predicate logic terminology, it is a monadic predicate (or rather, that entity in “reality” which corresponds to the predicate—a property, perhaps). Note that this is explicitly *not* a psychological item—Frege is quite resolute in his opposition to “psychologism”.¹⁰

Philosophers of language have antecedently believed that language “talks about the world” in that what one says is to be evaluated for truth or falsity¹¹ against what is happening in the world, not in anyone’s mind. When a person says *X*, s/he is not judged in accordance to how much that person believes that his/her concepts justify the assertion of *X*, but rather we just say the person has spoken falsely, if what they say does not actually match the world. Even if one has a stereotype of SNAKE that includes SLIMY, they cannot truthfully assert *Snakes are slimy* simply on that account. Even if *everybody* had a mistaken belief about some such feature of objects (maybe everyone’s mental idea of SNAKE includes SLIMY), it would not make the claim true. Present tense assertions describe “the way the world *is*”; future tense assertions describe “the way the world *will be*”. There is nothing here about the way the speaker believes it to be, believes it will become, or “conceives” it to be. For, if it were, then we’d say s/he was speaking truly even when saying something false—so long as s/he just believes it. (And while there are doubtless philosophical theories that promulgate such a viewpoint, they don’t include anyone in the philosophy of language.¹²)

This sort of consideration is motivated in part by some famous—and still accepted—doctrines brought out in works by Hillary Putnam (1975) on natural kind terms, by Saul Kripke (1980) on proper names, and by Tyler Burge (1979) on various terms that seem to rely on societal features rather than on an individual’s mental concepts. The moral of these works, and the lesson that has been adopted in philosophy of language, is nicely encapsulated by a phrase made famous by Putnam: “Meaning Ain’t in the Head”. Philosophers of language find it difficult to believe that there are respected (and reflective) theorists who seem to think otherwise.¹³

We can trace the difference between the “objectively oriented” philosophers and the “subjectively oriented” psychologists (and cognitive linguists) that are to be discussed in the next section by means of the *What did s/he say?* test.

Test 1 (What did s/he say?) *When a person utters a sentence, X, ask “What did s/he really say by uttering X?”*

¹⁰For a summary of Frege’s (and Husserl’s) arguments against “psychologism”, see for example Pelletier et al. (2008). These were directed against various attempts to think of mathematics and logic as “psychologistic”, but one might also see them as bolstering his position on thinking that meanings and concepts can’t be mental entities.

¹¹Or along other dimensions corresponding to speech acts other than assertion.

¹²Still talking generically!.

¹³I continue to speak generically!.

If, for example, the speaker utters “All snakes lay eggs”, then when the philosopher that we’ve been discussing applies the test, it would be claimed that the speaker *really said* that every (species of) snake procreates by laying eggs. This is to be evaluated by appeal to worldly facts, and if there are (species of) snakes that give birth to live young, then what the speaker said was false. If you take this view of *what is said*, you will always claim that these cases—where the speaker is sincere, but mistaken—show that meaning is *in the world* and that meaning is not a matter of what the speaker believed nor of what his/her concepts might be contained or was organized in accord with. But since that is what we say about cases when the speaker utters a falsehood, it seems equally to follow that when the speaker talks correctly, so too is his/her meaning a matter of what is in the world.

There are those on the other side, those who think that what the speaker *really said* was a fact about his/her concept of a snake, then the utterance “All snakes lay eggs” might not be false, depending on what features his/her concept SNAKE happened to have. It seems relevant here to note that in ordinary social interactions we distinguish between speaking falsely and lying. The former seems to track what the Objectivist has in mind by the **what did s/he say?** test, and therefore is about “reality”. The latter seems to involve the speaker’s mental state, and depend on whether the speaker is truly reporting what his/her concepts contain (with the intent to deceive). A more stark way of putting this point is that Objectivists take the **what did s/he say?** test to require an answer *about the world* whereas Subjectivists must take it to require an answer *about the speaker’s beliefs*. Furthermore, and although intuitions here can be a bit murky, it seems that if someone intentionally mis-reports his/her conceptual beliefs (say, s/he believes that snakes all lay eggs, but is trying to tell a falsehood) in saying “Some snakes give live birth”, then this is nonetheless a lie even though it is a truth.¹⁴ This intuition would be in accord with what the Subjectivist seems to hold about the **what did s/he say?** test, although I think one should instead identify it with **what did s/he intend?**

4 (Some) Cognitive Psychologists

There have been a number of nice reviews of the recent history of the notion of a concept within cognitive psychology, telling mostly the same story (Komatsu 1992; Laurence and Margolis 1999; Murphy 2002; Machery 2009; Barsalou 2010; from a more recent viewpoint see Hampton and Jönsson 2012; and from a recent philosophical perspective see Leben 2015). The narrative generally goes like this:

¹⁴About 70% of snake species are oviparous (they lay eggs). Of the other 30%, about two-thirds are ovoviviparous (eggs are hatched while within the body then expelled). The remainder are viviparous—they give live birth by nourishing the developing young through a placenta and yolk sac, a very unusual feature for a reptile. While ovoviviparous snakes, like the rattlesnake, are often said to give live birth, this is not true in the same sense that the viviparous snakes do it—which is akin to how mammals give live birth.

In the Classical Version of concepts, a concept was some mental item that classified items into those that fell under the concept versus those that didn't. The general idea was that a concept was a definition, although it was not always clear whether the theorists held that the classifying properties had to indicate *necessary* properties of the item in "the world" (and a sufficient number of them) or whether they might be able to "get along" with accidental properties so long as the classification ability worked in all the cases in which it was called upon. Would it matter whether our concept of orangutan or of baking soda listed exactly the biological classification properties¹⁵ or chemical features (NaHCO_2)—or instead listed a sufficient number of properties that any object that was encountered could be correctly classified as an orangutan or not, or as baking soda or not? [Well, yes it does matter, philosophers say. But that didn't seem to affect any of the psychological theories of the time.]

But data involving typicality judgments and category membership judgments were seen as overthrowing the Classical theory. Even those theorists who wanted to maintain that a concept had a "core" part and a "periphery" (sometimes called a "dictionary versus encyclopedic" distinction) found that there was no special place for such a distinction: it took subjects longer to judge that a chicken was a bird than that a robin was a bird. Presumably this difference is due to "non-core" aspects, such as size and flying ability, since they both are 100% birds, and everyone knows this. Thus there came to be a thought of replacing this with some sort of Prototype representation, one that attempted to accommodate the typicality and membership judgments in terms of "distance from the mental prototype". Here a prototype is something like a "most typical instance" or an "average" of all the instances encountered. But this in turn was refined into (or replaced by) a "family-resemblance" version of the theory. (The reference here to Wittgenstein was intentional, even if these theories were not, on the whole, all that well-grounded in Wittgenstein 1953.)

Another way to accommodate the evidence from typicality judgments was to think of a prototype as an *exemplar*: some particular example of the concept that one happens to fasten on. In addition, there will be a number of learned "rules" about how the exemplar could be modified and still represent a member of the relevant category. When this seemed not to adequately match the full range of typicality judgment data, a more careful version of Prototype theory was elaborated as having feature lists each one with associated probabilities. The concept DOG could have a feature COLOUR with many possible instances listed, each associated with a probability or likelihood. (Note now that there is no special distinction here between "core" and "periphery" of these features.) But computational complexity issues concerning the probabilities of *combinations* of features made these theories be untenable for some,

¹⁵The suggestion in the text of necessary or essential features defining scientific categories is perhaps most appropriate for non-biological concepts, or biological concepts as used by laypeople. But it is not the current viewpoint of professional biologists. Instead the current view is that *taxa*—species, genera, phyla—should be conceived genealogically, in terms of common ancestry with selected paradigms. Thus, for a professional biologist, "being an X" is a relational property of an individual organism, independent of both the superficial properties used for initial identification (amphisbaenias are not snakes despite being legless reptiles) and of explanatory properties (tuna are ray-finned fish despite being warm-blooded). Thanks to Allen Hazen for discussion on this topic.

and gave way to certain variants. If you think that every combination of values among the different feature lists and their probabilities could result in different outcomes for some further feature and its probability, and that this also depends on the probability of an interaction effect between groups of the features, it is pretty easy to see that this is not a tractable problem to characterize and nor to compute a result, if one tries to compute these in a straightforward manner. One way to extend this model could be to add further structure to the concepts, such as having many exemplars and a statement of the particular values each has on each feature dimension. One might then propose that a novel item is evaluated by sampling the stored exemplars and determining how close the newly encountered item is to the value generated from the sampling. But this seems to be no longer a pure prototype theory.

It seemed to many theorists that all these prototype theories have the shortcoming that there was no structure to the feature lists. Schemata theories were therefore devised that divide the various properties envisaged by prototype theories into separate “dimensions” that list the sorts of values they can embody, and can be related to other dimensions or values in those other dimensions. And still more recently “theory based” versions of this same idea have come to the fore with the underlying idea that concepts are not just learned with an already-existing background of other concepts, and therefore have to “fit in” with them, but their internal structure will have to embody previous knowledge about what sort of features are likely to be important in a new concept.

And yet even more recently we’ve seen “neo-empiricist” theories of concepts rejecting the assumption that there has to be *any* relationship among the features of a conceptual structure. Instead we get temporary ad hoc groupings from sensation and motor perception that fit whatever situation one is in. You are in a situation/context where it is important to recall the colour pattern of poisonous snakes, then your ad hoc concept of SNAKE will contain that sort of colour information. But if instead it is a situation where it is important to recall the type of constrictor snakes and their size, then your ad hoc concept of SNAKE will contain that information. The idea is that there is no continuous concept, no such thing as a concept of SNAKE in general. (See Barsalou 2008, 2010, 2017, and in a related framework, Prinz 2002.)

The general idea is that the meaning of some general term (or singular term, for that matter) is the mental concept it occasions. As we have just seen, some of the recent theories hold to the view that such concepts aren’t merely “retrieved” but are actively constructed on the spot as the linguistic and environmental input is received. The construction depending on features of the environment, broadly conceived. We will look at this “construction of meaning” in Sects. 5.2, 6.2 and 6.3 below. In the literature that discusses this—both in the psychological literature and in the cognitive linguistics literature (which is discussed in Sect. 5.2 below)—it is remarkable that focus is given to very simple lexical items, and sometimes to the “conceptual combination” apparent in two-word phrases. But language is “special” (in ways outlined below in Sect. 6.4). In particular it can create arbitrarily long noun phrases that, according to the present theory, should also designate concepts. But one wonders whether the mechanism of conceptual-combination-in-context is up to the job. Here is what Carlson (2010) says about two examples:

[N]ot only do we talk about flying geese and meowing cats and leaf-bearing trees, but we also talk about

- *Unpainted kitchen appliances that are just beginning to rust* need to be replaced within a period of two to three months by a qualified kitchen professional in order to prevent any possibility of bacterial contamination.
- *Friendly but slightly confused medical professionals without appropriate training who nonetheless have medical degrees from top-ranked teaching hospitals* have much to contribute to society beyond their incomplete medical expertise.

The principled point is that linguistic expressions can be unbounded in complexity... with the consequence that if these are mapped directly onto certain types of mental states that we are calling concepts, then we stand in need of a device ... that can produce the arbitrarily large number of corresponding brain states in order to provide these phrases with appropriate denotations.

Depending on how this “device” is defined—and if such a device is really possible to define—it might describe Ontological Compositionality or it might obey Functional Compositionality. Or it might obey neither. But until such a “device” is described, it is not possible to evaluate claims concerning whether the mental combination being presumed by these modern theories of concepts is or isn’t compositional.

5 (Some) Linguistic Theories

Scholz et al. (2016) list three “grand traditions” in modern (and not-so-modern) linguistics, which they label *Essentialism*, *Emergentism*, and *Externalism*. These grand traditions are characterized by their different attitudes towards seven different aspects of what a linguistic theory should be about or achieve. The seven aspects are:

- What are the primary data for theory construction in linguistics?
- What is the primary subject of a linguistic theory?
- What is the aim of a successful linguistic theory?
- What is linguistic structure?
- Besides accuracy, what is to be valued in a linguistic theory?
- What is the relation between child language and adult language?
- What is learned when a first language gets acquired?

According to (Scholz et al. 2016)’s categorization, these grand trends reflect a tendency within different groups of researchers to answer these seven questions in one of three ways. And so the members of a trend can be characterized by commitments to a number of central tendencies or foci. But as Scholz et al. are also at pains to remark, a wide variety of differing research programs could be carried out within each of the three trends... although as a sociological fact, members of each of the three trends seem all to focus on more or less the same directions as one another.

The three trends can be briefly explained by saying that (a) Externalists are most interested in “language as it occurs”, and the most visible research effort of this trend here is corpus analysis, (b) Emergentists are most interested in how language is used by people and how this is an outgrowth of other cognitive aspects of people, and probably the most visible research effort of this trend is Cognitive Linguistics/Semantics (but see footnote 17 below), and (c) Essentialists are most concerned to presume an internalized language faculty that can be characterized by an abstract, computationally-oriented model. The two most visible subgroups of this grand trend are the theories put forth by Chomsky over the decades and the formal semanticists. The three groups’ answers to the seven aspects of what a linguistic theory should be like are given in (Scholz et al. 2016)’s Table 1.

5.1 (Some) Externalists

As Scholz et al. (2016) characterize these three trends or tendencies of research goals, the Externalists are most typically exemplified by some computational corpus linguists, whose goals are to find underlying patterns that are discernible in actual usage when the sample is large enough to show it. The patterns might then be characterizable by some rules that could be used to give an overall theory of the language. Another version of Externalism holds to certain philosophical accounts of meaning in terms of use: for example the “inferentialism” of Brandom (1994, 2000) or the “language game” and “form of life” of Wittgenstein (1953). Many forms of linguistic behaviourism could also be characterized in this way, for example Skinner (1957) and in a different way Osgood et al. (1957). The computational approaches to linguistic theorizing in the Externalist vein tend to eschew the postulation of mental items, such as concepts, intentions, purposes, and the like, and instead take meaning to be the likelihood of co-occurrence with nearby words. As can be seen from the ‘AIM’ row of Table 1, it is not a central concern of Externalism to worry about generating semantic representations that obey semantic compositionality, nor with psychological data concerning speaker/hearer mental life during conversation. And hence they stand outside the sort of discussion this paper is concerned with, namely the interaction between linguistic features and mental items. So, I will ignore the work in the Externalist direction in the remainder of this essay, except for mentioning how some inferentialists view semantics.¹⁶

¹⁶Important references for current versions of Externalism are (Manning and Schütze 1999; Jurafsky and Martin 2008) and (Schubert 2015, esp. Sect. 9: “Statistical Natural Language Processing”). In this section, Schubert discusses how this is a new direction in computational linguistics that “is significant from a philosophical perspective, and not just a practical one.” Schubert discusses a wide range of areas where this ‘philosophical shift’ has been applied.

Table 1 Three approaches to the study of natural language. Printed with the kind permission of the Editor of *The Stanford Encyclopedia of Philosophy* and the authors of Scholz et al. (2016)

	Externalists	Emergentists	Essentialists
Primary phenomena:	Actual utterances as produced by language users	Facts of social cognition, interaction, and communication	Intuitions of grammaticality and literal meaning
Primary subject matter:	Language use; structural properties of expressions and languages	Linguistic communication, cognition, variation, and change	Abstract universal principles that explain the properties of specific languages
Aim:	To describe attested expression structure and interrelations, and predicting properties of unattested expressions	To explain structural properties of languages in terms of general cognitive mechanisms and communicative functions	To articulate universal principles and provide explanations for deep and cross-linguistically constant linguistic properties
Linguistic structure is:	A system of patterns, inferable from generally accessible, objective features of the use of language	A system of constructions that range from fixed idiomatic phrases to highly abstract productive types	A system of abstract conditions that may not be evident from the experience of typical language users
Value especially:	Accurate modeling of linguistic form that accords with empirical data and permits prediction concerning unconsidered cases	Cognitive, cultural, historical, and evolutionary explanations of phenomena found in linguistic communication systems	Highly abstract, covering-law explanations for properties of language as inferred from linguistic intuitions
Young children's language is:	A nascent form of language, very different from adult linguistic competence	A series of stages in an ontogenetic process of developing adult communicative competence	Very similar to adult linguistic competence though obscured by cognitive, articulatory, and lexical limits
What is acquired:	A grasp of the distributional properties of the constituents of expressions of a language	A mainly conventional and culturally transmitted system for linguistic communication	An internalized generative device that characterizes an infinite set of expressions

5.2 (Some) Emergentists

It is the linguists from the Emergentist group—so-called because they see language as developing out of and emerging from antecedently existing cognitive, social, or biological forerunners—who are the most self-conscious employers in linguistics of

“mental terminology” such as (mental) *concept*. One group—the self-named *cognitive linguists*—is explicit, indeed insistent, about their orientation, that psychological evidence is part and parcel of their own enterprise¹⁷:

...[W]hat is the precise meaning of *cognitive* in *Cognitive Linguistics*, and how does this meaning differ from the way in which other forms of linguistics conceive of themselves as being a cognitive discipline? ... Cognitive Linguistics is the study of language in its cognitive function, where *cognitive* refers to the crucial role of intermediate informational structures in our encounters with the world. Cognitive Linguistics is cognitive in the same way that cognitive psychology is: by assuming that our interaction with the world is mediated through informational structures in the mind. It is more specific than cognitive psychology, however, by focusing on natural language as a means for organizing, processing, and conveying that information. (Gaeraerts and Cuyckens 2007, pp. 4–5)

Sometimes, though, the use of different names by various authors within the framework for the mental entities that they assume to be the mental reality that language is describing or reflecting, can make a casual reader think that they may not be talking of the same thing that cognitive psychologists are. For instance, their use of *schema, constructions, spaces, categories, frames, prototypes, figure/ground*, among others, in explaining the inner workings of their theory may hide this. But about them, under whatever name, the thought is that the meaning or semantic value of lexical items (especially) are these mental items. Some programmatic but representative remarks about the standpoint of Cognitive Linguistics, as made by some of its practitioners, are:

The experience that users have with language shapes cognitive representations, which are built up through the application of general principles of human cognition to linguistic input. The structure that appears to underlie language use reflects the operation of these principles as they shape how individual speakers and hearers represent form and meaning and adapt these forms and meanings as they speak. (Bybee and McClelland 2005, p. 382)

Representation is perhaps *the* most important, and most contested, foundational concept in modern cognitive science. Cognitive Linguistics takes the view that linguistic structure is motivated by conceptual representation and communicative function, thereby placing the representational function of language at the center of its concerns. (Sinha 2007, p. 1280)

[The psychologist’s notion of] the basic level of categorization... is linguistically basic because basic-level terms tend to be morphologically simpler, to be acquired earlier by children, to be used as the unmarked choice for introducing referents into discourse, and to provide the raw material for extensions of the lexicon by means of metaphor, metonymy, and word formation. (Schmid 2007, p. 124)

Several general cognitive abilities can be brought under the broad title of “construal operations”. At issue is the fact that linguistic expressions do not, and cannot, designate a state of

¹⁷There are other schools or orientations within the broadly Emergentist category, such as Functionalism (Bybee 1998; Dirven and Fried 1987), Construction Grammar (Croft 2001; Goldberg 1995; Fried and Östman 2004; Östman and Fried 2004; Kay 1995), various sorts of Sociolinguistics (Halliday 1978; Hurford 2000; Labov 1966, 1975; von Savigny 1988; Tomasello 1998; Van Valin 2003), and Language Evolution (Burling 2005; Cheney and Seyfarth 2005; Greenhill et al. 2010; Li 2002; MacWhinney 2005; O’Grady 2008; Power 1998; Wildgen 2008; Worden 1998).

affairs as it “objectively” is; rather, the scene must be processed and conceptualized by the human mind. Construal operations include attentional processes, force-dynamic construals, metaphor, image schemas, and conceptual blending. (Taylor 2007, p. 574)

The semantic views held by Emergentists can perhaps be gleaned from the foregoing quotations, but in short, and at least when restricted to a lexical item, the view about meaning is:

Definition 3 (*Linguistic Emergentism*) The meaning of a lexical item is a (mental) concept.

Of course, Emergentists will have interesting and complex things to say about the nature of these mental concepts, and their mutability in context—where by ‘context’ is meant not only the longer sentence or discourse in which the lexical item occurs but also the situation when and where it was uttered, and also by the sum of background experiences the speaker has plus his/her beliefs concerning the background experiences of his/her audience. It is one of the fundamental views of Emergentism, or at least the Cognitive Linguistics subtype of Emergentism, that linguistic expressions are widely variable in their meanings, depending on the sorts of factors just mentioned. But then since they are also committed to the view that meanings are mental items, it follows that these concepts must also be able to vary along the same dimensions.

Evans (2009) states a general feature (‘literalism’) he perceives in “traditional linguistic theories”,¹⁸ and which he is concerned to overthrow. It can be seen that Evans’ view of compositionality, which he sees as inherent in these traditional theories, amounts to what I called Ontological Compositionality with its “monotonic”, “additive” view of composition, and not the Functional Compositionality that I had attributed to these other types of linguistic trends. In this I think Evans is wrong, and I think that the notion embraced by formal semanticists (who I shall introduce shortly) is Functional Compositionality. I also see this as one of the “disconnects” between cognitive linguistics/psychology and philosophy of language/formal semantics. A further feature of Cognitive Linguistics, emphasized here by Evans, is that there is no such thing as “the meaning of a word” (or other linguistic item): they are always infected with contextual aspects when they are used, and furthermore when they were first learned. A meaning, on this view, is something fleeting—it comes and then goes as the context (linguistic or situational) changes.¹⁹

[T]he view of word meaning adopted under literalism is that word meanings are assumed to be relatively fixed and stable. Put another way, the semantic primitives which make up a given word meaning can be identified independently of context. Once identified, word meanings are integrated, by applying the rules of the grammar, in order to provide sentence meaning. Literalism, then, assumes that the contribution of language to meaning construction is essentially additive in nature, positing grammatical principles which ensure that the

¹⁸By which he probably intends what I’ve called Essentialism.

¹⁹Actually, Evans tries to separate “senses stored in semantic memory” from those “meanings constructed ‘online’”. But the idea is that the former serve only to guide the construction of the latter, which is the meaning of the word-in-context.

semantic units which result are unable to change or delete the meanings of the units which are conjoined to form a larger semantic unit or expression. This restriction serves to make a larger expression, for instance a sentence, monotonic with respect to its component parts, where the term “monotonic” has to do with the view that the component parts retain their original meanings in the larger expression. Thus, the individual word meanings do not alter their meaning in the larger semantic units of which they form part. (Evans 2009, p. 6)

Later in the book, Evans returns to this thought:

[W]ords, I suggest, do not mean independently of context. Thus, the fundamental problem with literalism is that it attempts to artificially divorce (word) meaning from (situated meaning in) context of use. More precisely, literalism lives in something of a fool’s paradise. It holds that language users retain an idealized, timeless meaning for *open* which they neatly keep apart from the situated meanings of *open* which arise from its use in examples . . . The mistake that literalism makes, then, is in being reductionist and simplistic about meaning.

[W]ords are never meaningful independent of the utterance in which they are embedded, and the encyclopaedic knowledge and extra-linguistic context which guide how words embedded in an utterance should be interpreted. Indeed, evidence from the perspectives of social psychology, cognitive psychology, interactional sociolinguistics, cognitive linguistics, corpus linguistics, and computational linguistics reveals that the view that words constitute fixed, context-independent structures, and that meaning construction is appropriately modelled in terms of the straightforward approach to compositionality sketched above is untenable. (Evans 2009, pp. 21–22)

This viewpoint is not idiosyncratic to Evans. Rather, as he says, part of this general viewpoint of Cognitive Linguistics is that mental concepts, and hence the linguistic items that receive their meaning from these concepts, are determined not by some “core” dictionary meaning but rather are related to much more about the world of one’s acquaintance. Using other terminology, a concept is said to embody *encyclopedic knowledge* and to be augmented or restricted by situational features of the utterance. Many of the early Cognitive Linguists have also promulgated this view, but Langacker (1987; 2008) is perhaps the most theoretically-oriented of the early practitioners, and he was probably the most influential developer of “encyclopedic semantics”. The idea is that lexical items (and longer linguistic phrases) directly access concepts, and since these concepts are developed through interaction with the world in non-uniform ways, these linguistic items can’t be understood independently of the encyclopedic information held in these concepts. Semantic structure, then, is just *identical* to conceptual structure.

Evans is concerned throughout his works to view meaning not only as mental concepts, *à la* Langacker, but to show that these concepts are continually changing, as the following quotation from a later work illustrates:

Language works extremely well because it is part of a larger meaning-making complex. It’s dependent upon a suite of other capacities – in particular, a repository of thoughts and ideas, that we carry with us, in our minds – upon which language draws each time we open our mouths to speak. Language, together with this mental apparatus, co-conspires to produce meaning, enabling effective communication. (Evans 2015, p. 6)

(One notes that, even were all this true about how the co-conspiracy works, it is quite unclear how that will help with communication!)

This group of the Emergentists—the Cognitive Linguists—is one of the two groups of linguists that I wish to keep in focus during the remainder of this paper. We will return to a discussion of this group, with an eye to seeing whether and how it is similar to (some) work on concepts in cognitive psychology of the sort I described above. But first we continue our account of linguistic theoretical frameworks.

5.3 (Some) Essentialists

A large number of contemporary Essentialists who follow Chomsky's teaching on this matter claim that semantics and pragmatics are not a part of the study of language.

[T]he study of meaning and reference and of the use of language should be excluded from the field of linguistics ... [G]iven a linguistic theory, the concepts of grammar are constructed (so it seems) on the basis of primitive notions that are not semantic (where the grammar contains the phonology and syntax), but that the linguistic theory itself must be chosen so as to provide the best possible explanation of semantic phenomena, as well as others. (Chomsky 1977, p. 139)

It seems that other cognitive systems – in particular, our system of beliefs concerning things in the world and their behavior – play an essential part in our judgments of meaning and reference, in an extremely intricate manner, and it is not at all clear that much will remain if we try to separate the purely linguistic components of what in informal usage or even in technical discussion we call 'the meaning of linguistic expression.' (Chomsky 1977, p. 142)

In Chomsky's view, "it is possible that natural language has only syntax and pragmatics" (Chomsky 1995, p. 26); that is, only "internalist computations and performance systems that access them"; semantic theories are merely "part of an interface level" or "a form of syntax" (Chomsky 1992, p. 223). In his view, then, there is no distinction between semantics and pragmatics, and *neither* forms a part of linguistic theory.

So, in the Chomskian take on the matter, semantic topics such as the meaning of lexical items are a matter for theories other than linguistics to determine. It is a matter of the cognitive systems that describe our beliefs about things in the world which would determine what meaning (and reference) is. Without using the term *concept* in any technical sense, it would seem that Chomsky endorses the claim made in Definition (3) above, about Emergentism, and denies only that this should have anything to do with linguistics. (Of course, when an Emergentist hears this they will be astonished to hear the 'only' in such a sentence, since they believe that this is the fundamental starting point for linguistic theory!) Nonetheless, we can see that this form of Evans' "traditional linguistic theories" does *not* embody his "fool's paradise of literalism". Meanings and such are not at all a part of linguistic theory.

However, not every Essentialist agrees with Chomsky on this point. Many believe that every theory *should* incorporate a linguistic component—semantics—that yields meanings, in much the same way that most philosophers of language believe there to

be such a separate component of a theory of language. Often, although not always, this component amounts to a truth-theoretic account of the values of syntactically-characterized sentences. This typically involves a translation of the natural language sentence into some representation that is “intermediate” between natural language and a truth-theory—perhaps an augmented version of first-order logic, or perhaps a higher-order intensional language. And other times it involves a further level of syntax, LF (sometimes this is inaccurately called “logical form”), which itself is then interpreted by some truth- or information-theoretic method. The Essentialists who study semantics in such ways usually agree with Chomsky in seeing little role for pragmatics within linguistic theory. But their separation of semantics from pragmatics allows them to accord semantics a legitimacy within linguistics itself, and not “just” in psychology or sociology. This group are known by the general name of Formal Semanticists. And perhaps this is Evans’ group of “traditional linguists” who adopt “literalism”?

So for these Formal Semanticists, such features as the individual differences between speakers on their beliefs about “the world” will become a part of pragmatics, and not a part of semantics. And hence, *not* a part of linguistics. (But rather, of some social psychological theory, or perhaps a theory about an individual’s cognitive life.) This is a view they share with the philosophers of language who were surveyed earlier. Chierchia and McConnell-Ginet (2000) and Heim and Kratzer (1998) are leading textbooks in this general framework, and both adopt a “logic-oriented” approach to semantics that employs a possible world semantics and other formal tools. Their differences reside in their length (Chierchia and McConnell-Ginet is much longer than Heim and Kratzer) and to some extent in their differing underlying syntactic frameworks. (Chierchia and McConnell-Ginet 2000) put their case in the preface and beginning of the book as follows:

[W]e focus on what has come to be known as logical, truth-conditional, or model-theoretic semantics. This general approach to meaning was developed originally within the tradition of logic and the philosophy of language and over the last twenty years or so has been applied systematically to the study of meaning in natural languages, due especially to the work of Richard Montague. (Preface)

Whatever linguistic meaning is like, there must be some sort of compositional account of the interpretation of complex expressions as composed or constructed from the interpretations of their parts and thus ultimately from the interpretations of the (finitely many) simple expressions contained in them and of the syntactic structures in which they occur. . . . An important test of a semantic theory is set by compositionality. Can the theory generate the required interpretations for complex expressions from a specification of interpretations for the basic items? As we will see, explicit specification of how word meanings are combined to produce sentential meanings is not a trivial task. (pp. 6–7)

Somewhat later in the book, this general train of thought is revisited, arguing further that formal semantics is committed to Objectivism and is opposed to Subjectivism:

Meaning manifests itself in the systematic link between linguistic forms and things that we speak of or talk about. . . . Were languages not to provide for significance in this sense, the question of meaning would hardly arise.

Language enables us to talk about the world, to convey information to one another about ourselves and our surroundings in a reliable fashion. What properties of language and its uses underlie this remarkable fact? What allows language to serve as a guide to the world and to enable us to learn from what others have perceived (seen, heard, felt, smelled) without having to duplicate their perceptual experience ourselves? Informational significance does not require that language links to the world in ways that are predetermined by the physical structure of our environment. Nor does it require that environmental information is simply registered or received without active input from perceiving and thinking human minds. Yet it does probably require a regular and systematic correspondence between language and the shared environment, what is publicly accessible to many different human minds. If you are skeptical about informational significance, consider the use of language in giving directions, warnings, recipes, planning joint activities, describing events. Things occasionally misfire, but by and large such uses of language are remarkably effective. Language could not work at all in such ways were it not imbued with some kind of informational significance, being about matters in a public world. (pp. 11–12)

It is also appropriate here to mention the area of “Truth-Functional Pragmatics” (Recanati 2012). This is also an approach that I would like to bring under the wing of Formal Semantics, even though the practitioners think of themselves as allowing “pragmatic” effects to “intrude” on the semantic understanding of utterances. François Recanati is possibly the most influential member of this group, and although this approach seems to endorse non-compositionality (in the Functional sense), Recanati himself has given some ways to have it be “almost compositional” (in Recanati 2012; see also Pelletier 2013 for further ways). On the dimension of Cognitive-Linguistics-to-Formal-Semantics, truth-functional pragmatics is quite close to Formal Semantics, and I will include them into the Formal Semantics group, even though it also shares a number of features with Cognitive Linguistics (e.g., the importance of context in determining the *content* of what is said, and a penchant for claiming that this entails that the “true meaning” of what is said is full of “pragmatic enrichment” and “modulation”). For, the fact is that in these theories the final evaluation of an utterance depends on truth: that is, upon whether the meaning of what is said is to be evaluated against what is happening “in the world”. And this puts it pretty squarely in the camp of the Formal Semanticists.²⁰

This Formal Semantics part of the Essentialist camp is the second group of linguists that I wish to engage about the notion of compositionality, comparing their view(s) with those from philosophers of language, cognitive linguists (Emergentists), and cognitive psychologists.

²⁰Of course, there are differences within the group and thus between (some of) them and (some of) the earlier groups we distinguished. For the purposes of this paper, though, I will focus on their mutual endorsement of Atomism and through that commitment, their advocacy of semantic compositionality and its relatives.

6 (Some) Objectivists Versus (Some) Subjectivists

There is a similarity, as I have noted, between the philosophers of language and the formal semanticists. They both think of “meaning” as being *outside* of speakers, as *public*, as being *in common*. Despite whatever other differences there may be between and amongst them, I therefore will call this merged group *The Objectivists*.²¹

There is a different similarity that I noted between the cognitive psychologists and the cognitive linguists. They both think of “meaning” as something that is *inside* speakers, *private*, as being *personal*. Despite whatever other differences there may be between and amongst them, I therefore will call this merged group *The Subjectivists*.

Carlson (2010) ends his discussion of the difference in usage of the term ‘concept’ in the formal semantics literature and the experimental psychology literature with

Concepts, as least as studied by psychologists, have not found a natural place in the study of natural language semantics in the formal semantics tradition as practiced in the past thirty or so years. This does not point to any shortcomings in semantic theory, I emphasize, for there are excellent intellectual as well as practical reasons as to why this should be. However, it does leave us with something of a disconnect between what is going on in the mind – and I take concepts to be major features of it – and the way we talk about things. (p. 33)

I remarked earlier that philosophers of language—who also employ the word ‘concept’—think of it as being “in the world” and not “in the mind”. Another way of putting all this is to say that Objectivists think of a concept as something that can literally be apprehended or grasped by more than one person (and not as some sort of “abstraction from, or averaging of, the private possessions of many different people”). On the other hand, Subjectivists are committed to a view of concepts that are “private”, in the sense that it is literally true that a concept is in the mind of exactly one person. (Correspondingly, the sense in which different speakers have the same concept is one of “determining similarities between and among the different private concepts, either by abstraction or averaging”.)

This difference is one of the root causes for the disputes that dot the literature. It gives rise to one of the standard arguments used by Objectivists (mostly the philosophers of language) against the Subjectivists (Sect. 6.1). It also gives rise to one of the standard arguments used by Subjectivists (mostly the cognitive linguistics) against the Objectivists (Sect. 6.2). In Sect. 6.3 I point to another place where (some) Subjectivists and Objectivists battle: general terms. I argue in favour of the Objectivists, although I also point to a possible two-tier theory imagined in Sect. 7.3. In Sect. 6.4 I discuss the so-called “Master Argument(s)” in favour of semantic compositionality, bring out a feature that seems to have gone unnoticed by its proponents—that, as an argument from the Objectivist side, it undermines itself by its own presuppositions. I also point out that these Master Arguments do not really support semantic compositionality, but rather they are arguments for some sort of semantic Atomism. But semantic Atomism is compatible with some Subjectivist frameworks.

²¹Yes, there are various other groups that are also called ‘Objectivists’, but I’ll not be talking about any of them in this essay.

We start with a standard argument used by Objectivists against Subjectivists, to the effect that the Subjectivist notion of a concept implies an untenable position concerning “privacy”. If this argument is right it puts the Subjectivists in a very dire position indeed—they would be committed to the incoherence of all of what is believed about human interaction. Let’s see how that goes:

6.1 Concepts and Privacy

Concepts, on the Subjectivist account, are private—there is no literal sharing, no joint ownership, etc. There is thus no direct way to determine whether a speaker and his/her audience share anything whatsoever when it comes to a concept—even in the case where the concepts are occasioned by the same verbal signal (word). And even though one might think of behavioural correlates such as classifying the same objects and seeing how many they disagree with, this doesn’t give any reasonable guarantee of sameness of concept, as opposed to agreement in behaviour.

But perhaps the Subjectivist would say that agreement in (certain types of) behaviour is sufficient for our purposes here; that we don’t really need shared concepts, only shared behaviour of the relevant sort. To this an Objectivist counters by asking whether we really would like to claim literal identity of different people’s sensation of pain, whenever they agree to rate the same number out of 10 on a Numeric Rating Scale. The Objectivist thinks that no one, not even the most committed Subjectivist would agree to that. Indeed, once the Classical theory of concepts is thrown out, it is a part of the Subjectivist story that the concepts are developed by accommodating to the experiences that an individual has. Surely this guarantees that the concept that one person has will be both numerically and qualitatively different from the next person’s. And if we are supposing, as the current theorists do, that the meaning of a lexical item is the concept it occasions, then of course no two people will have the same meaning for any word. Subjectivists have not talked much about whether there are also different ways for concepts to combine in different people on account of their differing experiences, seemingly treating any combination as if it were a conjunction. But even if they were able to give us an account of how simple concepts can combine in whatever ways they do, and how the resulting complex concepts can themselves combine further, and yet further, until a “complete thought” (corresponding to a full declarative sentence) is formed—as I say, even if that is all granted—then since these combining rules or abilities will themselves be due to experience that is learned-by-an-individual, they will thus almost certainly differ from one person to another. In short: this aspect of the Subjectivist theory guarantees that there will be no understanding between any two people, at least if understanding amounts to having (qualitatively) the same concept.

It is tempting to say, and many Subjectivists have,²² that communication does not require “absolute identity” of the corresponding mental items but only that they be “similar enough”. But this can’t really be made out in a non-circular way, it seems to me. The hypothesis that two different minds are “similar enough” with respect to their understanding of some term—say, ‘democracy’—has no empirical content other than simply some antecedent belief that the two people *are* understanding what each other says when they talk, despite their differences in acquisition of the relevant concepts. But that was precisely the (alleged) fact that the “similarity” was hypothesized to explain. The very nature of the Subjectivist notion of concept makes it impossible to give an *independent* criterion of “similar enough”. Contrast this with the Objectivist claim that there literally is something—something *in reality*—that the participants both grasp. It may be correct to object that this is asserted without any independent evidence, but at least it is not a restatement of the issue, as I claimed (in the previous footnote) that the Subjectivist is committed to.

6.2 Concepts and Meaning

Although Objectivists seem to believe they have knock-down reasons for holding that “meaning just ain’t in the head”, as we outlined in Sect. 3, it seems equally obvious to Subjectivists that without taking into account the mental life of speakers (and possibly hearers, too) it is impossible to give an account of what meaning is. After all, isn’t it obvious that a speaker has some belief that s/he wishes to get across, or warning that s/he is trying to convey, or emotion that s/he is expressing, or And are these not all mental states? How could there be any meaning in an instance of language use that did not give any role whatsoever to these mental items? Isn’t it obvious that when a speaker says that *The object in the corner is a chair* that they

²²Here are two, chosen from some 200 years apart:

Men do not understand one another ... by mutually occasioning one another to produce exactly and completely the same concept; they do it by touching in one another the same link in the chain of their sensory ideas and internal conceptualizations, by striking the same note on their mental instrument, whereupon matching but not identical concepts are engendered in each. (von Humboldt 2000, p. 152)

Since the cumulative experience that makes up the conceptual system will differ between individuals, often considerably; so, the ability of two individuals to communicate effectively depends on there being a minimum degree of overlap between their ongoing representations and, therefore, the concepts from which these representations derive. The question then becomes: How much overlap is enough. Critically, understanding one another in communication does not require identical representations between speaker and listener. ... just representations that overlap enough to achieve the current communicative goal. (Connell and Lynott 2014, p. 400)

I think that anyone can see that neither account gives anything more than a restatement of the actual issue, not a solution.

are identifying some mental representation (a percept?) of some object taken to be in the corner, and are classifying it as falling under the concept CHAIR? Given all this, it certainly seems obvious that there is some relationship or other between concepts and meaning. Objectivism denies this, and is deficient for this reason, the Subjectivist says.

The intuitive and natural manner to characterize meaning in a Subjectivist framework is to say that the basic way for mental concepts to operate is by *representing the world*—at least the world as seen from the point of view of the holder or owner of those concepts. A linguistic item then will designate or mean or be brought forth by (or will bring forth) the relevant mental concept of the language user, and will thereby *represent* the relevant aspect of the world. This is a type of “two-level” or “dual aspect” or “two-tier” semantics, and we will canvass that sort of response in the section below (Sect. 7). For now we just remark that the sort of two-tier theory being alluded to also has analogies with theories mentioned in McNally and Boleda (2017) and Winter (2017), and perhaps shows a remarkable convergence of opinion among theorists with widely different starting points.

One formal approach to subjectivist theories of meaning is exemplified by so-called proof-theoretic semantics. Such theories have started by thinking of the meanings of the logical connectives as being given by the ways they can be introduced into a discourse and what sorts of “language moves” they can justify. For some technical reasons this led to adopting intuitionistic logic as the underlying framework for a theory of the use of such connectives (see especially Dummett 1991). Fuller versions of this idea try to generalize the range of such introduction and exploitation moves that can be given a formal exposition. Most of these theories have their formal roots in the intuitionistic type-theory of Martin-Löf (1980) and Prawitz (2006); see for example its development in natural language structures by Francez and Dyckhoff (2010); Francez et al. (2010); Francez and Ben-Avi (2015); Francez (2014, 2015). In this approach, meaning is taken to be a certain type of proof, namely a set of canonical derivability conditions—stated in natural deduction format from which the derivability conditions are based. The idea is that the proof system reflects the “use” of the sentences in the linguistic fragment under consideration and thus allows recovery of their entailment and assertability conditions. This formal conception of semantics is opposed to the more usual formal version in linguistics and philosophy of language of describing the meanings as truth-conditions in arbitrary models, which are taken to be ways the world actually *is*, or might be (the latter for intensional language). In proof theoretic theories there is no notion of “truth” or of “modelling the world”. There are only rules for the appropriate *use* of language. In more informal versions of proof theoretic semantics (Brandom 1994, 2000) this general strategy is called *inferentialism* and is seen as a “use-based theory of meaning”, which employs notions such as “language entry and exit rules”, and “language evaluation rules”.²³ (A nice description of both the formal and informal aspects of the topic is in

²³“Use theories of meaning” trace their ancestry to ordinary language philosophy, especially the sort that is usually attributed to the later Wittgenstein (but now moved to a higher level of formality). In the Wittgensteinian version—and generally carried into the modern versions—the issue of

Schroeder-Heister (2016). We shall not here evaluate the various arguments from the proof-theoretic and model-theoretic sides that allege to find shortcomings in the other approach. A lot of this back-and-forth mirrors the Objectivist–Subjectivist argumentation we have been canvassing, except done on a highly formal and mathematical level. Instead I will simply appeal to proof-theoretic semantics below as a possible formal account of the ways in which people can *use* the representations afforded by concepts to produce an “internal meaning” in our discussion of two-tier semantic theories of Sect. 7.3.)

6.3 *Specificity and Meaning in Particular Instances of Language Use*

Another feature of the dynamic between Objectivists and Subjectivists that holds back some sort of rapprochement is their differing underlying notions concerning just what meaning should *be* in particular cases. We’ve already seen the large-scale differences in the form of context-dependent versus independent, encyclopedic versus dictionary, and the like. But even in a specific case of meaning-in-a-very-definite-context there are important differences between the two camps. Here I outline the Objectivist side against the Subjectivist. The differences come in two forms.

6.3.1 *Generality*

Objectivists and Subjectivists differ on the “level of specificity” that should be a meaning in the various cases. It is not so clear to me why this should follow from their overall positions, but their writings leave little room for doubt that it is so. Objectivists have a notion that some meanings include others as special cases or as particular ways that the meaning can be manifested. For example, they might say that ‘The cover of this book is red’ can be made true in many different ways: it could be crimson, it could be cardinal red, it could be scarlet, it could be cerise, etc., etc. And, they hold, a person could truly say and mean that the book cover was red even all the while *knowing* (or believing) that it has some specific shade, say raspberry. Their Objectivist view is that *what was said* allows for instantiation in many different ways; the speaker’s particular knowledge of the shade is irrelevant to *what was said* (as is any potential hearer’s knowledge of the situation which might dictate the particular shade). Similarly, should some speaker claim ‘There is a red fruit in my refrigerator’, it wouldn’t matter to *what was said* that it was, or was known by the speaker to be, an apple with a red-delicious coloured exterior, a watermelon with a ruby-coloured interior, or a burgundy-coloured grape, etc., etc. The idea, again, is that language

(Footnote 23 continued)

Objective versus Subjective was said to rest on a common mistake, that of not according primacy to how (ordinary) people (ordinarily) use language.

contains *generality*, allowing us to talk without having to always and completely describe each and every detail of what is being asserted.

Presumably this attitude follows from the Objectivist's viewpoint that *what is said* (think of the **what is said?** Test (1) of Sect. 3) determines a certain group of possible states of the world—those possible states that would render the statement true (for assertions; *mutatis mutandis* for other speech acts). So long as the actual world is one of those possible states, the statement is true. Thus, for the Objectivists, the meaning is the set of possible ways that the sentence could be true. And with 'The cover of this book is red', any of hundreds of specific shades would be adequate.²⁴ This picture is nicely captured by Sainsbury (2001):

'Italian *F*' is satisfied by a satisfier of *F* which is of or pertaining to Italy. This is an unspecific but definite, unambiguous and complete meaning. An Italian book is one of or pertaining to Italy, and a book may pertain to Italy by being on my pile of books to take there, or by being about Italy, or by being manufactured in Italy, or by being written in the dominant language of Italy. . . . 'Feline care' is satisfied, in the same sense though not in the same way, by a vet tending a cat and by a cat tending the puppy to which she is acting as foster mother. Both vet and cat supply care of, or pertaining to, felines . . .

Sainsbury's main point in this is the contention that theorists should distinguish between different 'readings' or meanings of a sentence from different ways in which it could be made true. Just because (for example) *This is an Italian book* can be made true by this book's being about Italy, or written in Italian, or the book bought while visiting Italy, and so forth, it does not at all follow that this shows that there are somehow different 'meanings' involved with the sentence. And if these are all the same 'meaning', then a compositional semantics needs only to generate a meaning for *Italian book* that is a function of the meanings of *Italian* and *book*. It is not the job of semantic theory to generate all the different 'readings' of a sentence.

Some Subjectivists (especially in Cognitive Linguistics or where they follow the general (Wittgensteinian) doctrine that all experience is "of the particular" and that a desire for generality is "the root cause of all philosophical conundrums") seem to refuse to acknowledge that there is any reasonable notion of generality in meaning—whether it be in language or in concepts. The fact that someone maybe has in mind that they are attracted to the colour red in their footwear makes it be that when the person says "I'm going to go buy a pair of shoes this afternoon" s/he *means* that they are going to buy a pair of *red* shoes. The fact that some particular physical phenomenon can be described in differing ways is claimed to give rise to different "construals" (this is a technical term).²⁵ Such different construals are seen by these

²⁴A caveat should be registered here, concerning homography and homophony, where different words are spelled or pronounced the same. Different words can give rise to different meanings and hence divergent ways they could be made true. But the same word-sense just gives rise to one meaning with differing ways to be made true. See also footnote 9.

²⁵This technical usage seems to stem from Langacker (1990), who discusses (p. 61) a person looking at the night sky, describing what s/he sees as a *constellation*, a *cluster of stars*, as *specks of light in the sky*, etc. He says "such expressions are semantically distinct; they reflect the speaker's alternate construals of the scene. . . .".

Cognitive Linguists as being different meanings for one and the same piece of language; the apparent reason for this being a belief that the semantic interpretation of some linguistic item is *directly* a particular mental structure without any intervening level of representation that might itself describe or be true of many differing values, and that these mental representations are *specific*. (Perhaps not totally specific, one presumes, but certainly more specific than Objectivists think their non-mental concepts are.) Objectivists naturally find this viewpoint mysterious and hard to take seriously, thinking that this is just a case of one and the same piece of language having different ways that it could be made true—that is, as being a case of generality. But here are a few quotations to give the flavour of the view, according at least to the types of Subjectivists we are canvassing.

Expressions do not mean; they are prompts for us to construct meanings by working with processes we already know. In no sense is the meaning of [an] . . . utterance ‘right there in the words.’ When we understand an utterance, we in no sense are understanding ‘just what the words say’; the words themselves say nothing independent of the richly detailed knowledge and powerful cognitive processes we bring to bear. (Turner 1991, p. 206)

Thus, there is no “general meaning” for a word, only the specific meaning as it is used on some particular occasion.

In [some] cases, varying content appears to originate in background situations. When representing *sofa*, for example, people might include information about sofas occurring in living rooms. Similarly, when representing *truth*, people might include information about truth being relevant in courts of law. In general, considerable amounts of information in conceptual representations appear to describe the background situations in which concepts occur, including settings, related objects, other agents’, actions, and variety of internal states (e.g., goals, evaluations, affect, motivation, reward, mentalizing). . . . The concept of *chair*, for example, is likely to be associated in memory with background situations for kitchens, living rooms, classrooms, theaters, jets, and ski lifts. . . . Sitting in a jet chair, for example, is associated with a different setting, actions, and internal states than sitting in a living room chair. . . . As a consequence, when a concept is processed in a particular situation, the situation activates related background information that enters into task performance. In other words, the extrinsic properties active for a concept also vary across situations, not just its intrinsic properties. (Barsalou 2017)

And so each particular situation uses the word in the sense appropriate exactly to such a situation, and not in its general sense.

Objectivists complain that the Subjectivists are so enthralled with the *particularity* of an experience that they have blinded themselves to the fact that language is full of generality in the description and reporting of that experience and refuse to acknowledge that this is precisely what happens on particular occasions. Indeed, they say, that’s the whole point of general terms. You use X to say something, then it is what X is true of that counts. . . not some hidden (or even public) idiosyncratic interpretation of X that you might have. If I say that I met someone on the train yesterday, an Objectivist would claim that it is just silly to think that I “really” meant that I met a woman, just because I know that it is true that the person is a woman! Objectivists think the Subjectivist theory is completely bewitched by the view that what is being communicated is just *exactly* what one is aware of and what one’s par-

ticular background knowledge picks out in some specific case. That's just plain false, they say: it depends *exactly* on *what was said*.

The **what did s/he say?** test shows that it can't depend on what the speaker happens to believe, even when what the speaker believes *is* true: suppose a speaker says *Connie saw a cat yesterday, but Delia didn't*. And let us suppose that the speaker knows it was a black cat that Connie saw, and that she saw it in regular daylight, with her normal eyes, while paying attention. So the speaker knows that Connie is aware that she was seeing a black cat.

Notice now that what I just said in those circumstances—that Connie saw a cat yesterday but Delia didn't see one—certainly wouldn't be true if Connie saw a black cat and Delia saw a tabby! But such a circumstance *would* make that claim be true if the sentence were to mean what Subjectivists claim: that Delia didn't see the same black cat that Connie saw (or same colour of cat that Connie saw), as seems to follow from the Subjectivist position under discussion. For, Delia can surely see a tabby while she is not seeing any black cat! This shows, using the **what did s/he say?** test that the speaker did *not* mean that Connie saw a black cat—despite the fact that the speaker knows it was a black cat and despite the possible fact that the speaker's mental image of the incident included a black cat, and that s/he would assent to the claim that Connie's seen cat was black, and so on. That simply is *not what was said* and hence not what was meant on that occasion.

The way to get the correct intended meaning is to make the *a cat* not be any more specific than “being a cat”. Add any more and the sentence could be made true if Delia didn't see this extra material, even though she saw a cat. The view that every meaning is as specific as the speaker happened to have in mind would make lying or disseminating impossible. Diplomacy would be impossible. Romantic relations would be impossible. To Objectivists, it's as if Subjectivists about language have a morbid fear of things—even mental things, even linguistic things—that are not tied down to specifics.²⁶ One wonders what they might think about the concept SPECKLED HEN and the number of speckles it might have (Tye 2009).

On a more conciliatory note, one might wonder whether the following is the root cause of this dispute on generality. The truth-conditions of the *language* are general, even when the relevant concepts are particular in their nature. But which should be called the meaning? Objectivists point to the world, using arguments of the sort just canvassed. Subjectivists point to the mind. Given the radical differences in what is being “pointed to” by the two theories, they can't both be right about the same thing. So—is one right and the other wrong? or are they just talking about different things, calling them both ‘meaning’? Or—as we explore in Sect. 7.3—is it that meaning has two different aspects, while still being one item?

²⁶I'm speaking generically, of course.

6.3.2 Grounding Concepts

We continue now with the general issue of subjectivity, and the Subjectivist's apparent commitment to meaning—and to thought more generally—as being “private” and within the mind of a single person. This is obviously a position that none of our Subjectivists²⁷ wants to hold in such a blatant form, for the reasons that I have given above. The standard response by our Subjectivists is to invoke some notion of “groundedness”.²⁸

There are many such notions in their literature,²⁹ sometimes going by other names, such as ‘situated’, ‘embodied’, and ‘enactive’—with subtle differences in application. (‘Grounded’ is used both as a general cover term and also more specifically to cover a notion of “gained by experience”; ‘situated’ is often used to talk about how a concept is activated in a particular situation; ‘embodied’ seems to be used to avert to the fact that we are physical beings and thus live in the physical world; ‘enactive’ is used for a more neuropsychological causal pathway. I’ll continue with ‘grounded’ as a general term for all these.³⁰) The idea behind them all is to give the otherwise private meanings (and other mental items) a way to “touch reality”. (Perhaps a way to avoid the “Brain in a Vat” scenario to the effect that since we never have any interaction with “reality”, we can’t know or talk about it (Putnam 1981).) Employing one of these notions is pretty popular in all the forms of Subjectivism about meaning we have been looking at, but while at least some varieties might seem to Objectivists to be plausible, the high degree of “fuzziness” in the articulation of how it in detail works—as well as the imaginative terminology that has been invented to describe it—has not met with much in the line of conversion of Objectivists. To many such un-convinced Objectivists, it seems no more informative to be told that such meaning concepts are grounded, situated, embodied, or enacted, than it is to be told that inter-subjective validity is guaranteed by “our shared form of life”.

The general reception by Objectivists of the notion of grounding in this form seems to be: Subjectivists can talk and talk about how one’s mental concepts and other meaning-related mental structures are ‘grounded’ in experience, ‘situated’ in context, ‘embodied’ in one’s perceptual/motor systems, or ‘enacted’. But so long as this process—whatever it is—allows for differing results on the mental items for different people or the same person at different times, there has been no explanation of how communication of mental structures to one another can succeed.

²⁷Recall that I am speaking generically!.

²⁸A general overview of the area can be gleaned from (Shapiro 2014) and (Pecher and Zwaan 2005), or at a higher level of generality in (Wilson and Foglia 2016).

²⁹A negative view of the many different notions called “groundedness” is offered in (Wilson 2002). A positive view that nonetheless details many different meanings of the notion is in (Rohrer 2007, pp. 28–31).

³⁰See (Barsalou 2010, p. 619): “‘Grounded cognition’ reflects the assumption that cognition is typically grounded in multiple ways, including simulations, situated action, and, on occasion, bodily states.”.

6.4 *Compositionality Cuts both Ways*

Although the privacy topic is, I believe, the root cause for the differences between Objectivists and Subjectivists, the literature from either side does not seem to dwell on this topic. Instead, we find other disputes that are tangentially related to this, but this root cause is not often mentioned. For example, Objectivists have long pointed to a set of “arguments” that are supposed to establish compositionality³¹:

Argument from understanding We can *understand* an infinite number of novel sentences, so long as they employ words we already understand. We understand sentences and other combinations that we have never encountered. So, language must be ‘compositional’: it must start with a finite stock of words/morphemes with their meanings, and put these together in a finite number of different ways, but using an unlimited recursive method to arrive at the infinite number of understood sentences.

Argument from productivity/creativity We can *create* new sentences that we have never heard or used before, and we know that they are appropriate to the situation in which we use them. This can only happen if language is ‘compositionally’ organized, so that we learn some finite base of words and rules, but know how to combine them recursively so as to produce totally new descriptions with the intended meanings.

Argument from learnability We are finite creatures who are exposed to a finite amount of information concerning our language. Nonetheless we can *learn* a system that is capable of infinite expression of meanings. The only way this can happen is if what we learn has a finite basis of terms and rules, but the rules themselves allow for arbitrarily complex ‘composition’.

As I see it, there are three threads interwoven in the arguments:

1. that language is something special (infinite, or novel, or creative, or whatever) in its ability to express meanings,
2. that people manage to use/learn/understand language despite their being “finite”,
3. that one (the only known?) way to do this is if language exhibits a semantically compositional framework.

It is important to keep these arguments in mind when reviewing the Objectivist position. Note in particular something we will return to in Sect. 7: Thread (2) of this general position cites certain abilities or outcomes of mental operations in language speakers. So whatever else is in the Objectivist position about meaning being “in the world”, at least in their only direct arguments for compositionality, there is a presumed subjective aspect.³² And although many Subjectivists try to establish that the

³¹Well, it is admitted by all that these are “arguments to the best explanation” and do not *logically* entail the conclusion of semantic compositionality.

³²In the linguistics literature there is also, to be sure, some argumentation concerning the proper form of a linguistic theory; and in that milieu, semantic compositionality is praised for its clarity and transparency. But these “methodological” (aesthetic?) considerations don’t really count as direct arguments in favor of the position, in my mind.

information present in the world impinging on language learners/speakers will allow for the Thread (2) to occur without compositionality (even though Thread (1) is correct), it seems to me that they would be on more secure ground in arguments against Objectivism were they to focus more attentively on the Objectivist's "missing link" (that these arguments can't succeed in securing (3) if the underlying presupposition of Objectivism is correct).

I have just argued that the Objectivists have painted themselves into a corner by employing as their only justification for (Functional) semantic compositionality an argument that presupposes the falsity of Objectivism. But what is the alternative, if one wanted to somehow include the Subjectivist presupposition? Two thoughts come to mind, and both can be found in the Subjectivist writings (mostly in the Cognitive Linguistics works). However, they are often stated in such a way that it is not clear which of the two thoughts is being promoted. And to my mind, this is because the two can be seen as shading into one another along a particular dimension, which I will try to elucidate. As I see it, one of the two understandings is an impossible position for anyone to hold, but the other one is not at all impossible. But given that the two ends of this dimension shade into one another, a lot of the argumentation has been against the impossible end and it has been allowed by some to impugn the other end by a sort of guilty association. Even granted that it is difficult to stop the "slippery slope" from the one end to the other, I will try to emphasize an underlying difference so as to take a step in rehabilitating the reputation of the group of possible theories that inhabit the one end.

One end of the dimension along which these theories vary is given by such theorists as Schank (e.g., 1972) and Anna Wierzbicka & Cliff Goddard (e.g., Goddard and Wierzbicka 1994; Wierzbicka 1996; Goddard 2002). In these theories there is a finite set of mental primitives in terms of which all other concepts can be defined. As I explained in Sect. 2, theories that can define all larger wholes in terms of their parts are Atomistic in nature, and so this version of Subjectivism is Atomistic. As I also explained in that section, theories that are Atomistic can often be compositional: at this end of the dimension under discussion here, we see that it is maybe possible to see the meaning of any defined concept to be a function of the meanings of the parts of that concept, and so on, down to the basic primitive universals. But one might relax the idea of a small set of primitives and allow more and more primitive concepts (or primitive word-concepts). In fact, one might allow *every* word to be such a primitive, thereby allowing *every* word-concept to be used in defining the others. Depending on how this is carried out, it may be either Wholistic or Atomistic, and the resulting semantic theory might therefore be either holistic or atomistic. In the former case it cannot be compositional, as I pointed out in Sect. 2, since holistic theories embody a type of Wholism, and by definition a wholistic theory is not compositional. But in the latter case, it might be... depending on the details of the theory.

One version of this more inclusive picture is offered by Langacker (1987), and I interpret it as Atomistic even though it makes all word-concepts available as possible members of the meaning of any other word-concept. (Langacker is being cited here with approval by Evans 2009, p. 61):

What Langacker appears to have in mind is that the semantic material – informally the meaning – associated with a lexical form, i.e., a word, relates directly to the contents of conceptual structure. In principle, this conceptual structure relates to a diverse and sophisticated body of non-linguistic knowledge The meaning of *uncle*, on this view, is potentially a function of the vast body of encyclopaedic knowledge we have of what it means to be someone's uncle. In addition to the specific relationship holding between the child of uncle's sibling, this also includes detailed knowledge relating to marital relations, familial relations, the social status of uncles, the types of behaviours associated with uncles, as well as individual knowledge any given individual may have with respect to uncles they have known. Yet while this knowledge is encyclopaedic, it is for Langacker part of semantic structure, i.e., directly encoded by a lexical form. Langacker's argument is that there is no principled way of separating putative linguistic from non-linguistic semantic representation.

This sort of theory of the meaning of word-concepts is in the “middle” of my dimension, and despite that, I would wish to view it as being Atomistic even though it allows for meanings to form “circular” definitional chains due to the large amount of material allowed in the meanings of a word. At the far end of the dimension, maximally different from the Shank or Wierzbicka/Goddard theories, and forming the Wholistic version of these types of theories, are so-called conceptual role theories.

Conceptual Role theories can be found in very many versions of Cognitive Linguistics. Like the Langacker version I just outlined, these theories also do not have a small number of primitive concepts, but rather all word-concepts are allowed to appear in the definition or explanation of any other concept. What sets conceptual role theories apart from the Atomistic Langackerian theories, and makes them be Wholistic, is their employment of *semantic liaisons*. A semantic liaison is an implicative connection that one concept might have to another concept—typical liaisons are ones of semantic inclusion (e.g., that being a tiger *implies* being a feline), contrary semantic entailment (e.g., that being a mammal is *incompatible* with being cold-blooded), compatibility relations (e.g., that one object can be both a brother to someone and a cousin of someone). More generally, in conceptual role theories *all* such implicative connections that a concept A has are inherited by any concept that includes concept A as part of its definition. Thus, if concept A has concept B as a part of its definition, then concept A also has all the implicative relations that B embodies. It is usually thought that this means that every concept has all other concepts and their implicative relations as part of its meaning. So, the meaning of a lexical item will be the concept it occasions *plus all the “liaisons” it has with other concepts*. Here a liaison is seen as either some meaning-entailment or (less strictly) some other association that the concept has with other concepts (that is, all the encyclopedic associations that a person might have with respect to the concept plus *their* liaisons). The meaning then is seen as the collection of all these liaisons. And of course, that is true for all concepts, including for the ones that form the liaisons with the first concept. In other words, meaning now becomes a feature of the entire set or body of concepts, and insofar as one can attribute meaning to any one concept, it is

most accurately said that this concept's meaning is its place in the entire network of liaisons among the concepts. Or, in the terminology mentioned above in Sect. 2, this is Semantic Holism—the competitor to Functional Semantic Compositionality. But it is not just a competitor to Semantic Compositionality, it is a competitor to any theory in the Atomistic camp... even the Langacker-style theories surveyed just above. In my mind it is important to keep these Atomistic vs. Wholistic versions of this style of theory clearly separated, when we wish to evaluate candidates for a Subjectivist theory of meaning.

Many theorists—and not just the Objectivists—find meaning holism wanting,³³ since according to this view, it would be very likely that different people will have had different language-learning experiences, and therefore according to semantic holism will have different liaisons among their lexical items. In turn, this means they do not speak with the same meanings associated with their words as one another. So, it follows that they do not mean the same thing as one another even when they utter the same word, phrase, or sentence. But then, when they (think they) agree with one another, or disagree with one another, they are in fact not doing so. Any evidence one has that you disagree with your conversational partner is, according to the theory, equally good evidence that you are talking about different things, and not disagreeing at all. And should you (appear to) *agree* with someone whom you don't like, that is better evidence that you are talking about different things, since you are independently motivated to disagree with him or her.

It is also often pointed out that, not only is communication with others an impossibility according to the theory's own features, but also one cannot communicate with oneself from one year to the next, one day to the next, one hour to the next, ... For, in the intervening time period the person would have new experiences and (except in very special cases where the person is comatose) these will impact the class of liaisons that impinge on the meanings of his lexicon. Although one *thinks* one remembers that the bark of eucalyptus trees peels off yearly, there is no justification for this, according to the theory. Over the year when you last had that thought, many of your liaisons have been altered—or at least, it is most likely to have happened—and thus your last-year-thoughts do not necessarily have *any* connection with this-year-thoughts.

Indeed, a Subjectivist-holist can't change his mind about *anything*! For, the very act of getting new information makes it be a different thing, and hence it is *not* a case of changing one's mind about concept X. The old concept X is no longer there to have a new opinion about. Not only is it impossible to disagree with another person, as I remarked above, but it is impossible to disagree with the past. We think that we have learned that Anaxagoras was wrong when he said that the unevenness of the moon's surface is due to the mixture of earthy matter with cold. But Subjectivist holism is committed to claiming that we are in fact not at all disagreeing with him. We can't disagree. We can't agree either. We've just changed the subject.

³³ (Fodor and Lepore 1991) is perhaps the best-known opponent. It seems that one author is a Subjectivist and the other is an Objectivist about meaning.

At this point it is tempting for our Subjectivist-holist to retreat a bit from the view. But it seems that the relevant retreat would be to a position akin to the Langacker-style Atomistic Subjectivism of very complex, but separate, concepts—which is *not* holistic. And so it might be a candidate for the Subjectivist portion of a two-tier theory imagined in Sect. 7.3.

7 (Some) Desiderata for a Two-Tiered Semantic Theory

A two-tiered semantic theory is one that alleges an accommodation of both the Objectivist and the Subjectivist theories, or at least of their main aspects or main contentions concerning meaning. Some versions of such theories might keep the two tiers separate, in the sense that neither one has any particular effect on the other, so that the features of the subjective (mental) never affects anything about the features of the objective, and correspondingly, the objective features never affect the subjective ones. We might, for example, follow a once-standard usage and call the subjective features that we have been discussing the *narrow meaning* or the *narrow content* of beliefs and meaning; and call the objective features the *wide meaning* or *wide content* of language. We could then form a theory of narrow content of mental concepts and its effects on, as it may be, beliefs, desires, and actions. And a different theory of wide content for linguistic items that gave an account of the truth and falsity of utterances.

But even if such a disjunctive two-part theory could be successful, each part in its own realm, it wouldn't be very satisfying to either the Objectivists or the Subjectivists, who each think that their theory accounts for some very important aspects of the unified activity or concept of meaning. What is desired is an "interactive" theory that shows how the two are related. And to do this requires some point of contact: a *nexus* of interaction. This nexus needn't be anything especially exotic; it need only be some item that both theories can agree upon, but which leads each theory in its own way. I describe one possible nexus in Sect. 7.3.

There is a worry here, though, when one tries to integrate these two theories. (Lewis 1970, p. 19) puts the worry thus:

I distinguish two topics: first, the description of possible languages or grammars as abstract semantic systems whereby symbols are associated with aspects of the world; and, second, the description of the psychological and sociological facts whereby a particular one of these abstract semantic systems is the one used by a person or population. Only confusion comes of mixing these two topics.

Partee (1979) has also discussed some worries. She describes a group of phenomena that she thinks cannot be handled by (what I call) Subjectivist theories. And in considering what a Subjective-theorist might say in reacting to this challenge, she considers three general attitudes they might take, but thinks that they all are incapable of handling these (propositional-attitude-related) phenomena in a satisfactory manner. In the end she despairs of finding any satisfactory solution. She calls the

three approaches (i) the idealized approach, (ii) a realistic-egocentric approach, and (iii) a realistic-nonegocentric approach; and it is useful to appreciate these options and the shortcomings that they give rise to. The option (i) seems to require a systematic understanding of what the “non-ideal” mechanisms are, and how an idealized theory can nonetheless account for them by means of some type of “contextual or performance factors”, in the same way that friction and the like are contextual factors affecting ideal laws of motion. She thinks this approach would not realistically be possible in the realm of meaning. Approach (ii), she thinks, requires figuring out one person’s (one’s own?) psychological performance factors, allowing one thereby to have a semantics for one’s own mental economy. But the remaining part of such a theory then requires attributing those factors to other speakers, and this seems like an implausible direction to go, given the difference in relevant experiences we all have. The third approach is, in her estimation, just “giving up” on the problem. It amounts to saying that, unlike option (ii) where we discover our own subjective theory and attribute it to others, we instead just say that *there are* some such factors and every person employs their own, whatever they may be. We merely “add variables over psychological mechanisms incorporating performance factors and add variables over psychological interpretations of individual lexical items”, and then existentially quantify over these variables. She thinks that this is just a retreat from trying to account for both Subjectivism and Objectivism in a single theory. It amounts to saying that since we don’t know what the relevant features and factors are, we just put existential quantifiers over them.³⁴

(Partee 1979) closes with this remark

What I have tried to suggest is that the [cognitive] linguist’s concern for psychological representation may be relevant to every semanticist’s concern for an account of propositional attitudes. So far I don’t see how to achieve either goal; my only positive suggestion is that a good theory might be expected to achieve both at once.

I do agree with Partee’s “positive suggestion”. But I also don’t wish to engender Lewis’s “confusion by mixing the two topics” of Objectivism and Subjectivism; and rather, I would want to show that they each *can* play a role in a joint semantic theory.

One aspect of Lewis’s admonition not to mix the two types of theories is the conviction that attempts to provide the two tiers both from *within* one side of the Objectivist/Subjectivist divide is not the right way to go—this is where the confusion comes from. I will try to demonstrate why I think this in Sects. 7.1 and 7.2. A part of the story will be that, coming from within just one side of the dispute, there is no common *nexus point* upon which the two can agree. It seems to me that the way to avoid Lewis’s worry is to show how there can in fact be a common nexus, and from this common point of agreement each of the theories goes its own way and does not necessarily affect the inner workings of the other. Below, in Sect. 7.3, I mention some considerations that may alleviate the worry that with only a nexus for connection, the two theories could differ in all sorts of unimaginable ways.

³⁴Partee also thinks of Lewis’s positive approach in (Lewis 1970) as being like this.

7.1 Objectivist Two-Tiered Theories

Some apparently two-tiered theories are actually theories *within* an Objectivist framework. Kaplan (1989), for instance, distinguishes ‘context of utterance’ and ‘circumstance of utterance’, as well as *character* and *content*. A way of putting his distinctions is that the *character* of an utterance is “the meaning without taking into account the context of utterance” or perhaps “the meaning in the language”. In any particular context of utterance, the values of the various pronouns, demonstratives, times, etc., that occur in the utterance are fixed, and then the evaluation of the character at that context will yield a *content*. Now using this content, we can evaluate it at a particular “circumstance” (maybe a particular possible world, such as the actual one) and the result will be a truth value. Kaplan’s motivation for the distinction between character and content involves the interpretation of pronouns and other indexicals (perhaps extended to predicates and the like), and not, generally speaking, the Subjectivist-oriented concerns about the structure of concepts and how they lead to beliefs, desires, and other features that guide behaviour.

So this picture plainly has both tiers working on the Objectivist side of meaning: once the values of the contextual variables are determined, the resulting semantic content is pretty much what one would expect from an Objectivist—a function on a possible world or “circumstance” to a truth condition. The semantic value of the character of a sentence is a function whose value at any context is the propositional content at that context. That is, a function that takes contexts as arguments and returns content, content being a truth conditional statement, namely a function from contexts and circumstances to some truth value. (See Westerståhl (2012) for how to make Kaplan’s theory be compositional.)

Another theory that comes from the Objectivist side is Frege’s, although perhaps the undeveloped Subjectivist side perhaps *could* be developed in some non-Objectivist way. Frege’s semantic theory of meaning (*Bedeutung*) and sense (*Sinn*) contains a subtheory which has certain explanatory powers relevant to a two-tiered account, even though it has its roots in an Objectivist framework.³⁵ In the objective, physical realm are the usual members of the physical universe of ordinary (and scientifically discoverable) objects. In the objective, abstract realm are not only such items as numbers, but also concepts and functions—taken to be the referents of predicates, relations, and operations on objects. Singular terms of language mean (that is, *bedeuten* or denote) the ordinary objects and the number-like items. Monadic predicates mean the concepts, relational expressions mean the relations, and function terms mean the functions—all members of the abstract realm.

A person can *grasp*—a technical term—these items in the abstract realm and thereby consider and possibly understand them (in the ordinary senses of these terms). But this is always done under some *mode of presentation* of the thing thus grasped. Frege insists on this way of putting how objects are grasped under a mode of presentation because of his realization that one needs somehow to distinguish cog-

³⁵Well, taking Frege’s view of the “abstract”, non-physical, non-mental, “third-realm” repository of concepts and thoughts as being a version of Objectivism.

nitive states of a person when the person has two names for the same object. The sentences in (3) obviously have different cognitive statuses, and yet the two names mean (denote, *bedeuten*) the same object:

- (3) a. Allen Stewart Konigsberg is Woody Allen
 b. Allen Stewart Konigsberg is Allen Stewart Konigsberg

The mode of presentation ALLEN STEWART KONIGSBERG designates a person and so does the mode WOODY ALLEN. But being told that the person designated via ALLEN STEWART KONIGSBERG is that person who is designated via ALLEN STEWART KONIGSBERG is cognitively trivial, whereas being told that it is the same person as the one designated via WOODY ALLEN is (or can be) informative.

Frege does not give any account of how the subjective world ought to process this “grasping”, nor what would happen in the subjective world that causes us to recognize when modes of presentation are identical, or how certain conceptual features would lead to assent or belief or action. Nonetheless, this picture has guided those Frege-friendly Objectivists who want some (at least minimal) psychological account of (at least) the informativity of identity statements. Frege’s type of formal move seems relevant to having a two-tiered semantic theory: the Objective theory would see no difference in the semantic values of (3-a) and (3-b), but the Subjective theory would make it important.

7.2 *Subjectivist Two-Tiered Theories*

Subjectivist theories that attempt to also account for the “objective reality” aspect of language all have the feature that they replace the Objectivist slogan of *language* representing reality with the slogan that (some aspect of) people’s *mental life*—for instance, their mental concepts—represents reality. It would seem to take some fancy footwork to justify the replacement of an unexplained notion of representation with another unexplained notion also to be called a representation.

One way that some Subjectivists have responded to the challenge that they need to have two distinct tiers was to take the viewpoint that they can use the distinction between people’s judgments of membership in a category and their judgments of the typicality of membership in that category. The membership judgment could perhaps fill the role of the objective-tier membership in that category. But from the Objectivist point of view using the judgements of membership as a substitute for “really being a member of the category” is the sort of idea that will give rise to Lewis’s “confusion”. At the best it can be *evidence* for membership in the category, but is clearly “contaminated” (the Objectivist claims) by all sorts of other biases. We can see that this can’t really be the Objectivist’s “fact about reality” by considering some of the actual experimentation that tests this. For example, the fact that subjects will judge a middle-height person as “both tall and not tall” in strong preference both to “tall” and to “not tall” (Alxatib and Pelletier 2011) does not in the least tend to support any theory about the reality of people manifesting contradictory properties.

Attempts to separate the two types of judgments (membership and typicality) and use one as a proxy for “reference” just miss the point of Objectivism. Although we can see this idea as endorsing a two-tier theory, we can also see that the use of “membership judgment” as if it were real membership is not something that Objectivists would agree to. And this is because judgments of whatever type just aren’t the sort of “neutral” nexus point that true two-tier theory requires.

Classic examples of this sort of “only apparently two-tier” theory from the Subjectivist side, provided by the group of theories that separate the notions of typicality from membership in a category, are (Smith et al. 1974; Osherson and Smith 1981; Smith and Medin 1981; Smith and Osherson 1984; Armstrong et al. 1983; Bourne 1982), among many others that are more recent (e.g., Winter 2017). In this arena, the main emphasis seems to be an argument as to whether there really are two different processes that people engage in. For example, Armstrong et al. (1983) demonstrated that typicality judgments are given to concepts whose membership was categorial by anyone’s standards, for example, odd versus even numbers. Such a result is then taken by this group as evidence for a two-tier theory.³⁶ But notice, as I said just above, the presumption that this can be what Objectivists mean by “membership in a category” is just wrong. An Objectivist will insist that it instead is a person’s *belief* or *judgment* that it is or isn’t a member. So, this is not the type of theory that I wish to call a true two-tier theory. A part of the objection is that such theories do not provide any *nexus* of contact between the two realms that is agreed-upon by both. Here we see the Subjectivists take the *concept* of category-membership as the nexus; but in the Objectivist’s eyes, this just raises the issue of how such a mental conception be attached to the external world. In a way, this is the reverse of the Subjectivist’s critique of the Frege “mode of presentation” nexus as a way to solve the issue of co-referring proper names. A Subjectivist would not (or: not necessarily) agree that this is where the two realms should meet.

Earlier in this subsection I remarked on the “fancy footwork” that might be required to make replacement of one unexplained notion of representation by another seem like an advance. Chomsky’s “Minimalist Program” has sometimes been thought to exemplify this sort of footwork. I will display this by a few quotes from McGilvray (1998), who is here explicating the notion of meaning in the Minimalist Program. We will not pause to evaluate its accuracy as an account of Minimalism, but instead will use it as a further example of a two-tier theory of meaning that has both aspects on the Subjectivist side. The first point in McGilvray’s exposition is to argue that the “external world” cannot be relevant to meaning, on the grounds that it is not “near” where the linguistic computation is taking place.

[B]ecause [linguistic] computations must also be local – they must proceed over ‘neighbouring’ elements, or be local – practical necessity restricts it to items in the head. Extending the domain of a theory of a cognitive competence outside the head makes it responsible for far too much; a theory would have to cover so much that it would lose any serious chance of being constructible. This poses a problem for an intentionalist view of a cognitive capacity, one that insists on determining internal states by their relations to things outside the head.

³⁶The negative side of the argument has seemed to many to be settled by Hampton (1988b). See also Hampton (2017).

Intentional properties arise when an inner state or process is treated as a representation of something in the world, or as having referential or alethic properties.³⁷ (McGilvray 1998, p. 231)

Further to the thought that the external world can't participate in meaning, McGilvray later says

Although [linguistic structure] does not deal with the kinds of semantic issues that [Objectivists] have in mind (reference and truth), it does deal with meanings. Generally, I suggest that the domain of syntax includes all locally determined, intrinsic features of linguistic mental events. If so, syntactic sciences include any disciplines that describe and explain the elements of this domain. As we have seen, the locality condition excludes reference and truth. (McGilvray 1998, p. 243)

And finally, McGilvray tries to show how this “perspective” allows for an internalist viewpoint on linguistic meaning which nonetheless has certain “ties” with the external world by means of when language is employed “in the world.”

Chomsky's view of [the semantic value of a linguistic item] as “a perspective” was an adaptation for a general audience of the technical syntactically defined concept of [such a value] as “an interface”. He points to a way in which [these values] are both rich and anthropocentric, and suggests a way to conceive of how they work cognitively in “interpreting” the world: [these values] focus attention on selected aspects of the world as it is taken to be by *other cognitive systems* [my emphasis], and provide intricate and highly specialized perspectives from which to view them, crucially involving human interests and concerns even in the simplest cases. I have in mind by richness the idea that [these values] provide very fine-grained media for all sorts of enterprises, not only describing and explaining the world (although not as a science that uses vocabulary apart from natural language would), but writing and reading literature, chatting with friends, and so on. (McGilvray 1998, p. 256)

We note that the connection to objective meaning no longer counts as semantics, in McGilvray's presentation of Chomskian Minimalism. It instead becomes part of the way the (rather mysterious) entities that are the “intrinsic features of linguistic mental events” are processed by the non-linguistic mental realm. Appropriately, this would relegate both the Objectivist's semantic values and the Cognitive Linguist/Psychologist's meaning concepts to a non-semantic realm of (say) pragmatics.

This viewpoint then seems to have denied *any* nexus for Objectivists and Subjectivists to agree upon for the interpretation of language and meaning. What starts as the apparent promise of a method of accommodating both Objectivists and Subjectivists in the realm of meaning has become a theory that neither side will agree to.

³⁷Note McGilvray's use of ‘intensionalist’ here: it means that the mental item *is about* something else—something in the objective realm, for instance. But it could instead be about some other type of *mental* entity. This sort of intensionalism would be friendly to those Objectivists who want their representations to “be about items of reality”. Thus McGilvray is here arguing against Objectivism, and so is presumably attributing a form of Subjectivism to Chomsky's Minimalism.

7.3 Two-Tiered Theories with a Nexus

It is perhaps most clear from Frege's example of "grasping" that we would want the Objective Tier of our two-tiered theory to give an account of truth and possible truth (etc.) by means of features of the world (or of features of possible worlds that are objectively available in the actual world). And we would want the Subjective Tier of the theory to give an account of the way that the world has "presented itself" to an individual to account for aspects of that person's belief states, desire states, and intentional actions. (Some sort of a *Belief plus Desire entails Action* theory would give an example of the sort of Two-Tier theory I have in mind.) This is the sort of theory that Jerry Fodor seems to have had in mind during one of his earlier states of theorizing (Fodor 1980), where *wide* versus *narrow* content was introduced. But I would prefer a different method of "tying together" the Objective and Subjective than Fodor seems to have employed during this period of his thought on the topic.

From such a viewpoint, it also seems clear that the desired *nexus*—the place where the two conceptions of meaning meet—is the word, or more accurately, the *individual sense of a word*, for the sort of reasons outlined in footnote 9 of Sect. 3. On the Objectivist side of this theory, the meanings assigned to these senses (some sort of reference or function on possible worlds) are the atomic units of meaning that form the basis cases for compositional-like (atomistic) accounts of meaning. On the Subjectivist side of this theory, it is the individual understanding of the mental meanings of these senses that combine to form complexes which in turn can interact to form beliefs, give rise to desires, and produce actions. My argumentation in Sect. 6.4 was aimed at the conclusion that this subjective side also needed to be atomistic in its treatment of meaning—although perhaps not compositional.

One recent foray into this realm is taken in (McNally and Boleda 2017), although perhaps motivated by different concerns than the one identified here. McNally and Boleda are concerned with the fact that "in the absence of any specific context" some sentences show a specific and particularly pronounced understanding, whereas when enclosed in a wider linguistic context they can vary so much as to make the original "default" interpretation be difficult to see. They see two sources of information as available for a language understander: the referential and the conceptual. And the two sources may operate differently, the latter being especially sensitive to "affordances" provided by surrounding context. As they put it:

In the case of language, we take its connections to concepts and to the world to be distinct features, each of which facilitates a distinct process of concept composition. When the concepts contributed by two expressions in a grammatical phrase naturally suggest the ways in which they should be composed, and interlocutors avail themselves of such a suggestion, we can say the concept composition is *conceptually afforded*. Alternative, if specific, independently available information about the referent described by the phrase is used to guide the way in which the concepts in question are composed, we can say that the concept composition is *referentially afforded*.

This general picture is rather attractive, and especially in that it locates the relevant nexus to be linguistic items. (Although I would have preferred them to have

identified specifically that they were senses of lexical items.)³⁸ Language is—as it were—“neutral” in how it is to be understood, but different cues give rise to a distinction in ways of understanding, between a referential interpretation and a conceptual interpretation. Like the Frege theory, this seems to give us a way to picture how a two-tier theory could be developed, this time from a Subjectivist point of view. Below I will likewise gesture at a way of looking at a two-tier theory. I intend mine to more strongly separate the (Objective) referential component from the (Subjective) conceptual-connection component. But like McNally and Boleda (2017), I will just gesture at the theory and not give a lot of detail. In Sect. 8, I describe a number of phenomena that need to be dealt with in the subjective tier before any such theory can work.

In brief outline: I see the two-tier theory working somewhat as follows. A sense of an individual content word³⁹ means^o (*mean* in the objective tier) the sort of items favoured by the formal semanticists—individuals for singular terms, sets of individuals for simple nouns, properties⁴⁰ for adjectives, activities and processes for verbs, and so on as required by a truth theory of meaning (or as modified for a theory of truth at a possible world). They mean^s (subjective meaning) some mental concept (of the appropriate sort).

The objective tier is semantically compositional—or at least, is atomistic with some sort of assignment of values for meaning^o in a way that honours the general argumentation for atomicity of meaning. It would be immensely pleasing if the subjective tier could be similarly organized—that all the observed differences present in meanings of complex linguistic expressions could be traced to the differences between the meaning^o and meaning^s of individual lexical items. But it is a long-established result that lexical items do not form the only place where the difference in the meaning^s of complex concepts is to be found. They are also found in the ways that meanings^s combine.

I view the referential portion to be in fairly good shape in these regards (it being the focus of formal semanticists for many decades and all aimed at a similar “truth in the objective world” goal), but that the concept-portion needs to answer many topics that have not been dealt with in the Subjective literature (and hence can’t be directly incorporated into a two-tiered theory, yet). See Sect. 8 for discussion of what other further clarification of these matters is still called for.

The crucial point of a genuine two-tier theory is that there is some point of contact—what I have called a nexus—where each side of the theory can agree upon the features of that contact, but such that each side can decide on what aspect of these features their particular theory should feature. As I see the nexus, it resides in a sense

³⁸And one might also object to this version of a two-tier theory on the grounds that McNally and Boleda’s view of reference is questionable (they appeal to a distributional analysis of the closeness of a word to other words in a corpus, using a multi-dimensional vector analysis). But we will not dwell on that here.

³⁹I do not here wish to discuss the rather different workings of such non-content words as the logical connectives, the copula-like words, the quantifiers, the multi-faceted prepositions, and so on. Various comments about them are below, in Sect. 8.3.

⁴⁰Functions from individuals to truth values (Or: functions from possible worlds to functions from individuals to truth values.).

of a lexical item. Such an item will have two parts: a meaning^o and a meaning^s, corresponding to some objective features of reality and some relevant subjective properties of mentation. Not only are these properties of a lexical sense different, but what can be done with each of them differs, due to their differing properties.

One crucial feature to be investigated concerning these differences is how far the two theories of meaning can diverge from one another. Given that the meaning^o is dictated by the fixed features of reality, the question then seems to become whether the mental realm can veer willy-nilly away from the objective realm. As thus far described, the answer is yes—no restrictions have been placed on the mental representations. It is only by experimentation that we can find out the fine detail of the subjective realm. However, on an *a priori* level we might say that, since the meaning^s is supposed to give the agent's reason for action by the way it can construct the agent's beliefs, it seems that these beliefs (in the main, anyway) can't diverge too, too far from the Objective realm. If the beliefs are all screwy, the agent won't be able to get along in the real world—for long, anyway. And it will be likely that many actions and expressed beliefs will be at odds with others of theirs and of other people's. That is, when meaning^s conflicts in action-oriented ways with meaning^o, bad things can and will happen. And so a basic thought about meaning^s is that in an important sense they can't diverge much from meaning^o. Serious divergence between (say) the objective reality versus the subjective beliefs of the eating habits of grizzly bears might limit the gene pool to be just those who do not have a divergence here.

In less serious cases, for instance where a person's beliefs about snakes have SLIMY as a component, different outcomes might occur. If nothing intercedes, they just go on believing that snakes are slimy and their actions (such as snake-avoidance) is explained on the grounds that they have an aversion to sliminess. When it is explained to them that snakes *in reality* are not slimy, they may alter their SNAKE concept and change their actions; but they may also *not* accept such a claim and continue their previous actions; or finally, they may *claim* they accept the truth^o, and therefore answer that snakes are not slimy when asked, but *not* change their SNAKE concept and not change their other actions that involve snakes. All three of these reactions are possible. But the second and third reactions can't happen too often, since it would be evolutionarily a disaster were people not to change their internal concepts in concert with what occurs in reality. It is this force that keeps the meaning^s connected to meaning^o.⁴¹

Asher (2011) develops the general form of theory that I wish to promote. Central to his notion is that lexical items contain two parts: a contribution to objective semantics and what he calls a TYPE. The objective portion is especially simple, but the effect of TYPE is intricate and detailed, with many different aspects, including pre-supposition enablement, context sensitivity effects, sense transfer, metonymy, predicate transfers, discourse effects, explanation of *qua* phrases, dedicative phrases, and other very subtle meaning-topics.

⁴¹James Hampton (p.c.) remarks that evolutionary pressure may not always work this way. Religious beliefs may serve to strengthen social groups to greater cohesion (and success) while being relatively immune to falsification by meaning^o.

Here are some of Asher's introductory remarks on his vision:

The guiding idea... is that almost all words will have single and simple lexical entries. Words like the nouns *cat*, *lunch*, *book* or the verbs *kill*, *read*, and *master* denote simple properties or relations, and so accordingly the logical forms that specify their denotations are very simple. For instance, *cat* has the lexical entry $\lambda x \text{Cat}(x)$ ⁴²... If words have simple entries and make simple contributions to truth conditional content, they come with a rich amount of information about the TYPE assigned... within them. These TYPES will guide predication and be responsible for fine-grained differences in lexical meaning. When words are combined together to form clauses, sentences, and discourses, the TYPE associated with various terms will interact in complex ways. I will introduce operations of TYPE adjustment in response to TYPE mismatches between predicate and argument that correspond to the accommodation of a TYPE presupposition, or more generally speaking, the justification of such a presupposition.... The effects of these TYPE adjustments... is that the logical form for a clause will contain elements that are not present in the lambda terms for the constituent words themselves. Predication involves not only applying a function to an argument but also operations of adjustment corresponding to type presupposition justification. (p. 22)

By his notion of "TYPE adjustment" and "coercion" he includes a very wide range of examples that are familiar from the Cognitive Linguistics literature, such as

- (4) a. good lunch versus good children
 b. Mouse isn't very tasty unless you're a cat
 c. John started a cigarette/started a car/began the sonata/started the novel
 d. John liked the dress with the flowers/liked the garden with the flowers

And he mentions certain entailments as also falling under the sort of "coercion" he is concerned to capture, such as

- (5) John started the car *defeasibly implies* John started the engine of the car

In Asher's account, another aspect of the TYPE portion of lexical entries concerns how they interact with the TYPE-entries of other lexical items in a longer stretch of discourse. This interaction is governed by the underlying logic of the mental side: his adaption of classical type-theoretic logic to become a proof-theoretic semantics for the TYPE side of his theory, as I suggested in Sect. 6.2.

I've argued that TYPES are concepts, mind-dependent entities with fine-grained content. TYPES have an internal semantics that is given at least in part in terms of the rules by which they combine with other TYPES; they are proof-theoretic objects. TYPES can be associated with other TYPES – what we might think of as *traits*; each trait we associate with a TYPE is a constraint on the introduction rule for the TYPE.... TYPES as proof objects provide an internal semantics for natural language sentences and discourses that complements the external semantics given by ordinary intensions. We can even distinguish between different conceptualizations of the same physical object or of the same property, so in some respects the structure of types has the capacity to make finer distinctions in meaning than intensional semantics can. But the purpose of the internal semantics is not the same as that of the external semantics. Intensions are the soul of a theory of meaning – they are needed to determine truth, reference, and other *external*, semantic properties that link language to the world we talk about. TYPES and their adjustments are the heart of a theory of predication and

⁴²This lambda term is to be interpreted as "the feature of being an x such that x is a cat", or more simply, "the property of being a cat."

responsible for other properties of meaning. . . . This is semantics but it is an *internal* matter, something that speakers do “in their heads.” (pp. 44–45)

The one aspect I find missing from this otherwise compelling story is that there is no room here for individual variation of meanings. In particular, the mental realm—as captured by the TYPE dimension of the dual-aspect lexical items and carried along via the proof-theoretic semantics to full sentence-thoughts is the same for all people, and it furthermore (appears to) contain exactly the *objectively true* features of the encyclopedic facts in the word-TYPES when these are manipulated so as to give the meaning^s for the sentence-thought. It thus cannot deal with issues of false beliefs about items in the world; it will not give any account of why different subjects give individually different results on the many tasks about meaning that have been administered over the decades in cognitive psychology.

Although Asher’s theory is not described this way, perhaps the alteration I would advise is to make his theory be a theory for an *individual agent*, allowing this agent to have his/her own set of encyclopedic “facts” for every lexical sense.⁴³ Presumably there would also be individual differences in the underlying inference mechanisms of the proof-theoretic semantics also, another topic not envisaged in Asher’s picture.

But this now would leave us without any way to accommodate “real communication”, where the two conversationalists are “really” talking about the same (mental) things. While there can always be some check against “the objective facts” due to the objective portion of meaning in Asher’s theory, this will not help conversation that relies on having some mutual understanding of the TYPE-portion of meaning—the meaning^s.

I leave this problem hanging, as a very important topic still to be solved by any two-tiered theory. And I turn instead to other aspects that need to be dealt with in the subjective, TYPE portion of the theory.

8 (Some) Aspects of Current Theories of Mental Concepts that Require Attention

Objectivist theories of meaning are just fine—so far as they go, which is not very far when it comes to mental life. (And one might reasonably think this is a very major shortfall!) On the other hand, Subjectivist theories need some further work, even within just the mental side of meaning with which they’ve been concerned. This section is about a group of topics that need to be addressed more fully than they have been, in order that one can take the Subjectivist theory of meaning seriously enough to be included in the meaning^s portion of a two-tier theory of meaning. I mention three areas that seem to be “obvious problems” that need to be dealt with before it is possible to develop any deep mental theories that invoke concepts., including the

⁴³I might mention that Asher’s picture does not take individual senses as the basic items, but rather uses TYPE-coercion to generate the distinct senses as required in any particular sentence.

desired concept-oriented portion of any two-tiered theory of meaning of the sorts mentioned in Sect. 7.

The first area concerns the type of concepts (and the details concerning them) that such theories envisage, and it has several subparts. The second concerns a venerable puzzle that should be considered by anyone when they embark on proposing any account of meaning, whether Subjectivist or Objectivist. The third area requests Subjectivists to construct even a basic version of what is already present in Objectivist accounts. . . and then asks whether there is any reason to go to the trouble.

8.1 (Some) Messy Details for Conceptual Theories

I mention in this subsection a series of interrelated topics that seem to me to draw attention to the inadequate attention that is paid to the vast majority of linguistic data by theories of concepts and conceptual combination. These phenomena are seen as central in Objectivist accounts, but which seem not to figure in most Subjectivist accounts. This holds, I think, not only for the psychological theories that postulate concepts to be tested against certain behavioural responses by language users (where these theories don't pretend to test a full range of linguistic phenomena), but also for the cognitive linguists who are very aware of many of the deep and peculiar intricacies of natural language but who nevertheless use the overly-simple results from the psychological theories to account for them. (Sometimes just assuming that "more of the same" will accommodate the intricacies.)

I write this from the point of view of an Objectivist, formal semanticist. All the topics discussed below are thought to be adequately dealt with on the Objectivist side of a two-tier theory (at least by Objectivists⁴⁴). Presumably, Subjectivists would have their own lists of topics that Objectivists need to deal with before any two-tier account could be made to work. But someone else will have to provide that list.

8.1.1 Basic Problems with Singular Terms

Proper names and other singular terms do not seem to feature heavily in the Subjectivist theories, although it seems pretty clear that these theories would somehow wish to generate concepts in an individual's mind corresponding to them. But proper names and other singular terms have a number of well-studied peculiarities that seem to tell against mental versions of concepts for proper nouns.

As we have seen, Subjectivist accounts claim that concepts embody "encyclopedic knowledge" of the items of which they are concepts. Thus, even though NATIVE OF AFRICA is not part of the meaning of *lion*, it is a part of the concept LION, in the Subjectivist's view. The philosophers of language have advanced a number of arguments that have seemed to Objectivists to show conclusively that this cannot

⁴⁴Speaking generically.

be true for proper names. Although TEACHER OF ARISTOTLE might be part of the stereotype or encyclopedic knowledge one has of Plato, we could discover that it was all a mistake in the transcription of one ancient document, and that in fact Aristotle was never in Plato's Academy. Nonetheless, even after discovering this, when we use the name *Plato*—as I just did in the last sentence—I was referring to the same person as I always referred to. And, it is usually added, this could hold of *any* of the encyclopedic properties of Plato—indeed, of *all* the encyclopedic properties (Kripke 1980).⁴⁵

A theory of concepts should in one way or another be able to accommodate this finding. How can a singular term concept such as PLATO evolve so as to lose some, and possibly all, of the properties given by our encyclopedic knowledge?

Similar remarks could be made about other singular terms. Unless one wants to make the concept answering to *The current President of the United States* be the same as that corresponding to *The person I believe to be the current President of the United States*, one cannot include in the concept any stereotypical properties that are not essential. But if you *do* decide that these were the same concept, you could never discover that you were wrong! (Because they are the same concept, they said; so any change just continues to make them (it?) the same concept.)

8.1.2 Basic Problems with Singulars versus Plurals, and Quantified Noun Phrases

It is surprising that there is not more mention of the difference between concepts that correspond to singular noun phrases in natural language and concepts of the corresponding plurals. A survey of any of the literature will show that simple nouns like *dog* get associated with a concept that is multiply-instantiated—that is, this concept applies to (or represents) any dog. Intuitively it corresponds to a plural (or perhaps to an unmarked-for-number) term. But in either case it remains somewhat of a mystery how this can be employed to indicate a *particular* dog or to assert something true of a single dog, and at the same time to assert something that is true of some but not all the dogs. Clearly there needs to be some way to incorporate numerals and quantifiers in the theory. Surely we want to conceptually differentiate the sentences *A dog is in the yard* from *Two dogs are in the yard* from *Dogs are in the yard*. Yet it seems quite difficult to have this sort of modification of simple nouns when using only the type of conceptual combination that is usually described in Subjectivist writings.

⁴⁵This claim—that an object identified by a proper name *could* lack all the properties traditionally associated with the object and still be referred to by the proper name—follows from any version of the “historical-causal theory” of reference (of proper names). One intuition supporting the claim is that any of the initial identifying features in fact *needn't* have been true—they need only *seem* to be true to the initiator of the use of the proper name. And then those who adopt this name into their idiolect just assume these to be factual claims about the initial referent. But as this consideration shows, the truth needn't be a part of this assumed set of features.

Furthermore, natural language contains a myriad of different types of quantified noun phrases such as *all people* just as much as it does adjectival modifications of nouns such as *tall people*. This suggests that since conceptual combination is used to generate a concept TALL PEOPLE from the concepts TALL and PERSON,⁴⁶ something akin to it should be used to generate the concept ALL PEOPLE. But this would require a concept ALL, much along the lines of the concept TALL. And this has seemed to most philosophers an unlikely suggestion: quantifiers are higher-order, that is, they are properties of properties. Here ALL would need to be *a property of* PERSON, not *one of* PERSON's properties.

Without a theory that accommodates the quantifiers (and there are many more than just *some* and *all*), the theory of concepts is essentially unable to accommodate the vast majority of our language and our mental life.

8.1.3 Basic Problems with Noun-Oriented Concepts

It is striking, although perhaps explainable, that theories of concepts start by presuming concepts corresponding to certain lexical nouns: the concepts PET, CAT, MOUNTAIN, CABIN, . . . , and so on. As mentioned above, and as is ubiquitous in the literature, the history is that very much of this research program's time is now spent on trying to determine just what the nature of these concepts should be: feature lists, prototypes, stereotypes, schemata, dynamic whatevers, and so on. A test for correctness of the resulting theory is then developed, and it consists in how well "conceptual combination" works according to these tests. But the examples thereby considered for these tests, at least in the first instances—PET CAT, MOUNTAIN CABIN, CHERRY BLOSSOM, . . . —seem to be carefully-chosen noun-noun compounds, since they don't include things like PET ROCK, CABIN CAT, CABIN FEVER, . . . , for which there are intuitively plausible interpretations but where the method for arriving at these interpretations would seem to be in conflict with the methods established for the initial examples. The difference of interpretation between IRON GATE and WATER METER, obviously requires some distinction between "material substance of" and "recipient of" to formally yield for us a gate made of iron versus a meter that controls or measures water.⁴⁷

Furthermore, as presented in the literature, the various mechanisms for conceptual combination in the literature are—at least in the main and for the most part, as emphasized in (Hampton 2011b)—"symmetric": for example, one considers the feature lists from two prototypes and constructs some sort of new feature list corre-

⁴⁶Note the insertion here of some sort of plurality operator to form the combined concept! As I said, there needs to be some way for concept theory to generate the singular-plural distinction.

⁴⁷There are, of course, the hints from those Subjectivists who propose types of "metaphorical interpretations" and rely on magic for the formalism. And more plausibly, there are suggestions about how there can be underlying principles governing differing types of this sort of compounding. For example, there are nice principles stated in (Wisniewski 1997; Gagné 2002; Gagné and Spalding 2006; Spalding and Gagné 2008; Gagné et al. 2016), but I would prefer to see some sort of algorithm instead.

sponding to the combined phrase. From PET and FISH, one constructs a prototype for PET FISH. But as a symmetric procedure, this very same result should be generated from the phrase *fish pet* also, even though it is very difficult to give an interpretation for that noun-noun compound, as it occurs in natural language.

But matters are more complicated here, since some noun-noun compounds do in fact have interpretations for either order of the lexical nouns, but the interpretations are different: *grease gun* versus *gun grease*, *talk show* versus *show talk*, *truck delivery* versus *delivery truck*, *school teacher* versus *teacher school*, *flea circus* versus *circus flea*, for example. All this is one topic that should be addressed at the very beginning of any theory of concepts. Hampton (2011b) outlines various types of “non-intersective” conceptual combination, referring to some seminal work in (Wisniewski 1997). Word order (in English) gives rise to much of the “direction of interpretation”: the head of a noun-noun phrase tends to be on the right. Although this is correct (generally) for English orthography and verbal production, it is not at all clear why it should dictate how *the concepts themselves* should combine. Is it that the concepts GREASE and GUN can combine in both ways freely, but that one is associated with the linguistic phrase *grease gun* while the other is associated with the phrase *gun grease*? But now what is needed is some principle that determines why one of the concepts is associated with the one linguistic phrase while the other concept is associated with the other phrase. Is there anything at all intrinsic to the concepts that dictates one over the other? (For some thoughts on this dark matter, see Hampton 2011a.)

Of course, disjunctions of nouns occur in language, as well as conjunctions. So we would like a satisfying account of the type of concept one gets for *either a dog or a cat* and (more difficult, perhaps) *either an aggressive intruder or a figment of the imagination*.

Most theories of concepts also explicitly introduce concepts corresponding to adjectives, so that we have concepts like HAPPY, TALL, RED, etc. Now we can have conceptual combination to yield concepts corresponding to *happy pet*, *green cabin*, *red rock*, ...; but in these cases we *do not* have symmetry of conceptual combination—*pet happy*, *cabin green*, *rock red*, ... are not normally grammatical when the second item is an adjective, unlike (most) cases where they both are nouns. (In those cases the phrases are grammatical but may be semantically incoherent. This seems a natural thing to say, but one can wonder whether the Subjectivist account of concepts will allow a distinction between ungrammaticality and semantic incoherence.) This argues that concepts corresponding to adjectives have to be differentiated from those corresponding to nouns. And as a consequence, conceptual combination of concepts of adjectives would probably be different from that for nouns: the adjective phrase *rich*, *happy*, and *well-adjusted* maybe corresponds to a conceptually very different sort of concept from the concept corresponding to the noun phrase *men*, *women*, and *children*.

I do not suggest that *no* work has been done on these topics—disjunctions and negations were discussed long ago in (Hampton 1988a, 1997). But as I will mention in Sect. 8.3, a theory that explains the workings within the realm of concepts that is of a scope as broad as the Objectivist account still remains to be articulated.

8.1.4 Basic Problems with Verb-Oriented Concepts

One presumes that lexical verbs also correspond to mental concepts in this overall picture, and that since verb phrases can conjoin (*eats, drinks, and falls asleep* as well as *marries, works, breeds, and dies*); so presumably the corresponding verb-oriented concepts can conjoin. But as I will try to demonstrate in this section, it seems reasonable to suppose it will be a different type of concept combination than in the nominal side of things, because of the particular differences there are between the noun-concepts and the verb-concepts, as well as the higher degree of complexity that verb-concepts display.

Conceptual combination seems to have been designed to deal with what might be called “conjunctive (or: intersective) concepts”, even though the resulting feature lists (or whatever used to keep track of properties) is not a simple conjunction of the feature lists of the subparts. A part of the rationale for concentrating on this could be that English allows easily for conjoined noun phrases (*Bob and Sally, Men and women, Both small dogs and short snakes, All baseball players but only some Olympic high-jumpers, etc.*) And also that the interpretation of adjective-noun combinations has historically been seen as conjunctive (although modern formal semantics treats them otherwise). But it is not so clear to me how any similar technique will work with what might be called “disjunctive concepts”. After all, English pretty easily allows for disjunctive noun phrases: *Juan or Alice, Men or women, Either Americans or Canadians, etc.*⁴⁸

As remarked just above, verb phrases easily conjoin just as much as noun phrases do, but I claimed that the conceptual items corresponding to verbal items will require some different sort of combining method from the ones in the noun realm. Verb phrases can disjoin, like noun phrases can, for linguistic phrases such as *is either swimming or eating*. Another place where noun concepts and verb concepts appear to be different is in the case of linguistic conjunction. But given that the conceptual combining method for verb-concepts is different from that of noun-concepts, due to the inherently more complex nature of verbs, it seems likely that the method for forming disjunctive verb-concepts would have to be different from that of forming disjunctive noun-concepts.

A special difficulty for theories of concepts attends to those concepts that correspond to non-monic verbs. That is, to relational verbs—ones that take direct and indirect objects.) For example, the verb *kick* requires a direct object, and so any corresponding concept will require something corresponding to this. As with any verb-concept, there will need to be some notion of a subject of the kicking-concept. But in the case of transitive verbs, the corresponding concept will also require the notion of a recipient of the kicking. It seems that some as-yet unspecified notion of conceptual combination is required for this. (And a further unspecified one for

⁴⁸In the olden days there were various works on disjunctive concepts, both learning and using them (Furth 1963; Conant and Trabasso 1964; Seggie 1969; Bar-Hillel and Eifermann 1970; Trabasso et al. 1971; Newstead and Griggs 1983), but these have seemingly diminished in recent years, and the few more recent writings have focussed on reasoning with disjunctions. However, there are works by, especially, James Hampton on aspects of these topics: (1988a, 1997, 2007, 2012).

indirect objects.) As I have mentioned above, noun phrases are not only these simple lexical nouns nor do they merely include only the conjunctive and disjunctive nouns, but they also include proper names and quantified noun phrases. It is most unclear how the imagined-but-as-yet-unspecified-conceptual-combination that creates verbal concepts with direct and indirect object concepts really works, even in simple cases like *loves dogs*, but it seems quite a bit more difficult to imagine any such method of conceptual combination that would work with *loves all dogs* or with *loves some dogs* or with *loves many dogs*, or with *loves all dogs but only a few cats*, etc., and give the intuitively correct distinctions among them all. (And with indirect objects we can have verb phrases like *gave many children numerous gifts*. It is again most unclear how such complex concepts can be built up, especially since the underlying concept for the verbs and noun phrases are unspecified as yet.) A further related difficulty arises with those verbs that can also be used in an ergative manner as in (6-b):

- (6) a. Mary boiled the water
 b. The water boiled
- (7) a. Kim tripped Sandy
 b. Sandy tripped

Yet another difficulty that seems not often addressed is the difference that English (and other languages) shows between “stative” and “eventive” verb phrases. (*Knows Japanese* is a stative verb phrase whereas *is speaking Japanese* is eventive, reporting an activity. Other types of eventive verbs indicate achievements or accomplishments. One linguistic test is that eventive verbs allow the progressive (in English), whereas stative verbs don’t: **John is knowing Japanese* is not correct, but *John is speaking Japanese* is.) This distinction also brings up the difference between *ability* versus *activity* senses. And that is another difficulty for concept theories. When FLY is conceptually combined with TWEETY,⁴⁹ is the result a representation of the natural language *Tweety can fly* or *Tweety is flying*? There is a very large literature in the intersection of philosophy of language and formal semantics that deals with issues involved with this distinction, and these ought to be reflected in theories of concepts (e.g., Parsons 1990; Schein 1993; Casati and Varzi 2015).

8.1.5 Basic Problems About Other Noun- and Verb-Oriented Concepts

The fact that the various major categories of syntax allow for conjunctive and disjunctive combinations, and that the Subjectivist theory thereby requires some type of concepts corresponding to them, leads one to wonder whether there can be *negative* noun-concepts, or *conditional* noun-concepts, according to theories of mental constructions. Similarly, one wonders about negative or conditional verb-concepts.

Although negation in English can apply pretty easily to verbs/verb phrases, and to adjectival/adverbial phrases with only a slight increase in difficulty, they apply with

⁴⁹ Assuming that the concept theorists have a way of representing proper names as concepts.

rather more difficulty to noun phrases. Can there be a concept NON-CAT? Maybe it's more plausible to have a concept NON-STARTER? But then there are more complex noun phrases, and if concept-theories are extended to handle arbitrary noun conjunction/disjunction, it ought to be extended with negation so as to handle the concept associated with *short adults but not tall children* or *neither oil nor coal*, and the like. One wonders, then, just how negative noun-concepts are going to be handled. My own intuition is that it will require very different applications to the different sorts of concepts (that is, it will apply differently to different grammatical categories). It would be good to see a careful account that will accommodate all this.

Since English easily allows for conjunctive and disjunctive noun phrases, adjectival phrases, adverbial phrases, and verb phrases, and allows easily for negative verb phrases, adjectival phrases, and adverbial phrases, it seems that the usual “logical connectives” can apply to any of the usual linguistic categories other than nouns. But then one wonders about non-standard connectives like conditionals: English allows conditionals in all the linguistic categories (although it is rather more stilted in the noun case): *is happy if he is eating* (conditional verb phrase), *if not tall then at least not short* (conditional negations and adjectives); *morally if not legally* (conditional adverbial phrase), *crocodiles, if not all reptiles* (conditional noun phrase, with negation). So it seems that there needs to be some sort of conditional-forming conceptual combination. Needless to say, this doesn't at all look like the usual conceptual combination.

8.2 *Concepts, Sentences, and Bradley's Regress*

We observe that there has not been any account of the interpretation of *sentences* given by Subjectivists. That is, there is as yet no account of a (mental) *judgment*. Even granting a singular concept SANDY and a verbal concept SWIMS, there is not yet any judgment that Sandy is swimming.⁵⁰

This is related to another question that concept-theorists seem not to have addressed about an old philosophical puzzle and to a more modern apparent paradox. The puzzle is: how can mere *words* combine to form a *sentence*? That is, what is the difference between a list of words and a complete sentence? How can it be that a list of words becomes something different—something that expresses an entire proposition? All sentences, even very simple ones, give rise to this puzzle: Consider

My list of words:

[Word #1] Kim

[Word #2] runs

My Assertion:

Kim runs.

⁵⁰Or that Sandy *can* swim, depending on the account given as an interpretation of the concept SWIM.

As various writers in the past have claimed, the former is “inert” and “is a mere congeries”, whereas the latter actually “does something” in making an assertion that describes a feature of Kim.

Note that a structurally identical puzzle arises if ‘Word’ is replaced by ‘Concept’ and ‘My assertion’ is replaced by ‘My judgment’. Both Objectivists and Subjectivists are subject to this puzzle, and they both should give some sort of answer. However, Subjectivists seem not to have thought of a difference between the noun-phrase conceptual combination {SANDY; SWIM} and the sentential conceptual combination {SANDY IS SWIMMING}. Nor has it seemed to occur to them that in order to give a concept-analysis of some noun phrase like *a man who is tall*, there needs to be a sentential analysis of the embedded relative clause (as well as an analysis of the contribution of ‘a’, as mentioned above in Sect. 8.1.2).

This puzzle is at least as old as Plato, and has provoked considerable discussion over the centuries. One type of answer might be to postulate a copula-like intermediary, call it ‘exemplification’, with the idea that it is this intermediary that causes the other items to merge into a full sentence. But as one can pretty easily see, this just leads to some sort of infinite regress: for, now there needs to be a further explanation of how it is that ‘Kim’ (or the person corresponding to that name) can combine with this newly postulated object to then be able to combine with ‘runs’ (or the object corresponding to that predicate). This is known as Bradley’s Regress, or The Unity of the Proposition/Sentence problem.

Objectivists have proposed various answers to the Unity problem, one of the most popular being some variant on the metaphor of *saturated/unsaturated* from Frege (1892a). The idea is that singular terms are such that they are designed to designate “independent” or “saturated” objects, whereas predicate terms are designed to designate “unfilled” or “unsaturated” items, namely “concepts”. (Recall that for Frege and the Objectivists, a concept is a non-mental, abstract entity.) When these two different kinds of entities “meet”, the saturated object “fills in” the unsaturated object, producing a new, saturated object—a complete sentence (which expresses a “thought”, in Fregean terminology... where “thought” is also understood as a non-mental, abstract item).

Well, whatever one thinks of such metaphorical accounts (and something akin to this is adopted in Formal Semantics’ model-theoretic theories), it is at least some *sort* of explanation, showing that there might be some wiggle room in Objectivist theories that would provide an appropriate description of what is going on.

However, it seems that Subjectivist accounts are not in such a good place, at least not if they hold that lexical concepts designate mental items like concepts. For it seems that this would presume there to be a concept of INSTANTIATION—sort of a mental version of the linguistic *exemplification*, and now it appears that the Bradley regress can start. Subjectivists of the sort we have been considering will respond by claiming that *conceptual combination* is how they handle this topic. But it should be noted that conceptual combination in these works is defined for noun-noun and adjective-noun combinations (perhaps also for adverb-verb combinations?). But this won’t work for sentence generation: perhaps the conceptual combination of DOG and

TAIL will yield a concept that somehow contains dogs with tails, but it isn't a concept corresponding to *Dogs have tails*.

The Bradley Regress puzzle should encourage Subjectivists to engage in serious theory construction for such a further mental operation. But it has to be an *operation*, and not a concept. For, a concept will only reinstate the Bradley Regress. So, besides conceptual combination, and whatever other related operations are required to accommodate the types of other combination I mentioned in the previous subsection, it is absolutely necessary that there be some further type of operation that forms "complete mental thoughts" out of concepts. And what should these be called? It seems to me that it would be much too confusing to simply reuse the term 'concept', so I will start using 'thought'. This maintains the parallel of usages with the (Fregean) Objectivists: they use 'concept' to refer to an objectively existing, abstract object; they use 'thought' to refer to an objectively existing, abstract object that can be evaluated in a truth-valued manner (or other manner for different speech acts than assertion). Subjectivists in the proposed usage employ 'concept' to refer to subjective items of individual minds that correspond to items of the non-sentential syntactic categories; they use 'thought' to refer to the mental complex that can be evaluated as a complete judgment.

8.3 *REAL Negation, Conjunction, Conditional, and Disjunction*

We've seen that English easily allows noun phrases, adjectival phrases, adverbial phrases, and verb phrases all to be conjunctive and disjunctive. English also allows for negative verb phrases, adjectival phrases, and adverbial phrases. And there is also some opening for conditional phrases of these various linguistic sorts. But this is not—at least, is not according to most philosophical and formal semantic theories—the major use of conjunction, disjunction, and negation. For that use, say the Objectivists, we negate, conjoin or disjoin sentences.⁵¹ But there is no provision within the Subjectivist conception of a mental regime of concepts for this type of combination. (For one thing, there has not been any explanation of how to generate thoughts from simpler concepts.)

But suppose we grant the Subjectivists that they can generate thoughts—that is, that they have some sort of concept combination that will generate judgments rather than just merged noun-like concepts. And further suppose they have some operations to generate concepts that are negations, conjunctions, and disjunctions of noun-, verb-, adjective-, and adverb-concepts. They still haven't done what most formal semanticists see as the central point of these notions: the ability to form thoughts which correspond to complex statements. The sort of work that is in the literature seems aimed at somewhat different targets, such as *describing* the various effects that different of these constructions have in specific cases as opposed to large-scale

⁵¹Or propositions, if one is in the market for using abstract entities rather than mental concepts.

rules that would apply universally. Those problems seem simply not to be the ones that are addressed, and so it is not so surprising that the investigations don't yield such overarching theoretical proposals.

Instead there are "local" accounts. But note, for the present case under discussion, that negation, conjunction, disjunction, and conditional cannot be concepts, under the pain of running again into the Bradley regress. But also it is not so easy to deal with these complex sentential operations as it is to deal with the corresponding ones operating on concepts. One example is discussed in Pelletier (1998), where the idea was that negation might be an operation that transforms "regular" concepts into some other type of concept. For example, the conceptual operation of negation might transform a verb-concept like IS TALL into a disjunctive concept like IS EITHER SHORT OR MIDDLE-HEIGHT. Or, it might transform the verb-concept IS RED into IS SOME COLOUR THAT IS CONTRARY TO RED, which would presuppose that colour concepts (and maybe all concepts generally) fall into "incompatibility classes", where nothing can manifest any two different concepts in that incompatibility class. And there are other possibilities.

This has the effect of making all sentence negations into predicate negations, and it is far from obvious that this sort of move can really be maintained. It is also far from obvious whether there is anything similar that can be done with the other "thought-combining concepts".

But even if there were some such way to endorse all the earlier additions to the mental world of concepts, and even if there were some mental analogue to the Objectivist notion of "saturation" as a way to solve the Unity of the Thought issue, it seems that the Subjectivists would just be mirroring the Objectivist account, but with a very complex and otherwise unmotivated set of operations. And then, such rumination makes one want to ask: Why bother? Even if it could all be done, would there have been any advance at all in explaining how natural language describes the world? Or even what is possibly more to the point from the Subjectivist point of view: has there been any advance at all in explaining how people use natural language to describe the world? It is hard to see what it might be.

It seems to me that the only reasonable answer from the Subjectivist camp to such rumination would have to emphasize the role of concepts in a *Belief+Desire*⇒*Action* framework. It is in this realm that we can see people making use of concepts that do not correctly mirror reality. ("She eats brown eggs rather than white eggs because she thinks they have less cholesterol"). It seems that an appropriate account of this subject's concepts would somehow explain such an objectively strange behaviour by using some such account of his/her beliefs and desires.) From the Objectivist camp, presumably the only appropriate answer to such rumination would be to show how "objective truth" can alter conceptual connections. And I think it is in this interaction of the two realms that the correct account of meaning lies, including why both are required.

9 (Some) Concluding Remarks

I have tried in this paper to lay out the (sometimes) subtle interactions that are inherent in a number of distinct commitments held by many theorists who study the semantics of language. A major division among these theorists is between those who think to study language by an investigation of how language describes or attaches to “the world” versus those who think of it primarily in terms of how it describes “mental structures” of users of the language. This is what I called the Objectivist versus Subjectivist distinction. This difference in viewpoint impinges on several topics of interest. A simple but confusing topic concerns the difference in intent when the two groups employ terms like ‘concept’ and ‘thought’, where Objectivists, following Frege, think of these as “public”, objective entities that can be “grasped” by many different language users, while Subjectivists think of them as “private” and existing only in the mind of a single language user. (Hence, any talk of a shared concept then becomes an issue of the similarity between the different concepts and thoughts possessed by the different speakers.)

Although it might appear on the surface that this is the entire difference between the philosophers-formal semanticists (the Objectivists) and the cognitive linguists-psychologists (the Subjectivists), there are further ramifications that may not be so obvious from just that simple difference. For, it is this that makes the Subjectivists think that meanings are not compositional, because of data about such topics as conceptual combination (for the cognitive psychologists) and the choice of meaning specificity by a speaker and hearer (for the cognitive linguists). And on the other hand, the separation of a speaker’s mental life from “meaning in the world” makes the philosopher/semanticist group able to hold on to semantic compositionality even in the face of the psychological data. They say that since mental concepts are not compositional, that is yet further evidence (besides what they believe they have shown) that such concepts can’t be meanings—since semantic compositionality is true.

This difference furthermore opens the door to distinct accounts of how meanings are to be specified. In the Objectivist version of this, meanings are organized in what I called an Atomistic manner. There are smallest units of meaning (words, or whatever) and these smallest units combine in one or another way to form the meaning of larger and larger units. The dispute over compositionality within the Objectivist camp then becomes whether this combination follows a notion of Functional Compositionality, or whether there can be exceptions. Note that this dispute is entirely about “meaning in terms of the objective world” and has no direct reflex into the subjective contents of a person’s mind.

But when it comes to the Subjectivist side of things, issues of compositionality can take two different forms. One is a kind of mental “mirror” of the Objectivist viewpoint, where there are smallest units that bear meaning—lexical items, let’s say. But these lexical items mean concepts and these concepts have a very rich structure, replete with “encyclopedic facts”. Nonetheless, the notion of Functional Compositionality *could* still apply to this picture, if only the meanings of more complex phrases could be described in terms of a function applied to the very complex mean-

ings of the simpler parts. As I mentioned, most researchers think this is not possible, although in my mind the jury is still out on this until the Subjectivist theory of complex concepts is filled out more fully. As it is currently given, the majority of research concerns the concept-combination of two particularly simple noun-based concepts. And for them it does indeed seem to be difficult to give a compositional treatment—although one should always keep in mind that functional compositionality is a particularly powerful tool with many unexpected strengths to capture under its aegis all sorts of phenomena in the objective realm. (Many of these are recounted in Pelletier 2016.) Perhaps the same strengths can be employed in the mental realm, especially when the “meaning” of concepts is explicated in an atomistic way, as outlined in Sect. 6.4. But it is difficult to know in the absence of a more developed theory of complex concepts.

The other way that a Subjectivist theory of meaning might go is to say that a part of each concept’s meaning is all the liaisons that concept has with other concepts. This is different from the other form of complex concepts, because here each concept has every other related concept as a part of its meaning. And since this holds of these related concepts also, it follows that what is related to them becomes a part of the original concept’s meaning. As one can see, this makes conceptual meaning be a “network” or “web”, where every concept involves every other concept, and no concept has any meaning of its own.

The former of the two Subjectivist theories is—like the Objectivist theories—Atomistic in nature. While the latter theory is Wholistic, leading to a holistic theory of meaning. To my mind, there are a number of devastating objections to any version of this latter theory. And so I rejected this version of Subjectivism, but not the other version, which is the general sort of account that I see being plausible for one half of a two-tier semantic theory.

It is not surprising, though, that Subjectivists of both types claim that semantic compositionality is false. Holistic semantic theories are, by definition, not Atomistic, and hence stand no chance of being compositional. The other type of Subjectivism, where concepts corresponding to lexical items contain a wealth of encyclopedic information, could be compositional if these presumed complex meanings happened to be organized appropriately. But at least so far as that idea has been developed by this point in time, they do not appear to be compositional. For the most part, the theoretical apparatus used for determining the meaning of complex mental items is so underspecified that it is in fact unable to create such complex items, much less create them in a way that would obey the strictures of semantic compositionality. But on the other hand, the Objectivists have not done two related things. They have not said how the lexical items that are at the terminal nodes of a syntactic tree are assigned their meanings other than to say that these meanings are “from the world”; and they have not described what relationship there is between this “meaning in the world” and a language-user’s mental life or concepts. (Neither for the lexical items nor for larger linguistic units like sentences.)

Do the Objectivists need to give such descriptions? After all, they might say, the enterprise of formal semantics is pretty much defined as articulating how sentences encode information about the world, and it has no particular place for human menta-

tion, except perhaps in the case of finding semantic treatments for attitude reports.⁵² And the task of a formal semantics is to explain how complex meanings are constructed from simple meanings. For that purpose one needs to start with the simple meanings as a given.

Although both of those claims are true, at least as the task is conceived by the Objectivists, they seem also to have forgotten their rationales for finding semantic compositionality attractive in the first place: the arguments from understandability, from productivity, from learnability, and from creativity. Certainly all of these arguments make the claim that the justification of semantic compositionality lies in the way it facilitates mental processing. And thus, even if the picture that “language is about the world” is correct, these arguments nevertheless presuppose or require that language also embody features of one’s mental life.

A natural place for the Objectivists to find this missing connection is to claim that the terminal nodes of their meaning structures are given by mental items, such as concepts. And then they can go ahead and—from this starting point—be entirely compositional. In fact, I’m sure that most of the philosophers-semanticists who give any thought to the issue probably believe something along those lines. But let’s look at what this would entail.

If the psychological results are correct—and they have been replicated over and over, at least in regards to simple binary combination—that mental concepts allow both for differences in degree of typicality of the concept and also for differences in degree of membership (Hampton and Jönsson 2012, pp. 386–87), then this would have to be a feature of the meanings at the terminal nodes of the semantic structure. The lexical item ‘furniture’, for instance, would have to admit *in the world* of different degrees of truth (and not merely different degrees of typicality) for ‘This chair is furniture’ and ‘This clock is furniture’. I think philosophers-semanticists would not countenance this.⁵³

In response to all these competing pressures, I surveyed the prospects for a two-tiered semantic theory: one aspect being the Objectivist’s “connection to the world” and the other aspect being the Subjectivist’s account of mental activity. As it turns out, some of the works that claim to provide a two-tier theory actually are within just one of the Subjectivist or Objectivist camp, in that both of their tiers have their foundational starting point, or common nexus, entirely within one camp. I also mentioned a few theories that appear to be genuinely two-tiered, but noted that they require further input, especially from the Subjectivist side of the theory. To agree to anything like the two-tiered theory I outlined, one needs to accept two claims: First, that linguistic utterances are to be judged as to how well they “fit” the (external) world. One way to put this, and I did, is to say that *what is meant*^o by an utterance concerns the actual world, and not any conceptual representation. The two come apart when

⁵²Such as *John wants to marry a princess*, the ambiguity of which (a specific princess versus any princess) seems to invite an analysis that would involve John’s mental states. And also, since human mentation is actually occurring in the physical world, it must find a proper place even in Objective semantics as something to be described.

⁵³I’m still speaking generically!.

the representation is incorrect in one way or another, and evaluation—whether the utterance is true or is false (in the case of declaratives) or otherwise appropriately felicitous (in the case of other speech acts)—is provided by the fit with the world. And secondly, that the appropriate explanation of human actions of all sorts of types is due to the meaning^s that an agent uses to *conceptualize* the world: how its beliefs and desires lead to actions that are in accordance with those beliefs and desires, and not to how the beliefs and desires match up with the way the world (really) is.

So the upshot of this casual survey of work from linguists, philosophers of language, and cognitive psychologists seems to be twofold. First, that formal semanticists and philosophers of language make theories that have “meaning” be a feature of how language “fits the world”—usually some sort of intensionalized, possible-world-oriented truth-theory, or (in other hands) an information-theoretic dynamic framework. In doing this they have no obvious place to attach this theory to “meaning” as some feature about how a speaker-hearer dynamic can possibly work. This seems to be a very undesirable state of affairs for a theory of meaning.

On the second hand, the cognitive linguists and the cognitive psychologists seem to have left out the way language is intended to “describe the world”, as they emphasize instead the way that an individual’s mental concepts work in classifying other aspects of their mental life. The attempt to make it be grounded, embodied, situated, or enactive appear to lack convincing detail. In particular, it is not so clear that any expedient from this general point of view *can* succeed, at least according to the Objectivists, since it continues to employ mental “information” (sensations, beliefs about the speech situation, and impressions about ones intentions). It will therefore always be “within the mental world” and not touch “reality”.

Only a two-tier theory can do both jobs.

Acknowledgements Thanks to Greg Carlson for discussions and for comments on an earlier draft. Thanks also to the editors of this volume, James Hampton and Yoad Winter, for their searching critiques of a slightly later draft. Their comments revealed numerous shortcomings in that version, but unfortunately, I have not been able to respond to them all. However, I hope to have ameliorated at least some of their concerns! Some of the alterations came about by discussions with my polymath colleague, Allen Hazen.

References

- Alxatib, S., & Pelletier, F. J. (2011). The psychology of vagueness: Borderline cases and contradictions. *Mind and Language*, 26, 287–326.
- Armstrong, S., Gleitman, L., & Gleitman, H. (1983). What some concepts might not be. *Cognition*, 13, 263–308.
- Asher, N. (2011). *Lexical meaning in context: A web of words*. Cambridge: Cambridge UP.
- Atkins, S., & Rundell, M. (2008). *The Oxford guide to practical lexicography*. Oxford: Oxford UP.
- Bar-Hillel, Y., & Eifermann, R. (1970). Who is afraid of disjunctive concepts? A case study in the genesis of pseudo-problems. *Foundations of Language*, 6, 463–472.
- Barsalou, L. (2008). Grounded cognition. *Annual Review of Psychology*, 59, 617–645.

- Barsalou, L. (2010). Grounded cognition: Past, present, and future. *Topics in Cognitive Science*, 2, 716–724.
- Barsalou, L. (2017). Issues for psychologically plausible theories of conceptual combination. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and concepts in linguistics and psychology* (pp. 223–244). London: Springer.
- Boodin, J. (1939). *The social mind: Foundations of social philosophy*. New York: Macmillan Co.
- Bourne, L. (1982). Typicality effects in logically defined categories. *Memory and Cognition*, 10, 3–9.
- Brandom, R. (1994). *Making it explicit: Reasoning, representing, and discursive commitment*. Cambridge, MA: Harvard UP.
- Brandom, R. (2000). *Articulating reasons: An introduction to inferentialism*. Cambridge, MA: Harvard UP.
- Burge, T. (1979). Individualism and the mental. *Midwest Studies in Philosophy*, 4, 73–121.
- Burling, R. (2005). *The talking ape: How language evolved*. Oxford: Oxford University Press.
- Bybee, J. (1998). A functionalist approach to grammar and its evolution. *Evolution of Communication*, 2, 249–278.
- Bybee, J., & McClelland, J. L. (2005). Alternatives to the combinatorial paradigm of linguistic theory based on domain general principles of human cognition. *The Linguistic Review*, 22, 381–410.
- Carlson, G. (2010). Generics and concepts. In F. Pelletier (Ed.), *Kinds, things, and stuff* (pp. 16–35). New York: Oxford UP.
- Casati, R., & Varzi, A. (2015). Events. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/win2015/entries/events/>.
- Cheney, D. L., & Seyfarth, R. M. (2005). Constraints and preadaptations in the earliest stages of language evolution. *The Linguistic Review*, 22, 135–159.
- Chierchia, G., & McConnell-Ginet, S. (2000). *Meaning and grammar: An introduction to semantics* (2nd ed.). Cambridge, MA: MIT Press.
- Chomsky, N. (1977). *Essays on form and interpretation*. New York, NY: North-Holland.
- Chomsky, N. (1992). Explaining language use. *Philosophical Topics*, 20, 205–231.
- Chomsky, N. (1995). Language and nature. *Mind*, 104, 1–61.
- Conant, M., & Trabasso, T. (1964). Conjunctive and disjunctive concept formation under equal-information conditions. *Journal of Experimental Psychology*, 67, 250–255.
- Connell, L., & Lynott, D. (2014). Principles of representation: Why you can't represent the same concept twice. *Topics in Cognitive Science*, 6, 390–406.
- Croft, W. (2001). *Radical construction grammar*. Oxford: Oxford UP.
- Dirven, R., & Fried, V. (Eds.). (1987). *Functionalism in linguistics*. Amsterdam: John Benjamins.
- Dummett, M. (1991). *The logical basis of metaphysics*. London: Duckworth.
- Evans, V. (2009). *How words mean: Lexical concepts, cognitive models, and meaning construction*. Oxford: Oxford UP.
- Evans, V. (2015). *The crucible of language: How language and mind create meaning*. Cambridge, UK: Cambridge UP.
- Fodor, J. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.
- Fodor, J. (1980). Methodological solipsism considered as a research strategy in cognitive science. *Behavioral and Brain Sciences*, 3, 63–73.
- Fodor, J. (1983). *The modularity of mind: An essay on faculty psychology*. Cambridge, MA: MIT Press.
- Fodor, J. (1998). *Concepts: where cognitive science went wrong*. Cambridge, MA: MIT Press.
- Fodor, J. (2000). *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge, MA: MIT Press.
- Fodor, J. (2001). Language, thought, and compositionality. *Mind and Language*, 16, 1–15.
- Fodor, J., & Lepore, E. (1991). Why meaning (probably) isn't conceptual role. *Mind and Language*, 6, 328–343.
- Fodor, J., & Pylyshyn, Z. (2015). *Concepts without meaning*. Cambridge, MA: MIT Press.

- Francez, N. (2014). Proof-theoretic semantics for contextual domain restriction. *Journal of Language Modelling*, 2, 249–283.
- Francez, N. (2015). *Proof-theoretic semantics*. London: College Publications.
- Francez, N., & Ben-Avi, G. (2015). A proof-theoretic reconstruction of generalized quantifiers. *Journal of Semantics*, 32, 313–371.
- Francez, N., & Dyckhoff, R. (2010). Proof-theoretic semantics for a natural language fragment. *Linguistics and Philosophy*, 33, 447–477.
- Francez, N., Dyckhoff, R., & Ben-Avi, G. (2010). Proof-theoretic semantics for a sub-sentential phrases. *Studia Logica*, 94, 381–401.
- Frege, G. (1891, January). Funktion und Begriff. Vortrag, gehalten in der Sitzung der Jenischen Gesellschaft für Medizin und Naturwissenschaft. Translated as “Function and Concept”.
- Frege, G. (1892a). Über Begriff und Gegenstand. *Vierteljahresschrift für wissenschaftliche Philosophie* 16, 192–205. Translated as “On Concept and Object”.
- Frege, G. (1892b). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100, 25–60. Translated as “On Sense and Reference” (or “On Sense and Meaning”, by different translators).
- Fried, M., & Östman, J.-O. (Eds.). (2004). *Construction grammar in a cross-language perspective*. Amsterdam: John Benjamins.
- Furth, H. (1963). Classification transfer with disjunctive concepts as a function of verbal training and set. *Journal of Psychology*, 55, 477–485.
- Gaeraerts, D., & Cuyckens, H. (2007). Introducing cognitive linguistics. In D. Gaeraerts & H. Cuyckens (Eds.), *Oxford handbook of cognitive linguistics* (pp. 1–21). Oxford: Oxford UP.
- Gagné, C. (2002). Lexical and relational influences on the processing of novel compounds. *Brain and Language*, 81, 723–735.
- Gagné, C., & Spalding, T. (2006). Conceptual combination: Implications for the mental lexicon. In G. Libben & G. Jarema (Eds.), *The representation and processing of compound words* (pp. 145–168). New York: Oxford UP.
- Gagné, C., Spalding, T., & Kostelecky, M. (2016). Conceptual combination, property inclusion, and the Aristotelian-Thomistic view of concepts. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and Concepts in Linguistics and Psychology* (pp. 245–268). London: Springer.
- Goddard, C. (2002). The search for the shared semantic core of all languages. In *Goddard and Wierzbicka* (Vol 1, pp. 5–40).
- Goddard, C., & Wierzbicka, A. (1994). *Semantic and lexical universals*. Philadelphia: Benjamins.
- Goddard, C., & Wierzbicka, A. (2002). *Meaning and universal grammar: Theory and empirical findings* (Vols. 1 & 2). Philadelphia: John Benjamins.
- Goldberg, A. (1995). *Constructions: A Construction grammar approach to argument structure*. Chicago: University of Chicago Press.
- Greenhill, S. J., Atkinson, Q. D., Meade, A., & Gray, R. D. (2010). The shape and tempo of language evolution. *Proceedings of the Royal Society, B*, 277, 2443–2450.
- Grice, P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard UP. Contains Grice’s articles on language, and his later thoughts about his earlier works.
- Halliday, M. (1978). *Language as social semiotic: The social interpretation of language and meaning*. University Park, MD: University of Maryland Press.
- Hampton, J. A. (1979). Polymorphic concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 18, 441–461.
- Hampton, J. A. (1982). A demonstration of intransitivity in natural categories. *Cognition*, 12, 151–164.
- Hampton, J. A. (1988a). Disjunction of natural concepts. *Memory and Cognition*, 16, 579–591.
- Hampton, J. A. (1988b). Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology*, 14, 12–32.
- Hampton, J. A. (1997). Conceptual combination: Conjunction and negation of natural concepts. *Memory and Cognition*, 25, 888–909.

- Hampton, J. A. (2007). Typicality, graded membership, and vagueness. *Cognitive Science*, 31, 355–384.
- Hampton, J. A. (2011a). Concepts and natural language. In R. Belohlavek & G. Klir (Eds.), *Concepts and fuzzy logic* (pp. 233–258). Cambridge, MA: MIT Press.
- Hampton, J. A. (2011b). Conceptual combinations and fuzzy logic. In R. Belohlavek & G. Klir (Eds.), *Concepts and fuzzy logic* (pp. 209–232). Cambridge, MA: MIT Press.
- Hampton, J. A. (2017). Compositionality and concepts. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and concepts in linguistics and psychology* (pp. 95–122). London: Springer.
- Hampton, J. A., & Jönsson, M. (2012). Typicality and compositionally: The logic of combining vague concepts. In M. Werning, W. Hintzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 385–402). Oxford: Oxford UP.
- Heim, I., & Kratzer, A. (1998). *Semantics in generative grammar*. Oxford: Blackwell.
- Hinzen, W., & Sheehan, M. (2013). *The philosophy of universal grammar*. Oxford: Oxford UP.
- Hurfurd, J. R. (2000). Social transmission favours linguistic generalization. In C. Knight, M. Studdert-Kennedy, & J. R. Hurfurd (Eds.), *The evolutionary emergence of language: Social function and the origins of linguistic form* (pp. 219–230). Cambridge: Cambridge University Press.
- Johnson, M. (2014). Compositionality. Internet Encyclopedia of Philosophy. <http://www.iep.utm.edu/composit>.
- Jurafsky, D., & Martin, J. (2008). *Speech and language processing* (2nd ed.). Upper Saddle River, NJ: Prentice-Hall.
- Kaplan, D. (1989). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–566). Oxford: Oxford UP.
- Kay, P. (1995). Construction grammar. In J. Verschueren, J.-O. Östman, & J. Blommaert (Eds.), *Handbook of pragmatics* (pp. 171–177). Amsterdam: John Benjamins.
- Komatsu, L. (1992). Recent views of conceptual structure. *Psychological Bulletin*, 112, 500–526.
- Krifka, M., Pelletier, F. J., Carlson, G., ter Meulen, A., Chierchia, G., & Link, G. (1995). Genericity: An introduction. In G. Carlson & F. J. Pelletier (Eds.), *The generic book* (pp. 1–124). Chicago: University of Chicago Press.
- Kripke, S. (1980). *Naming and necessity*. Cambridge MA: Harvard UP.
- Labov, W. (1966). *The social stratification of English in New York City*. Washington, DC: Center for Applied Linguistics. 2nd edition Cambridge University Press, 2006.
- Labov, W. (1975). Empirical foundations of linguistic theory. In R. Austerlitz (Ed.), *The scope of American linguistics* (pp. 77–133). Lisse: Peter de Ridder.
- Langacker, R. (1987). *Foundations of cognitive grammar. Vol. 1: Theoretical prerequisites*. Stanford: Stanford UP. Chapter 12: Composition.
- Langacker, R. (1990). *Concept, image, and symbol: The cognitive basis of grammar*. Berlin: Mouton de Gruyter.
- Langacker, R. (2008). *Cognitive grammar: A basic introduction*. Oxford: Oxford UP.
- Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In S. Laurence & E. Margolis (Eds.), *Concepts: Core readings* (pp. 3–82). Cambridge, MA: MIT Press.
- Leben, D. (2015). Neoclassical concepts. *Mind and Language*, 30, 44–69.
- Lewis, D. (1970). General semantics. *Synthese*, 22, 18–67.
- Li, C. (2002). Missing links, issues and hypotheses in the evolutionary origin of language. In T. Givón & B. Malle (Eds.), *The evolution of language out of pre-language* (pp. 83–106). Amsterdam: John Benjamins.
- Locke, J. (1690). *Essay concerning human understanding*. Oxford: Clarendon Press. Edited, with forward, by Peter Niddich; published 1975.
- Machery, E. (2009). *Doing without concepts*. Oxford: Oxford UP.
- MacWhinney, B. (2005). The emergence of grammar from perspective taking. In D. Pecher & R. A. Zwaan (Eds.), *Grounding cognition: The role of perception and action in memory, language, and thinking* (pp. 198–223). Cambridge: Cambridge University Press.

- Manning, C., & Schütze, H. (1999). *Foundations of statistical natural language processing*. Cambridge, MA: MIT Press.
- Martin-Löf, P. (1980). *Intuitionistic type theory*. Naples: Bibliopolis.
- McGilvray, J. (1998). Meanings are syntactically individuated and found in the head. *Mind and Language*, 13, 225–280.
- McNally, L., & Boleda, G. (2017). Conceptual vs. referential affordance in concept composition. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and concepts in linguistics and psychology* (pp. 245–268). London: Springer.
- Murphy, G. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Newstead, S., & Griggs, R. (1983). The language and thought of disjunction. In J. S. Evans (Ed.), *Thinking and reasoning: psychological approaches* (pp. 76–106). London: Routledge & Kegan Paul. Reprinted 2013 by Taylor & Francis, in their Psychology Revivals series.
- O’Grady, W. (2008). The emergentist program. *Lingua*, 118, 447–464.
- Osgood, C., Suci, G., & Tannenbaum, P. (1957). *The measurement of meaning*. Urbana, IL: University of Illinois Press.
- Osherson, D., & Smith, E. (1981). On the adequacy of prototype theory as a theory of concepts. *Cognition*, 11, 35–58.
- Östman, J.-O., & Fried, M. (Eds.) (2004). *Construction grammars*. John Benjamins.
- Pagin, P. (2009). Compositionality, understanding, and proofs. *Mind*, 118, 713–737.
- Pagin, P. (2012). Communication and the Complexity of Semantics. In M. Werning, W. Hinzen & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 510–529). Oxford: Oxford UP
- Parsons, T. (1990). *Events in the semantics of English: A study in subatomic semantics*. Cambridge, MA: MIT Press.
- Partee, B. (1979). Semantics—mathematics or psychology? In R. Bäuerle, U. Egli, & A. von Stechow (Eds.), *Semantics from different points of view* (pp. 1–14). Berlin: Springer.
- Pecher, D., & Zwaan, R. A. (Eds.). (2005). *Grounding cognition: The role of perception and action in memory, language, and thinking*. Cambridge: Cambridge University Press.
- Pelletier, F. J. (1998). Thinking of ‘not’. In X. Arrazola, K. Kortá, & F. J. Pelletier (Eds.), *Discourse, interaction, and communication* (pp. 37–47). Dordrecht: Kluwer.
- Pelletier, F. J. (2012). Holism and compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 149–174). Oxford: Oxford UP.
- Pelletier, F. J. (2013). Context, compositionality, and brevity. In L. Goldstein (Ed.), *Brevity* (pp. 178–197). Oxford: Oxford UP.
- Pelletier, F. J. (2016). Semantic compositionality. In M. Aronoff (Ed.), *Oxford research encyclopedia of linguistics*. Oxford: Oxford UP.
- Pelletier, F. J., Elio, R., & Hanson, P. (2008). Is logic all in our heads? From naturalism to psychologism. *Studia Logica*, 88, 3–66.
- Power, C. (1998). ‘Old wives’ tales’: The gossip hypothesis and the reliability of cheap signals. In J. R. Hurford, M. Studdert-Kennedy, & C. Knight (Eds.), *Approaches to the evolution of language* (pp. 111–129). Cambridge: Cambridge University Press.
- Prawitz, D. (2006). Meaning approached via proofs. *Synthese*, 148, 507–524.
- Prinz, J. (2002). *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: MIT Press.
- Putnam, H. (1975). The meaning of ‘meaning’. In K. Gunderson (Ed.), *Mind, language and reality* (pp. 215–271). Minneapolis: University of Minnesota Press.
- Putnam, H. (1981). Brains in a vat. In H. Putnam (Ed.), *Reason, truth and history* (pp. 1–21). Cambridge, UK: Cambridge UP.
- Recanati, F. (2012). Compositionality, flexibility, and context dependence. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 175–191). Oxford: Oxford UP.
- Rohrer, T. (2007). Embodiment and experimentalism. In D. Geeraerts & H. Cuyckens (Eds.), *The Oxford handbook of cognitive linguistics* (pp. 25–47). Oxford: Oxford UP.

- Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, 104, 192–232.
- Sainsbury, M. (2001). Two ways to smoke a cigarette. *Ratio*, 14, 386–406.
- Schank, R. C. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, 3, 532–631.
- Schein, B. (1993). *Plurals and events*. Cambridge, MA: MIT Press.
- Schmid, H.-J. (2007). Entrenchment, salience, and basic levels. In D. Geeraerts & H. Cuyckens (Eds.), *The Oxford handbook of cognitive linguistics* (pp. 117–138). Oxford: Oxford UP.
- Scholz, B. C., Pelletier, F. J., & Pullum, G. K. (2016). Philosophy of linguistics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/win2016/entries/linguistics/>.
- Schroeder-Heister, P. (2016). Proof-theoretic semantics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/win2016/entries/proof-theoretic-semantics/>.
- Schubert, L. (2015). Computational linguistics. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/spr2015/entries/computational-linguistics/>.
- Seggie, J. (1969). Levels of learning involved in conjunctive and disjunctive concepts. *Australian Journal of Psychology*, 21, 325–333.
- Shapiro, L. (2014). *The Routledge handbook of embodied cognition*. Oxford: Routledge.
- Sinha, C. (2007). Cognitive linguistics, psychology, and cognitive science. In D. Geeraerts & H. Cuyckens (Eds.), *The Oxford handbook of cognitive linguistics* (pp. 1266–1294). Oxford: Oxford UP.
- Skinner, B. (1957). *Verbal behavior*. Acton, MA: Copley Publishing.
- Smith, E., & Medin, D. (1981). *Categories and concepts*. Cambridge, MA: Harvard UP.
- Smith, E., & Osherson, D. (1984). Conceptual combination with prototype concepts. *Cognitive Science*, 8, 337–361.
- Smith, E., Osherson, D., Rips, L., & Keane, M. (1988). Combining prototypes: A selective modification model. *Cognitive Science*, 12, 485–527.
- Smith, E., Shone, E., & Rips, L. (1974). Structure and process in semantic memory: A featural model for semantic decisions. *Psychological Review*, 81, 214–241.
- Smuts, J. (1926). *Holism and evolution*. London: MacMillan.
- Spalding, T., & Gagné, C. (2008). CARIN theory reanalysis reanalyzed: A comment on Maguire, Devereux, Costello, and Cater. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 1573–1578.
- Szabó, Z. (2012). The case for compositionality. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 64–80). Oxford: Oxford UP.
- Taylor, J. (2007). Cognitive linguistics and autonomous linguistics. In D. Geeraerts & H. Cuyckens (Eds.), *The Oxford handbook of cognitive linguistics* (pp. 566–588). Oxford: Oxford UP.
- Tomasello, M. (1998). Introduction. In M. Tomasello (Ed.), *The new psychology of language: Cognitive and functional approaches to language structure*. Mahwah, NJ: Lawrence Erlbaum.
- Trabasso, T., Rollins, H., & Shaughnessy, E. (1971). Storage and verification stages in processing concepts. *Cognitive Psychology*, 2, 239–289.
- Turner, M. (1991). *Reading minds: The study of English in the age of cognitive science*. Princeton, NJ: Princeton UP.
- Tye, M. (2009). A new look at the speckled hen. *Analysis*, 69, 258–263.
- Van Valin, R. (2003). Functional linguistics. In M. Aronoff & J. Rees-Miller (Eds.), *The handbook of linguistics* (pp. 319–366). Oxford: Blackwell.
- von Humboldt, W. (2000). Character of languages; poetry and prose. In M. Losonsky (Ed.), *Humboldt on language* (pp. 148–181). Cambridge, UK: Cambridge UP.
- von Savigny, E. (1988). *The social foundations of meaning*. Berlin: Springer.
- Westerståhl, D. (2012). Compositionality in Kaplan-style semantics. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford handbook of compositionality* (pp. 192–219). Oxford: Oxford UP.

- Wierzbicka, A. (1996). *Semantics: Primes and universals*. Oxford: Oxford UP.
- Wildgen, W. (2008). Sketch of an evolutionary grammar based on comparative biolinguistics. In L. Röska-Hardy & E. Neumann-Held (Eds.), *Learning from Animals? Examining the nature of human uniqueness* (pp. 45–59). London: Psychology Press.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9, 219–248.
- Wilson, R. A., & Foglia, L. (2016). Embodied cognition. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <http://plato.stanford.edu/archives/win2016/entries/embodied-cognition/>.
- Winter, Y. (2017). Critical typicality: Truth judgements and compositionality with plurals and other gradable concepts. In J. A. Hampton & Y. Winter (Eds.), *Compositionality and concepts in linguistics and psychology* (pp. 163–190). London: Springer.
- Wisniewski, E. (1997). When concepts combine. *Psychonomic Bulletin & Review*, 4, 167–183.
- Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: Basil Blackwell. Translated by G.E.M. Anscombe.
- Worden, R. (1998). The evolution of language from social intelligence. In J. R. Hurford, M. Studdert-Kennedy, & C. Knight (Eds.), *Approaches to the evolution of language* (pp. 148–166). Cambridge: Cambridge University Press.

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

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

