

JOŠKO ŽANIĆ \*

## WHY CONCEPTUAL SEMANTICS IS PREFERABLE TO THE TRUTH-CONDITIONAL APPROACH IN THE STUDY OF MEANING

**SUMMARY:** In this paper, I argue that Conceptual Semantics, a framework in the study of meaning developed by Ray Jackendoff and Steven Pinker, is preferable to the dominant truth-conditional approach. The structure of the paper is as follows. After a brief introduction, in Section 2, I present briefly the dominant approach. In Section 3, I outline what I take to be the main challenge to it, namely radical contextualism. In Section 4, I present the conceptualist approach and show why I think it successfully avoids this challenge, while offering satisfying explanations. I conclude with the point that Conceptual Semantics is also more naturalistic than Truth-Conditional Semantics, which I take to be a desirable goal.

**KEYWORDS:** conceptual semantics, truth-conditional semantics, meaning, contextualism, naturalism.

### 1. Introduction

In this paper I argue that Conceptual Semantics, a framework in the study of meaning developed by Ray Jackendoff and Steven Pinker, is preferable to the dominant truth-conditional approach. The structure of the paper is as follows. In Section 2, I present briefly the dominant approach. In Section 3 I outline what I

---

\* Independent researcher. E-mail: josko\_zanic@yahoo.com. ORCID: 0000-0002-0145-1255.

take to be the main challenge to it. In Section 4, I present the conceptualist approach and show why I think it successfully avoids this challenge, while offering satisfying explanations. In this final Section, I attempt to show that Conceptual Semantics is also more naturalistic than Truth-Conditional Semantics, which I take to be a desirable goal. So, let us begin.<sup>1</sup>

## 2. The Truth-Conditional Project

Truth-Conditional Semantics has been the dominant approach in the study of meaning for many decades now.<sup>2,3</sup> The basic requirement it imposes on itself, in line with a constraint introduced by Davidson (1984), is that the theory of meaning has to provide an explication of meaning for every possible sentence in a language. The basic idea is that it does this by giving the truth conditions of the sentence, namely the circumstances under which the sentence would be true. Sentences represent situations, so it seems reasonable to assume that specifying the situation whose obtaining would make it true gives the meaning of the sentence. Meaning is thus based on truth, on this approach.

The approach comes in two main varieties, the truth-theoretic (cf. the textbook Larson & Segal 1995, more recently Lepore & Ludwig 2007; it is based on the work of Davidson 1984) and the model-theoretic (cf. the textbook Chierchia & McConnell-Ginet 2000; the approach is based on the work of Montague 1974). For a sentence like ‘Snow is white’, the first, truth-theoretic variety issues in meaning-specifying theorems such as the following (in line with Tarski’s 1944 conception of truth):

- (1) ‘Snow is white’ is true iff snow is white,

The second, model-theoretic variety offers theorems such as the following:

- (2)  $[\text{Snow is white}]^{M_1} = 1$  iff  $[\text{snow}]^{M_1} \in [\text{white}]^{M_1}$

Whereas the first theorem is more or less self-explanatory, the second says something like the following: the sentence ‘Snow is white’ is true in model  $M_1$  if and only if the denotation of ‘snow’ in the model belongs to the extension of ‘white’ (the set of white things) in the model.

---

<sup>1</sup> A reviewer asked in which perspective I’m claiming CS is preferable. The perspective is semantics, considered as a part of cognitive science.

<sup>2</sup> The most recent, albeit popular, introduction is Schlenker (2022).

<sup>3</sup> There exists, however, another very lively approach which arose as a reaction to Truth-Conditional Semantics, namely Cognitive Linguistics (cf. Croft & Cruse 2004). The founding fathers are Lakoff, Langacker, Fauconnier, Fillmore, and Talmy. The problem I have with this approach is that it favors embodied meanings, which I feel adds nothing to the enterprise, because the body still has to be represented by the brain. And also, Jackendoff’s approach is partly formalized, a desirable goal, I take it.

The theorem in (1) could be accused of being trivial; however, in employing a distinction between objectlanguage and metalanguage, inherited from Tarski, it is not really so. The metalanguage, in which the meaning of the sentence is given, could be changed, e. g. to French, and this would dispel the air of triviality. The theorem in (2) specifies meaning relative to a model, and it has sometimes been asserted that this makes the truth-conditions relative, which is taken to be a drawback. However, it has been replied on behalf of this version of the approach (e. g. Glanzberg 2014, p. 266) that the reference to models can be eliminated. In any case, what the two versions of the approach offer is not really very different, especially for the purposes of this paper.

Both versions, crucially, specify meanings of sentences *compositionally*: they start with meaning-axioms for individual words, and then show how the meaning, i.e. the truth-conditions, of the sentence follow from the meanings of the words, and their mode of combination ('Snow is white' means what it does because 'snow' means what it does, 'white' means what it does, and they are combined in a predicative fashion). Since the sentences of a language are potentially infinite, this is, of course, the only way it could be done.

A parameter of variation within this approach, not exactly aligned with the difference between the two versions, is whether a variety of it is meant to be psychologically realistic. This is expressly the goal of Larson & Segal (1995), who aim to model the semantic knowledge of speakers. It needn't, however, be a goal of this approach, which can just be content with "getting the entailment relations right" (G. Chierchia, pers. comm.).

Truth-Conditional Semantics (henceforth TCS), of either the truth-theoretic or the model-theoretic variety, has been subjected to many criticisms. Lycan (2000: chaps. 9 and 10) summarizes some of the main ones:

1. TCS applies only to 'descriptive, fact-stating language'.
2. The theorems it delivers (especially for sentences more complex than the one in (1) and (2)) are too complex and technical to be empirically testable.
3. There are "ferocious" technical problems that emerge once one considers sentences with deictical elements in them.
4. TCS has trouble distinguishing co-extensive terms that are not mutually synonymous.
5. Some sentences behave non-truth-functionally.
6. The central notion employed in TCS, truth, has not been sufficiently clarified.

Of course, replies are available to each of these objections, many of which Lycan provides himself. Accounts of non-descriptive language have been proposed within this tradition (cf. Lepore & Ludwig 2007: ch. 12). Deixis can be handled in a Kaplanian manner (cf. Chierchia and McConnell-Ginet 2000: ch. 6). The model-theoretic variety handles co-extensive terms well, because it assigns

to them different intensions (functions from worlds to extensions). And so on. I assume therefore that the listed objections can be replied to, although it remains an open question whether these replies are fully adequate. In the next Section, I will focus on what I take to be the crucial objection to TCS. Before that, let me briefly present Lycan's (2010) "direct" argument for TCS (of the model-theoretic variety).<sup>4</sup>

Lycan offers the following argument:

- (1A) A sentence's meaning taken together with a totality of fact determines the sentence's truth value.
- (2A) Therefore, a sentence meaning is at least a function from possible worlds to truth values.
- (3A) Such a function is a truth condition.
- (4A) Therefore, a sentence meaning is at least a truth condition.

Now, as Lycan himself recognizes (cf. fn. 18), the step from (1A) to (2A) is "not strictly valid", but he offers it as the "obvious explanation" of (1A). Whether we can accept a not strictly valid step in a deductive argument is questionable. However, I will point out something else: the "totality of fact" that determines the truth value of a sentence has not been adequately construed by TCS – it usually excludes the intents and purposes of speakers. This is the topic of the next Section.

### 3. The Challenge

The challenge I have in mind, what I take to be the main challenge to TCS, has been labeled 'contextualism' or 'radical contextualism'. I will rehearse here two of the most famous examples that are adduced by contextualists (many other examples are provided by Recanati 2004, 2010). Here is the first example (from Travis 1997, cf. also his 2008):

- (3) The leaves are green.

Travis (1997, p. 89) elaborates this way:

Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those'. But now Pia speaks falsehood.

---

<sup>4</sup> It should be pointed out that what Lycan sees as the main opponent of TCS is use-theoretic, or inferentialist, semantics. He doesn't address the conceptualist approach which I defend in this article.

The other example I take from Searle (1980). Here are the relevant sentences:

- (4) He cut the grass.
- (5) He cut the cake.

Searle (pp. 222-3) elaborates:

The sort of thing that constitutes cutting the grass is quite different from, e. g., the sort of thing that constitutes cutting a cake. One way to see this is to imagine what constitutes obeying the order to cut something. If someone tells me to cut the grass and I rush out and stab it with a knife, or if I am ordered to cut the cake and I run over it with a lawnmower, in each case I will have failed to obey the order.

Lest you think that the meaning of ‘cut’ in (4) and (5) is fixed by its syntactic object, Searle (pp. 224-5) continues this way:

It is easy to imagine circumstances in which ‘cut’ in ‘cut the grass’ would have the same interpretation it has in ‘cut the cake’, even though none of the semantic contents of the words has changed. Suppose you and I ran a sod farm where we sell strips of grass turf to people who want a lawn in a hurry [...] Suppose I say to you, ‘Cut half an acre of grass for this customer’; I might mean not that you should *mow* it, but that you should slice it into strips as you could cut a cake or a loaf of bread.

Here is Travis (1997, p. 87) on the moral of these examples:

[...] it is intrinsically part of what expressions of (say) English mean that any English (or whatever) sentence may, on one speaking of it or another, have any of indefinitely many different truth conditions, and that any English (or whatever) expression may, meaning what it does, make any of many different contributions to truth conditions of wholes in which it figures as a part.

The contextualist examples, called by their opponents Cappelen & Lepore (2005) ‘context-shifting arguments’, show therefore that the truth conditions of a sentence are fixed only in a specific context, depending often on the intents and purposes of speakers. ‘Green’ can make a different contribution to the truth conditions of (3), depending on whether we are interested in surface appearance or in botany. So the truth conditions of (3) must be given either as:

- (6) ‘The leaves are green’ is true iff the leaves appear green,
- or
- (7) ‘The leaves are green’ is true iff the leaves are green, botanically speaking.

And, of course, the context could be further varied to produce further variation in truth conditions. The same goes for the ‘cut’ example.

As opposed to Lycan's (4A) claim above therefore (and to the TCS program in general), meaning (of the standing, i. e. context-independent, variety) is *not* a truth condition, or, put differently, it doesn't determine a truth condition. I take this to mean, and this is the main point of this Section, that standing meaning is always to a certain extent *schematic*, and it does not suffice to determine truth conditions.

Now, the anti-contextualists, such as Cappelen & Lepore or, more recently, Borg (2012) have replied to contextualist claims, and attempted to defuse them. It is not my goal to review this whole debate here, but to agree with those authors, such as, recently, Picazo (2022), who take the position that the claims of contextualists have not been successfully defused. Here is just an example of what I mean: It could be argued that the variation in truth conditions in the above examples is foreseen by hidden material attached to 'green' or 'cut'. The idea would be that each of these expressions carries a hidden parameter, namely for ways of being green or ways of cutting, and the context just chooses a value on this parameter. Now, in all likelihood, this is not the case. As already hinted, the context could be varied further, to provide new interpretations for the relevant sentences, and it is extremely likely that not all of these variations could be foreseen in the lexical meaning of the words we are considering. The "hidden-parameter-approach" cannot, therefore, tame the variability that the contextualist argues comes with language. In any case, contextualism seems to be in a better position dialectically: it just needs *one* example where the truth conditions of a sentence cannot be deduced algorithmically from the lexical meanings of the words and the way they are combined, whereas TCS needs to refute *every* case the contextualists might come up with.

If contextualism is true, and taking as the main goal of TCS to provide standing truth conditions of sentences in a language, I conclude that TCS is in trouble. What TCS construes as meaning, namely truth conditions, cannot be provided out of a specific context. Standing meaning, as TCS sees it (namely as context-independent truth conditions), in many cases doesn't exist. This is where Conceptual Semantics comes in.

#### 4.1. A First Argument for the Conceptualist Alternative

Conceptual Semantics (Jackendoff 1983, 1990, 2002, more recently 2011, 2021; Pinker 1989, 2007) attempts to find central, recombinable conceptual elements that make up the meaning of words and sentences. It is in terms of these that the standing meaning of expressions is given in this theory. I will rehearse here an example from Pinker (2007, pp. 102-7), based partly on the work of Beth Levin. I will modify the example, because I think it contains an omission.<sup>5</sup> The point will be that, although meaning is too schematic to be capturable in terms of (standing) truth-conditions and contributions to truth-conditions, it *is*

---

<sup>5</sup> It leaves out the 'impact' component, see below.

capturable in terms of these recombinable conceptual elements, often called meaning components, which (mostly) don't vary with the context.

Take the verb 'cut' that we encountered in the previous Section. Pinker observes that there are certain syntactic environments which it can and cannot appear in, as opposed to some other verbs. Compare the following alternations (the unacceptable sentences are marked with an asterisk):

#### The Conative Alternation

- (7) Mabel cut the rope. / Mabel cut at the rope.
- (8) Vince hit the dog. / Vince hit at the dog.
- (9) Nancy touched the cat. / \*Nancy touched at the cat.
- (10) John broke the window. / \* John broke at the window.

The explanation of this is that the construction specifies which conceptual components the verb appearing in it can and cannot have; this makes for acceptable or unacceptable sentences, depending on the verb we insert. The conative alternation seems to be of the following kind – it allows the conceptual components [motion] and [impact], but disallows the verbs which lack any one or both of these components. Since 'cut' and 'hit' arguably contain these components in their lexical representations, they can appear in this construction. Since 'touch' lacks [impact] it cannot appear in the alternation, and neither can 'break' because it lacks [motion].

To jump ahead a bit, and introduce two more semantic components, namely [contact] and [effect], the semantic representations of our four verbs can be given in the following manner:

- 'cut': [+motion, +contact, +impact, +effect]
- 'hit': [+motion, +contact, +impact, -effect]
- 'break': [-motion, -contact, -impact, +effect]
- 'touch': [+motion, +contact, -impact, -effect]

This will explain their (un)acceptability in some further alternations:

#### Possessor-Raising

(acceptable component: [contact])

- (11) Sam cut Brian's arm. / Sam cut Brian on the arm.
- (12) Miriam hit the dog's leg. / Miriam hit the dog on the leg.
- (13) Miriam hit the dog's leg. / Miriam hit the dog on the leg.
- (14) John broke Thomas's leg. / \*John broke Thomas on the leg. (unacceptable because 'break' lacks [contact]).

## Contact Locative

(requires [impact], disallows [effect])

- (15) I hit the ball with the bat. / I hit the bat against the wall.  
 (16) I cut the rope with the knife. / \*I cut the knife against the rope (unacceptable because 'cut' has [effect] as component).  
 (17) I broke the glass with the hammer. / \*I broke the hammer against the glass (unacceptable for the same reason).  
 (18) I touched the cat with my hand. / \*I touched my hand against the cat (unacceptable because 'touch' lacks [impact]).

Here is one more construction:

## The Anticausative

(requires [effect], disallows [motion] and [contact])

- (19) John broke the glass. / The glass broke.  
 (20) John cut the rope. / \*The rope cut.  
 (21) Mary hit John. / \*John hit. (intended meaning: 'was hit').  
 (22) Mary touched John. / \*John touched. (intended meaning: 'was touched').

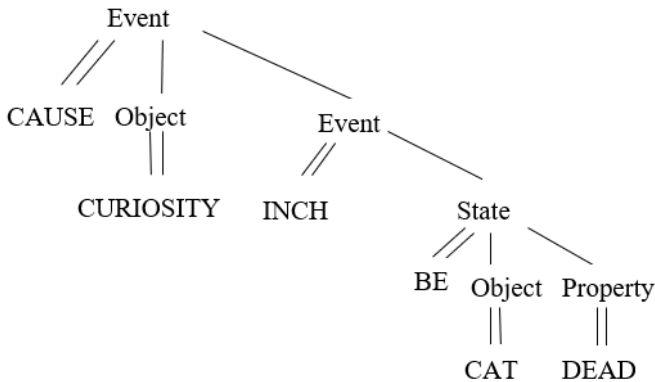
This is how Conceptual Semantics proceeds, therefore. It attempts to extract the basic meaning components of words by looking at syntactic environments they can and cannot appear in. It specifies (part of) the standing meaning of 'cut' as [+motion, +contact, +impact, +effect]. This is something that doesn't vary in context. The meaning effects observed by Searle for 'cut' are modulations (as Recanati would put it) that come from context, but the meaning components stay the same irrespective of context (not always though – see below).

Conceptual Semantics sees meaning (which it identifies with conceptualization) as a hierarchical combinatorial structure, an algebraic recursive system. It specifies the meaning of sentences (out of context) in terms of basic conceptual components (functions) and the arguments they take. The most frequent of these functions are: (i) the core functions: BE(X, Y), GO(X,Y), STAY(X,Y), etc.; (ii) the aspectual functions: PERF(X) (for 'perfective') and INCH(X) (for 'inchoative'); and the causative functions: CAUSE(X,Y), LET(X,Y), HELP(X,Y), etc. If we take a sentence such as

- (23) Curiosity killed the cat.

Conceptual Semantics specifies its meaning in the following fashion:

(24)



This structure can be explained in the following way. It contains conceptual functions, written in capitals and marked with double lines, which take arguments. Arguments are the basic ontological categories: Object, Property, State, Event.<sup>6</sup> The structure is best explained by going from the bottom. At the bottom we have an object and a property, namely a cat and being dead, which are the arguments of the function BE, yielding a state: the state of the cat being dead. When the conceptual function INCH (=come to be) is applied to this state, we get an event: one in which the cat came to be dead, i. e. died. The uppermost function, CAUSE, takes an object (curiosity, hypothesized to be conceptualized as an abstract object) and the whole aforementioned event as arguments, to deliver the event which is the meaning of the sentence (an event in CS is a conceptualized change of state). Ultimately, the meaning of the sentence is given as ‘Curiosity caused the cat to come to be dead’. Thereby its meaning is decomposed into an algebra of more basic, recombinable, units. The conceptual function CAUSE will also appear in the meanings of e. g. ‘break’ and, more controversially, ‘cut’. We see also that what corresponds to a sentence at a basic level is an event (or a state), not a truth value. Truth, for Conceptual Semantics, is an evaluation of a sentence with respect to (an understanding of) a situation, but it is not at the basis of meaning.<sup>7</sup> Conceptual semantics is a thoroughly mentalistic theory, which does not base meaning on truth and reference (although it has much to say about these, cf. Jackendoff 2002: ch. 10, see also his 2011).

Conceptual Semantics seeks to discover, therefore, the basic units of thought out of which the meanings of words and sentences are built. The standing mean-

---

<sup>6</sup> CURIOSITY, CAT and DEAD are also written in capitals, because they can be treated as zero-argument functions, which directly deliver the corresponding object or property as value.

<sup>7</sup> A similar project is pursued by Pietroski (2018), but I will not compare the two approaches here, since this comparison would require a whole another paper.

ings of words and sentences are given in terms of an algebraic combination of these basic units. Conceptual Semantics can also handle contextual enrichment, however. It sees this as operations on the basic conceptual structure. On this see Jackendoff (2002: ch. 12.2), who says that ‘the understanding of sentences is a rich interaction between grammar, independent well-formedness conditions on conceptual structure, and the construal of context’ (p. 388). There are situations, therefore, where a conceptual transformation is induced by context, but corresponds to nothing in syntactic or phonological structure (e. g. ‘She jumped until the bell rang’, where ‘jump’ has to be construed as iterative). I believe this is also problematic for TCS.

Fodor (cf. his 1998) would object to the analysis of (23) as (24) on the grounds that I can cause somebody to die without killing them. For instance, if I pay A to kill B and A succeeds, then I have caused B to die, but not killed them. On behalf of Conceptual Semantics, I would reply here by invoking Pustejovsky’s (1995) *qualia structure*, which Jackendoff incorporates into the theory (cf. his 2002: ch. 11.9). Qualia structure is the breaking down of a word’s meaning components into classes: formal, constitutive, agentive, and telic (following Aristotle). The analysis in (24) gives us the formal quale of ‘kill’, the organization of its functions and arguments. However, there are other qualia in the meaning of this verb: for example, the constitutive quale might be specified as [directly], and the agentive quale as [intentionally]. When we add these qualia to the analysis, claiming that ‘kill’ means ‘cause to die directly and intentionally’, we get a fuller analysis, which seems to take care of Fodor’s objection. Of course, this is only the prototypical meaning of the verb (Conceptual Semantics accepts the results of psychological research on prototypicality, cf. Murphy 2002: ch. 2). We can get figurative meanings by eliminating the constitutive or the agentive quale: if we eliminate the former, we get figurative ‘killing’ as in the counterexample (we can still say ‘I killed B’, only figuratively), and if we eliminate the latter, we get the meaning occurring in the sentence ‘The doctors killed the patient’.

The components identified for ‘cut’ and the other verbs above can be seen as constitutive qualia, specifying the structural composition of the meanings of these verbs. The formal quale of ‘cut’ might be given as [effect a linear separation in an object by means of an edged instrument] (cf. Recanati 2010, p. 40).

One needn’t take Conceptual Semantics to be completely inimical to TCS. Indeed, if, as Napoletano (2019) claims, specifying truth conditions is not always crucial to TCS (which, admittedly, sounds counterintuitive), Conceptual Semantics can integrate many of the results of TCS (by reformulating them in its own terminology). It is only when TCS insists on giving context-independent truth conditions, that I think that the conceptualist approach is preferable, because it focuses on standing, schematic meanings. And, in accepting prototypicality, this approach can also integrate many insights of Cognitive Linguistics (mentioned in footnote 3 above), which is in many ways based on the prototypicality insight.

Particularly its rich research on polysemy (cf. Croft & Cruse 2004: ch. 5) can be integrated into Conceptual Semantics (but less easily into TCS).

Glanzberg (2014) claims that TCS gets disquotational and manifests a lack of explanatory power when it bumps up against concepts. So, for instance, informative analyses, in mathematical terms (employing scales), can be given of gradable adjectives, but when we get to the concept that the scale is a scale of, e. g. TALL, we can only go disquotational, and say that it expresses a dimension of tallness. However, as seen above, mathematical (or quasi-mathematical) analyses *can* be given of concepts, so I think that the study of concepts can be fully integrated into semantics. This is the main business of Conceptual Semantics.

A reviewer asks how exactly Conceptual Semantics deals with contextualist examples. I will answer the question by looking at the “cut” examples, since “cut” has been mentioned already several times in the paper. As we’ve seen from Searle’s discussion, “cut”, at the grain level at which it is analyzed in TCS, can mean several different things depending on the context, and these meaning are not easily predictable. This generates the variation in truth-conditions and ultimately the truth value in the situation. What CS does is to *change the grain level*, i.e. go deeper, using Pustejovsky’s qualia structure. The formal qualia of “cut” would be something like [effect a linear separation in an object by means of an edged instrument]. The constitutive qualia would be the following: [+motion, +contact, +impact, +effect]. Surely there are more qualia to find, especially the agentive and telic qualia, but, when found, they do not change with the context. The point is that, if we change the grain level of analysis, contextual effects mostly disappear and we are allowed to go deeper and deeper into the meaning structures of the words and sentences. CS does, however, allow for context effects. E.g. in a sentence “She jumped till the morning”, an iterative sense is coerced into the meaning of “jump”. Now, CS gives a good explanation of this: what happens here is that, given that “jump” is already conceptualized as a point-event, possessing an inherent border, the constraints on conceptual structure don’t allow it to receive another border introduced by “till the morning”. So, at a higher level at analysis, there can be contextual effects, but these are well explained as operations on the underlying conceptual structure, which is stable.

#### 4.2. A Second Argument for the Conceptualist Alternative

Here is what I take to be another advantage of Conceptual Semantics. The problem I have in mind for TCS is discussed by Calzavarini (2021), although it has also been identified by Lycan (2000: pp. 156-8). As mentioned in Section 2, some versions of TCS aim to be psychologically realistic. Calzavarini takes this to apply to TCS in general, and formulates the requirement thus: ‘A theory of meaning should be conceived as a model of a normal speaker’s semantic competence’ (p. 43). The problem is that, if one accepts certain assumptions, what this theory takes to be the general format of meaning-specification, namely truth

conditions, are world-involving in a way that creates a tension with the aim of modeling the psychology of speakers.

The challenge comes from externalism (originated by Putnam 1975). On the currently dominant, externalist picture, the semantic values (the reference) of our words are not fully determined by what's in the head, but also by the way the environment is. If in some world, otherwise identical to ours, the liquid that flows in rivers and lakes, falls from the sky, is drinkable etc., doesn't have the usual chemical composition  $H_2O$ , but rather XYZ, then, in that world, when speakers utter 'water', even if they are not aware of the chemical composition of the relevant liquid in their world, they are referring to XYZ, and not to  $H_2O$ . Their words refer to what they were in causal contact with, irrespective of the fact that their 'stereotype' of water (colorless, odorless, potable liquid) is the same as that of speakers on Earth.

This has an effect on truth conditions. The truth conditions of 'John drank a glass of water', given for a speaker on Earth, have to be formulated as follows:

(25) John drank a glass of  $H_2O$ .

But for a speaker in the other world, so-called Twin Earth, they have to be formulated thus:

(25) John drank a glass of XYZ.

In this sense, truth conditions are world-involving. However, by assumption, John and Twin-John are not aware of the chemical composition of water. The truth conditional theory therefore cannot successfully model their semantic competence.

Conceptual Semantics, which also aims to model semantic competence, doesn't have this problem, because it gives meanings in terms of conceptual structures, which are in the head. Of course, concepts can also be given a 'wide' construal, claiming that the concept expressed by the word 'water' here and on Twin Earth is not the same concept. However, they *needn't* be given such a construal, and, indeed, they aren't given it in Conceptual Semantics. In this framework, concepts are pieces of mental information structure, located squarely in the head. Such concepts, and not ones individuated widely, are the ones that perform an important function of concepts, namely guiding behavior (on the interest of Conceptual Semantics in connecting to perception and action, see next Section). Given this construal, and having in mind the goal of modeling semantic competence, Conceptual Semantics again seems to come out on top. With that point I conclude this Section.

### 4.3. A Third Argument for the Conceptualist Alternative: Naturalism

I take a theory to be naturalistic if it is well-placed to be integrated with the more basic natural sciences; alternatively, but relatedly, I take it to be naturalistic if it shows (at least to some extent) that the phenomenon it studies is part of the natural order. I also take naturalism to be a desirable goal for theories in the humanities and social sciences, because it leads to an unification of knowledge.

It is my claim in this penultimate section that Conceptual Semantics is more naturalistic than TCS, in both of the senses mentioned above: it is better placed to be integrated with the natural sciences, and it is more successful in connecting meaning to the natural order. I will argue for the two points in turn.

As for integrating semantics with the natural sciences, this integration would presumably have to go through psychology. I claim that Conceptual Semantics is more psychologically realistic than TCS. It has already been pointed out that Conceptual Semantics accepts the prototypicality results from psychology, and these are robust (see again Murphy 2002). It has been shown, for instance, that some instances of a concept are seen by speakers as more typical than others (e. g. an apple is a more typical fruit than a lemon), and this manifests itself in speed of categorization, order of learning in childhood, etc. The same effects are observed for words (the word ‘apple’ will occur to subjects more quickly than the word ‘lemon’ when asked to name a fruit). How these results are reflected in Conceptual Semantics was shown above on the example of the verb ‘kill’.

I am not aware of TCS connecting to prototypicality results (it cannot, since it is wedded to the idea of necessary and sufficient conditions), or in general to much psychological work at all. Larson & Segal (1995), who explicitly aim to model speaker’s semantic competence, neither cite much psychology, nor attempt to show that their derivations are psychologically realistic.

The next step of integration is integration with neuroscience. It has to be shown how meanings are instantiated in the brain. This project is very much under way, although the results are mixed for now (for a recent review see Frisby et al. 2023). Within this project, semantic representations ‘can be defined as cognitive and neural states that express conceptual structure and support semantic retrieval/inference’ (Frisby et al. 2023, p. 259). It has to be figured out whether the ‘neuro-semantic’ code is homogeneous (units of the code all adopt the same activation when the represented information is present) or heterogeneous (different units express the same information differently), consistent (a given piece of information is always expressed with the same activity change) or inconsistent, whether units encode a piece of semantic information independently or conjointly, whether, within an individual, a given semantic element is anatomically contiguous or dispersed, etc. (pp. 263-5). The methods used are functional magnetic resonance imaging, electroencephalography, magnetoencephalography and electrocorticography. Within this field, Conceptual Semantics belongs to the amodal theories of concepts, as opposed to grounded theories, which see concepts as directly grounded in perception and action (Conceptual Semantics does

however connect concepts to perception and action, see below). In any case, the prospects seem good for integrating this theory of meaning with neuroscience as a more basic natural science.

TCS doesn't seem to be well placed for such an integration. I am not aware of any attempts to show how truth conditions are instantiated in the brain. Indeed, since truth conditions are, as was pointed out above, world-involving, they cannot really be instantiated solely in the brain.

So much for the first sense of 'naturalistic', namely the possibility of integration with the more basic sciences. As for the second sense, showing the relevant phenomenon to be part of the natural order, Conceptual Semantics again seems to be at an advantage compared to TCS. Jackendoff proposes 'The Parallel Architecture' (cf. his 2002: ch. 5) as a hypothesis about how meaning fits into the mind/brain more generally. According to this framework, meaning (i.e., conceptual structures) interfaces with syntax and phonology on one end and with perception and action on the other. Mediating the link to perception and action is what Jackendoff calls Spatial Structure, namely the representation of shapes and the layout of objects in the environment and their modes of motion. To use a concept CAT to recognize cats in the environment, one therefore relies on a visual prototype of cats stored in Spatial Structure. Again, I am not aware of TCS offering such a worked out model of how meaning connects to perception and action.

Of course, within the philosophy of TCS, the naturalization of meaning was supposed to go through the naturalization of reference. However, this project doesn't seem to have been successful (for a review see Loewer 1997, whose negative assessment is, I believe, still valid).

With this, I conclude my argument that Conceptual Semantics is preferable to TCS as a framework within which to study meaning.

## 5. Concluding Remarks

In this paper I have claimed that Conceptual Semantics is preferable to Truth-Conditional Semantics in the study of meaning. I basically have three arguments for this. The first argument was that, while TCS is greatly endangered by contextualism, CS successfully avoids this challenge, while offering satisfying explanations. The second argument was that, relying as it does on truth conditions, which are arguably outside of the head, TCS doesn't model well the speaker's knowledge of meaning, while CS, as explicitly mentalistic, does. The final argument is that CS is better placed to integrate with the natural sciences, which I take to be a desirable goal.

## REFERENCES

- Borg, E. (2012). *Pursuing Meaning*. Oxford: Oxford University Press.
- Calzavarini, F. (2021). Truth-Conditional Cognitivism and the Lexical Problem. *Topoi*, 40(1), 43–54.
- Chierchia, G., McConnell-Ginet, S. (2000). *Meaning and Grammar* (2nd ed.). Cambridge, MA: MIT Press.
- Cappelen, H., Lepore, E. (2005). *Insensitive Semantics*. Oxford: Blackwell.
- Croft, W., Cruse, D. A. (2004). *Cognitive Linguistics*. Cambridge: CUP.
- Davidson, D. (1984). *Inquiries Into Truth and Interpretation*. Oxford: OUP.
- Fodor, J. A. (1998). *Concepts*. Oxford: Oxford University Press.
- Frisby, S. L., Halai, A. D., Cox, C. R., Lambon Ralph, M. A., Rogers, T. T. (2023). Decoding Semantic Representations in Mind and Brain. *Trends in Cognitive Sciences*, 27(3), 258–281.
- Glanzberg, M. (2014). Explanation and Partiality in Semantic Theory. In A. Burgess, B. Sherman (Eds.), *Metasemantics* (pp. 259–292). Oxford: OUP.
- Jackendoff, R. (1983). *Semantics and Cognition*. Cambridge, MA: MIT Press.
- Jackendoff, R. (1990). *Semantic Structures*. Cambridge, MA: MIT Press.
- Jackendoff, R. (2002). *Foundations of Language*. Oxford: OUP.
- Jackendoff, R. (2011). Conceptual Semantics. In C. Maienborn, K. von Stechow, P. Portner (Eds.), *Semantics: An International Handbook of Natural Language Meaning* (pp. 688–709). Berlin/Boston: de Gruyter.
- Jackendoff (2021). Conceptual Semantics and Its Implications for Philosophy of Language. In P. Stalmaszczyk (Ed.), *The Cambridge Handbook of the Philosophy of Language* (pp. 491–514). Cambridge: CUP.
- Larson, R., Segal, G. (1995). *Knowledge of Meaning*. Cambridge, MA: MIT Press.
- Lepore, E., Ludwig, K. (2007). *Donald Davidson's Truth-Theoretic Semantics*. Oxford: Oxford University Press.
- Loewer, B. (1997). A Guide to Naturalizing Semantics. In B. Hale, C. Wright (Eds.), *A Companion to the Philosophy of Language* (pp. 108–126). Oxford: Blackwell.
- Lycan, W. G. (2000). *Philosophy of Language*. London: Routledge.
- Lycan, W. G. (2010). Direct Arguments for the Truth-Condition Theory of Meaning. *Topoi*, 29(2), 99–108.
- Montague, R. (1974). *Formal Philosophy*. New Haven, CT: YUP.
- Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Napoletano, T. (2019). How Important Are Truth-Conditions for Truth-Conditional Semantics? *Linguistics and Philosophy*, 42(6), 541–575.
- Picazo, C. (2022). Are Utterance Truth-Conditions Systematically Determined? *Inquiry*, 65(8), 1020–1041.
- Pietroski, P. M. (2018). *Conjoining Meanings*. Oxford: Oxford University Press.
- Pinker, S. (1989). *Learnability and Cognition*. Cambridge, MA: MIT Press.
- Pinker, S. (2007). *The Stuff of Thought*. London: Allen Lane.
- Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.

- Putnam, H. (1975). The Meaning of 'Meaning'. In his *Mind, Language, and Reality* (Philosophical papers, Vol. 2, pp. 215–271). Cambridge: CUP.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-Conditional Pragmatics*. Oxford: OUP.
- Schlenker, P. (2022). *What It All Means*. Cambridge, MA: MIT Press.
- Searle, J. (1980). The Background of Meaning. In J. Searle, F. Kiefer, M. Bierwisch (Eds.), *Speech Act Theory and Pragmatics* (pp. 221–232). Dordrecht: Reidel.
- Tarski, A. (1944). The Semantic Conception of Truth. In L. Linsky (Ed.), *Semantics and the Philosophy of Language* (pp.13–49). Urbana: The University of Illinois Press.
- Travis, C. (1997). Pragmatics. In B. Hale, C. Wright (Eds.), *A Companion to the Philosophy of Language* (pp. 87–107). Oxford: Blackwell.
- Travis, C. (2008). *Occasion-Sensitivity*. Oxford: Oxford University Press.