

## **Can Epistemic Analyticity Explain A Priority?,**

Forthcoming in Dodd/Zardini (eds.) (2019), *The Apriori*, OUP

Magdalena Balcerak Jackson (Miami)

One of the most important motivations for the philosophical interest in the notion of analyticity has always been the prospect of using such a notion to explain our acquisition of a priori knowledge of logical, of mathematical, and of philosophical claims. While Quinean doubts have convinced many that it is hard to spell out a coherent notion of analyticity in metaphysical terms, many philosophers hold on to the idea that an epistemic notion of analyticity is coherent and non-empty, and are optimistic that it can provide good grounds for explaining a priori knowledge. In this paper, I will argue that this optimism is misguided. My argument is not based on a wholesale Quinean rejection of the analytic/synthetic distinction, or on skepticism about the a priori. Rather, I will present a dilemma for understanding-based accounts of the a priori, the upshot of which is the following: either the underlying notion of understanding is independently well-understood and well-supported, but only relatively few and non-interesting claims turn out to be epistemically analytic and therefore knowable a priori; or else the notion of understanding is thick and demanding, but turns out to rely on a priori knowledge (or closely related notions), and therefore cannot explain it.

### **1. Analyticity and Apriority**

According to the standard conception, the analytic-synthetic distinction is a distinction between two different kinds of truths. Synthetic truths are true partly in virtue of what they mean and partly in virtue of how the world is. Analytic truths, in contrast, are true just in virtue of meaning alone. This is a metaphysical conception of the analytic/synthetic distinction; analytic truths are those whose truth is grounded or explained in a certain way. Following Quine, many philosophers nowadays doubt that this metaphysical conception of analyticity is coherent.<sup>1</sup> However, Paul Boghossian has convincingly shown that we can also think of analyticity in epistemic rather than metaphysical terms. On the epistemic conception of analyticity, a true thought is analytic if and only if understanding it – that is understanding all the concepts involved and the

---

<sup>1</sup> This might be wrong. Gillian Russell has defended a metaphysical conception of analyticity in detail in Russell (2008).

ways they are combined – puts one in a position to know it. Many philosophers who are influenced by broadly Quinean worries about metaphysical analyticity nevertheless continue to see epistemic analyticity as a viable notion worth taking seriously.

One main reason for this continued interest in epistemic analyticity comes from the very traditional idea that we can use analyticity to explain how we can come to know certain truths a priori.<sup>2</sup> As Jerry Fodor and Ernie Lepore write:

*So maybe if there is a synthetic/analytic distinction, we could explain why the necessary truths, or at least some of the necessary truths, are knowable by anybody who knows a language that can express them. (...) It would be ever so nice to understand how a priori knowledge is possible.*<sup>3</sup>

Similarly Georges Rey takes it to be an obvious truth about the history of analytic philosophers that,

*Beginning with Frege, many philosophers have hoped that the apparent necessity and a priori status of claims of logic, mathematics and much of philosophy would prove to be due to these claims being analytic. (...) Why should philosophers be interested in what seems like a purely linguistic notion? Because especially in the first half of the Twentieth Century, many philosophers thought it could perform crucial epistemic work.*<sup>4</sup>

We seem to be justified in accepting claims about logic, mathematics and at least some kinds of claims in philosophy without relying on sense experience to provide the justification. How is this possible? The proponent of metaphysical analyticity has an answer for at least some of these cases. She can say that these claims are true in virtue of meaning alone, and so there is no need to look out into the world to find reasons to accept them. But if the notion of metaphysical analyticity is incoherent, this explanation of a priori knowledge fails. But can an explanation be given by relying on the epistemic notion of analyticity instead? Many philosophers think so. According to the proponent

---

<sup>2</sup> The interest in analyticity is often extended to the explanation of necessary truths, as the following quotes indicate. Nowadays, most philosophers reject a straightforward connection between necessity and a priority. Even if I believe that there is a notion of necessity that goes hand in hand with a priority, in the context of this paper I will restrict my attention to the project of explaining the possibility of knowledge of a priori truths, that is truths that can be known a priori, by appeal to analyticity.

<sup>3</sup> Fodor/Lepore (1996), 114.

<sup>4</sup> Rey (2003/2013).

of epistemic analyticity we can know at least some kinds of claims a priori because merely understanding these claims puts us in a position to know them. It is a tricky question what exactly it takes to be in a position to know a certain claim. Are we in a position to know something only if we will or would come to know it immediately upon attending to it? Or is it a matter of being able to come to know it upon sufficient reflection, or perhaps only upon ideal reflection? Fortunately, the answer to these questions do not matter for the purposes of this paper. I will simply assume that there is a good account of what it means to be in a position to know, and proceed from there. The crucial idea behind the explanation of a priori knowledge via epistemic analyticity is simply that, if we can come to know certain claims purely on the basis of understanding them, our knowledge of them need not epistemically depend on perception or introspection, or more generally on the cognitive capacities that are sources of a posteriori justification. So, at least for those claims that are epistemically analytic, it is no mystery that we can come to know them a priori; we can have knowledge that is based purely in our understanding of them.

Given this, it should be clear that a comprehensive and explanatory account of epistemic analyticity is subject to an important requirement: it must elucidate a notion of understanding that can serve as a plausible foundation for epistemic analyticity. What is it to understand concepts and thoughts? The discussion about epistemic analyticity to date has largely focused on the question whether there are any epistemically analytic truths, and on related questions about whether particular truths qualify as epistemically analytic. This might be taken to suggest that philosophers mostly agree on the question of what understanding is, at least in broad outline. But this is not the case. As we are about to see, disagreement about the existence of epistemically analytic truths can in fact be rooted in deep disagreement about what understanding is.

## **2. Are There Any Epistemically Analytic Truths?**

In recent prominent and influential work, Timothy Williamson argues that there simply are no epistemically analytic truths or entailments. “Nothing,” he says, “is epistemically available simply on the basis of linguistic and conceptual competence.”<sup>5</sup> Williamson intends his arguments to have implications for meta-philosophical views: If he can

---

<sup>5</sup> Williamson (2007), page.

successfully show that there are no epistemically analytic truths, then he has eliminated what he takes to be the best version of the view of (a priori) philosophy according to which philosophers aim to discover conceptual truths in some interesting sense of the term.<sup>6</sup> Williamson's strategy is to suggest that, for any thought *T* such that it is *prima facie* plausible that understanding *T* necessarily requires assenting to *T*, we can create a counterexample: a subject *S* who understands *T*, but who – even after careful, lengthy reflection – fails to assent to it. Williamson demonstrates a recipe for creating counterexamples by using simple logical truths such as “Every vixen is a vixen”, which many take to be among the most promising candidates for epistemic analyticity. Williamson asks us to consider two subjects, Peter and Stephen. Both grew up in our very own linguistic community and both are successful communicators and users of the English language. But each has a set of peculiar beliefs. Peter believes that any thought of the form that every *F* is *G* is only true if there is at least one *F*. He also believes that there are in fact no vixens. Stephen believes that every thought of the form that every *F* is *G* is only true if there are no borderline cases of *F* (or *G*). And he believes that there are in fact borderline cases of vixenhood. Williamson argues that Peter and Stephen both understand the thought “Every vixen is a vixen” even though they fail to assent to it. So understanding “Every vixen is a vixen” does not necessarily entail that one assents to it.

In a recent paper, Brendan Balcerak Jackson and I argue that Williamson makes an important mistake. He assumes that we should regard epistemically analytic truths as truths for which understanding-assent guarantees hold. An understanding-assent guarantee holds for a thought *T* if and only if whoever understands *T* assents to it. The point of the Peter and Stephen cases is to make it plausible that there aren't any understanding-assent guarantees. It is natural to read such a construal into Boghossian's work. But philosophers who hope to base philosophical methodology on our conceptual capacities do not need to rely on understanding-assent guarantees. It is perfectly possible to base the epistemology of a discipline or field of enquiry on certain cognitive capacities without assuming that the possession of these capacities provide any guarantee that their possessor comes to know the truths to which they provide access. Even if mathematical competence does not *guarantee* assent to – and ultimately knowledge of – a complicated mathematical theorem, it does not follow that something other than mathematical

---

<sup>6</sup> At least in Williamson (2007) Williamson also argues against alternative construals of philosophy as a conceptual enterprise.

competence must be at work when an excellent mathematician comes to know the theorem. And even if understanding does not guarantee assent to – and ultimately knowledge of – a complicated analytic truth, it might still be that nothing but understanding is at work in helping somebody who does acquire knowledge of the truth.<sup>7</sup>

Why doesn't Williamson consider this possibility? We suspect that ultimately the disagreement with Williamson bottoms out in a fundamental difference in how we think about understanding itself. Williamson endorses a social externalist, minimalist conception according to which standing in the right sort of causal relationships to other competent users of the language is all it takes to make one an understander in the full sense of the term.<sup>8</sup> By contrast, we push for a more substantial conception, according to which understanding is a complex competence that includes not merely dispositions to apply concepts in actual circumstances, but also the ability to think about their application in hypothetical circumstances, the ability to manipulate them in basic reasoning, and so on. On this way of thinking, understanding is a set of capacities to be exercised. They can be exercised well or poorly, and their exercise can have different results in different cases.

The important lesson, then, is that the conception of understanding one favors matters a great deal for where one stands on the question of whether there are any epistemically analytic truths, and on questions about how many and what sorts of truths these might be.

### **3. The Dilemma Concerning Understanding**

Let me recapitulate where we stand. Many philosophers who seek a non-mysterious explanation of the a priori put their hopes on epistemic analyticity, which, unlike its metaphysical cousin, does not seem to suffer from the dangers of incoherence. But the question of how populous the space of epistemically analytic truths is – and indeed whether it has any population at all – hangs on our conception of understanding, which is something that philosophers do not agree upon, and are very unlikely to start.

---

<sup>7</sup> For the full presentation of the argument see Balcerak Jackson / Balcerak Jackson (2012).

<sup>8</sup> Williamson (2007), 90-91.

This is the point at which my main argument takes hold. The argument will take the form of a dilemma. To make progress on questions about epistemic analyticity we need to settle on a way to think about understanding. We can construe understanding in either what I will call a *thin* way or in what I will call a *thick* way. If we construe it thinly – such that having linguistic competence with a word as it is used within a relevant linguistic community suffices for understanding the concept the word expresses – we get an empirically well-grounded conception of analyticity. But this conception is only able to account for trivial a priori knowledge. If we construe understanding thickly – as everything that a rational subject needs in order to engage in some form of rational enterprise of conceptual analysis – then we plausibly do get non-trivial analytic truths, but it becomes unclear whether analyticity can still be taken as explanatorily prior to a priority.

### 3.1. Thin Understanding

Let's take a closer look at the first horn of the dilemma: accepting a thin conception of understanding. The thin conception of understanding links understanding to linguistic competence as studied by linguists. To a first approximation, thinly understanding concepts and thoughts is a matter of being linguistically competent with expressions of a specific natural language that are conventionally used to express them.<sup>9</sup> The motivation for the thin conception is two-fold. First, it is hard to deny that for a subject who is a speaker of a natural language, being linguistically competent with its expressions provides him with a kind of grasp or understanding of the concepts they express. Second, the notion of understanding is in need of precisification, as we have seen, and tying it to the cognitive capacities investigated by linguists provides us with a non-arbitrary way of doing so. It is not just philosophers occupying some more or less random spot in the space of theoretical possibilities.

Of course, there are also difficulties with this approach. Our initial interest in understanding is primarily as a grasp of concepts and thoughts, not words and sentences.

---

<sup>9</sup> This view is compatible with the restriction that to understand some concepts, such as phenomenal concepts, more than linguistic competence with the words expressing those concepts is needed. The view should also allow for concepts that are not expressed by any words with which the subject is competent.

Concepts and thoughts are mental representations, words and sentences linguistic entities. And although there is plausibly a tight relationship between linguistic expressions and the thoughts they express, the relationship is non-trivial. For the purpose of this paper, I will simply assume that for speakers of a natural language there is a relatively straight-forward way to translate back and forth between talk of their competence with words and sentences in the natural language and talk about their competence with concepts and thoughts of our ‘language of thought’<sup>10</sup>. If this assumption turns out to be wildly mistaken, then we have reason to believe that the strategy of deferring to linguistic research when determining what understanding amounts to is not even *prima facie* a good strategy to pursue. This will only bolster the conclusion to be reached here, that the thin conception of understanding cannot support a robust account of a priori knowledge.

For the linguist, competence with a language has many inter-related components: phonological, syntactic, and semantic. What is most relevant for our purposes, however, are those branches that focus on the semantic aspects of competence. Within semantics, there are two branches with their own methodologies: lexical semantics and compositional semantics. The former studies the meanings of individual words of natural languages, while the latter studies how words and other expressions combine to form further meaningful expressions. Both branches of semantics tend to tell us about linguistic competence only in an indirect way. What they directly do is investigate hypotheses about meaning: about the meanings of various expressions and how they are determined, about systematic relationships among meanings and the rules underlying them, and so on. To get from such claims to claims about linguistic competence we need to assume a bridge principle, such as the principle that to be linguistically competent with an expression is to know or grasp the facts about meaning being described, or perhaps to know or grasp certain aspects of the semantic theory being offered. Still, the basic idea of the thin conception is that we can look at what semantics tells us about the meaning of individual and complex expressions in order to extract a characterization of linguistic competence, and so ultimately of understanding.

---

<sup>10</sup> The simple quotation marks are intended to signify that I do not wish to commit myself, or the reader, to the thesis that there is a language of thought literally understood (as Fodor as others did), but merely refer to the general idea that it is by means of concepts and the thoughts composed out of them that we reason and deliberate and that it is concepts and thought that we often report using words and sentences.

Can we use compositional and/or lexical semantics in this way to provide a conception of understanding that might help account for a priori knowledge? In the next two subsections I argue that we cannot.

### *Compositional Semantics*

Compositional semantics is the study of how the meaning of complex expressions systematically depend on the meanings of their parts (such as words). The dominant paradigm is broadly truth-conditional. The goal of compositional semantics thus understood is to systematically characterize how the truth-conditions of a sentence are a function of the meanings of its constituents and the way these constituents are put together. These characterizations are typically provided using a formal, artificial language such as First Order Logic or the Lambda Calculus. Theorizing in compositional semantics is not so much interested in merely listing the truth conditions of each sentence of the language under study. Rather, it aims to provide a systematic explanation of various phenomena that can be observed when we pay close attention to how competent speakers use language. These phenomena include the intuitive judgments of competent speakers about the acceptability of certain complex expressions, about entailment relations between sentences, and about ambiguity. For example, consider the following:

- (1) Colorless green ideas sleep furiously
- (2) Grace laughed loudly
- (3) Grace laughed
- (4) Every dog barks
- (5) Every spotted dog barks
- (6) She saw the man with one eye

A competent user of English will readily judge that (1) is unacceptable (though grammatical), that (2) entails (3), that (4) entails (5), and that (6) is ambiguous between a reading on which the seeing was performed with only one eye, and a reading on which the man seen had one eye. Moreover, speakers' judgments like these generalize across a wide range of similar linguistic constructions. For example, there is a large class of

sentence pairs structurally analogous to (2) and (3) for which the entailment judgment holds.

Semantic explanations of such data paradigmatically involve identifying the logical forms of the sentences involved; this often requires identifying covert elements that are absent from the surface sentence structure. For example, many semanticists would identify the logical form of (2) roughly as follows:

(2A)  $\exists e[\text{Agent}(e, \text{Grace}) \ \& \ \text{Laughed}(e) \ \& \ \text{Loud}(e)]$

(2A) represents the meaning of (2) where ‘Laughed’ and ‘Loud’ are “predicates of events and ‘Agent’ expresses a thematic relation that holds between an event done by someone (or something) and the relevant person (or thing)”<sup>11</sup>. Similarly, many semanticists would represent (3) as follows:

(3A)  $\exists e [\text{Agent}(e, \text{Grace}) \ \& \ \text{Laughed}(e)]$

The inference from (2A) to (3A) is simply an instance of conjunction reduction within the scope of an existential quantifier. So, if (2A) and (3A) capture the logical forms of (2) and (3), and this is something that competent speakers know or grasp, then we have a ready explanation of their intuitive judgment that (2) entails (3).<sup>12</sup>

Let us grant that compositional semantics helps elucidate linguistic competence in the way just illustrated. What consequences does this have for epistemic analyticity on the thin conception of understanding tied to linguistic competence? Are there any thoughts that we in a position to know merely in virtue of being linguistically competent with their constituents and their compositional structure?

It would seem that there are, at least if the explanation of speakers’ judgments about the entailment from (2) to (3) – or others like it – is correct. The core of that explanation is that a speaker’s competence with (2) and (3) allows her to recognize that the former

---

<sup>11</sup> Pietroski (2000), page (online 2). Similar proposals about logical form can be found in e.g. Davidson (1967), Parsons (1994), Higginbotham (1985), Schein (1994)

<sup>12</sup> More discussion of the role of compositional semantics in accounting for linguistic competence, along with references, can be found in Pietroski (2000).

entails the latter. (Note that logical competence alone will not do here: the surface forms of (2) and (3) do not even contain any logical constants.) If this is correct, merely understanding (2) and (3) in the thin sense puts one in a position to know that the entailment holds. The same holds for many other cases, since the entailment from (2) to (3) is one instance of a pattern that has many others, such as the following:

- (7) Abby quickly ran; so Abby ran.
- (8) Abby loudly chewed; so Abby chewed.

The semantic explanation provided for (2) and (3) above applies just as well to these cases, because that explanation does not rely on any hypothesis about the specific meaning of ‘laugh’ and ‘loudly’; it relies only on a hypothesis about the semantic significance of adverbial modification. If understanding these sentences amounts to being linguistically competent with them, and linguistic competence puts you in a position to know that the entailments in this pattern hold, then they will qualify as epistemically analytic – and so (we are supposing) knowable a priori.

The pattern above is an example of what semanticists sometimes call a *structural* entailment pattern, a pattern whose validity is due to the logical forms or compositional structures of the sentences, regardless of the specific meanings of the words involved.<sup>13</sup> There are numerous patterns of structural entailments that have been observed by semanticists:

- (9) Ray broke the glass; so the glass broke.
- (10) Julien boiled the water; so the water boiled
- (11) Julien painted the fence white; so Julien painted the fence
- (12) Ray wiped the table clean; so Ray wiped the table
- (13) John thinks that Bob likes himself; so John thinks that Bob likes Bob.
- (14) John is easy to please; so it is easy to please John.
- (15) I heard Pat sing, so Pat sang.

---

<sup>13</sup> Structural entailment patterns and their semantic explanation are discussed in detail in Pietroski (2008): *Small Verbs*. They are also discussed in Balcerak Jackson, B. (2017), who argues in Balcerak Jackson, B. (2009) for their role in defending analyticity against Williamson’s attack. Many of the examples here are directly borrowed or adapted from Balcerak Jackson (2009) or Pietroski (2008).

Each of these patterns has many, many instances, and on our current assumptions all of them plausibly qualify as epistemically analytic and so a priori knowable.

Still, while the search for a semantic theory that best explains all these patterns might be fascinating and challenging work for semantics, from the point of view of epistemology the results are rather trivial and uninteresting. The a priori knowledge we get in this way – knowledge of structural entailment relations – certainly seems to contain little in the way of philosophical discoveries or insights. A conception of epistemic analyticity based only on what compositional semantics tells us about linguistic competence is unlikely to yield an account of much of the a priori knowledge we had originally hoped to capture.

### ***Lexical Semantics***

The fact that compositional semantics fails to yield philosophically interesting a priori knowledge surely has something to do with the fact that it abstracts away from the meanings of individual words. As we saw, the explanation for competent speakers' knowledge of the entailment from (2) to (3) works just as well for a wide range of other entailments involving words that have very different meanings than 'laugh' and 'loudly'. So perhaps we should look to lexical semantics – the study of word meanings – to find a way of grounding a priori knowledge in linguistic competence.

One of the problems with looking to contemporary lexical semantics, however, is that there is little agreement between linguists on how to pursue it. According to the linguist Dirk Geeraerts, whose book provides both a systematic overview of lexical semantics and a description of its historical development, lexical semantics is a complicated landscape of mountains and rivers with “broad vales where the main streams of research flow, branch off into side valleys and even smaller dales where theories are refined and specific topics pursued.”<sup>14</sup> To put it more bluntly, lexical semantics has no agreed upon and generally accepted methodology.

Nevertheless, Geeraerts identifies two very broad and very loosely unified main streams within contemporary lexical semantics: a *restrained* approach, which seeks to “restrain

---

<sup>14</sup> Geeraerts (2010), Introduction.

semantic description either by maintaining some form of distinction between meaning and cognition, and/or between meaning and use, or by subjecting semantic descriptions to the requirements of a formal representation”<sup>15</sup>, and a more *maximalist* approach that “is wary of the distinction between semantic knowledge and encyclopedic knowledge, that is to say, of the belief in an autonomous level of linguistic structure that is strictly separated from cognition in the broader sense”<sup>16</sup>.

A maximalist approach, in which the distinction between linguistic and encyclopedic knowledge is systematically jettisoned, cannot deliver the epistemological foundation we are looking for. It is not hard to see why. Typical examples of the maximalist approach, such as prototype theory and other branches of cognitive semantics, seek to investigate how our language use relates to the psychological act of categorizing perceived objects in the world.<sup>17</sup> Cognitive semanticists are not mainly interested in which objects our words refer to, but in how our words are related to the ways we think about objects, including ways that are influenced by our individual and social backgrounds. Here is a simplified prototype-theoretical example. The word “bird” as used by me is associated with a fuzzy but ordered collection of representations of birds. Swallows and the puffin Oona are represented in the center, chickens and penguins are further away from the center, and Sebastian the ibis, the University of Miami mascot, is in the far periphery. While some organizational features of this collection are likely to reflect a shared bird-prototype between me and fellow English speakers, others are reflective of my personal experience as a member of the UM faculty and an avid watcher of the television show *Puffin Rock*. It should be obvious that knowledge or grasp of prototypes, or other experientially grounded representations the cognitive semanticist postulates as being associated with words, goes well beyond what it takes to count as a competent user of the word. And it should be equally obvious that the cognitive semanticist’s systematic rejection of the distinction between linguistic and encyclopedic knowledge makes their approach to meaning unsuitable to serve as the basis for an account of a priori knowledge.

---

<sup>15</sup> Geeraerts (2010), page (online242).

<sup>16</sup> Ibid., page (online 244).

<sup>17</sup> Proponents include Eleanor Rosch, the founder of prototype theory, George Lakoff or Peter Gaerdenfors.

In contrast, the restrained approach to lexical semantics takes its starting point from the structuralist assumption that the job of semantics is to “describe the structure of a sentence in isolation from its possible setting in linguistic discourse (written or verbal), or in non-linguistic contexts (social or physical).”<sup>18</sup> Given this assumption, the restrained approach is – at first glance – a better contender to elucidate meanings of words in a sense that could be a useful basis for characterizing linguistic competence with them. It explicitly demands a separation of the semantic study of meaning from the pragmatic and from the encyclopedic. To see whether this first impression is confirmed, however, we need to take a closer look at the restrained approach, and at the kinds of information about the meanings of words that it delivers (or attempts to deliver). We will focus on two examples of the restrained approach: Anna Wierzbicka’s decompositional approach and the WordNet database project.

Wierzbicka’s *natural semantic metalanguage* tries to reduce the semantics of all lexical items to a highly restricted set of semantic primitives. The semantic primitives are supposed to express a set of universally shared basic concepts, and so the theory requires a thorough cross-linguistic investigation to isolate only those concepts that re-appear again and again across different linguistic communities. The result is a list of semantic primitives represented by their English “exponents,” that is by those English words whose meanings are or approximate the universal concepts identified. As of this writing, the list comprises 65 primitives, including concepts such as I, YOU, GOOD, SMALL, SEE, HAPPEN, NOT, AFTER or BELOW. The purpose of the list is to provide the basis for “reductive paraphrases” of the concepts expressed by all the other words of a natural language like English solely in terms of the semantic primitives.<sup>19</sup> For example, here is Wierzbicka’s reductive paraphrase of ‘disappointment’:

X is disappointed =  
X feels something  
sometimes a person thinks something like this:  
something good will happen  
I want this  
after this, this person thinks something like this:  
I know now: this will not happen  
because of this, this person feels something bad

---

<sup>18</sup> Fodor/Katz (1963), 173.

<sup>19</sup> For a recent comprehensive introduction into the program see Goddard/Wierzbicka (2014).

X feels like this<sup>20</sup>

This is an ambitious program, and there are obvious worries about it. The first is that it is doubtful that the current list of 65 or so primitives is indeed cross-linguistically universal. Second, the proposed reductive paraphrases are very often inadequate because they hopelessly oversimplify, or because they exclude some obvious referents of the term. For example, the definition of ‘disappointed’ presented above includes the condition that the subject has (or had) the thought that something good will happen. But one can be disappointed by an outcome even if one merely hoped for it and did not believe that it was even likely. One can even be disappointed by outcomes that one had no particular attitudes about before events unfolded. This kind of problem is pervasive.<sup>21</sup> And it is largely due to its restrained approach: it seems wildly optimistic to think that a list of even 65 primitives will be enough to provide for plausible reductive paraphrases of every word of English. While proponents of natural metalanguage semantics spend a great deal of effort gathering empirical confirmation for the list of primitives, they spend comparatively very little time testing their reductive paraphrases. Without such testing, though, it is very hard to say what Wierzbicka’s postulated definitions really tell us about linguistic competence with the words defined. Given these serious methodological flaws, the enterprise seems more interesting as an intellectual exercise than as a rigorous study of linguistic competence. If the motivation behind a thin notion of understanding is a close connection to a well-established and widely-respected linguistic research method, we should be hesitant to defer to natural meta-language semantics.

Instead of trying to propose definitions, the large, electronic database WordNet attempts to capture paradigmatic lexical relations between words of English. It groups synonymous words, such as ‘car’ and ‘automobile’ into unordered synsets. And, it represents relations like hyponymy (super-subordinate relations) and meronymy (part-whole relations) between noun-synsets, and relations like antonymy between adjective-synsets.<sup>22</sup> A search for ‘bird’ in WordNet<sup>23</sup> returns information about synsets including this:

---

<sup>20</sup> Wierzbicka (1992), 549.

<sup>21</sup> See example concerning Wierzbicka’s definition of ‘fruit’ as discussed Geeraerts (2010), page

<sup>22</sup> More information can be found in Fellbaum (2005).

<sup>23</sup> You can try out WordNet at <http://wordnet.princeton.edu/>.

(S1)(n) **bird** (warm-blooded egg-laying vertebrates characterized by feathers and forelimbs modified as wings)

(S2)(n) [dame](#), [doll](#), [wench](#), [skirt](#), [chick](#), **bird** (informal terms for a (young) woman)

and information about hyponymy relations involving synset S1 including this:

**bird; vertebrate; chordate; animal; organism; living thing; whole/unit; object; physical entity; entity**

So, WordNet can be used as a dictionary and as a thesaurus, but also as a tool to calculate various degrees of semantic similarity and relatedness that result from the encoded data about specific synsets associated with words of English. Does the lexical information about a word of English contained in WordNet inform us in any way about what a user of the word needs to know in order to count as linguistically competent with respect to that word? Both the definitional gloss contained in the synset list associated with 'bird' and the list of categories standing in a hierarchical hyponymy relation to 'bird' contain too much. A speaker surely does not count as linguistically incompetent with the term 'bird' if they have no knowledge of the fact that birds are vertebrates. WordNet might be a plausible representation of some important semantic relations between words of English, but it does not even seem to be designed to tell us what sorts of information must be grasped by someone to be competent with the words in its database.<sup>24</sup>

At first glance, it might seem quite plausible that lexical semantics, as the study of word meaning, is the empirical discipline that is most likely to contribute to a thin conception of understanding that can be of use for epistemic analyticity. But we have seen that this is not so, for two significant reasons. The first is that the lack of agreement amongst lexical semanticists how to pursue it, and serious flaws in the methodology of specific prominent streams of lexical semantics, should make us wary of any deferential reliance on its results at this point.<sup>25</sup> The second reason is more fundamental: The semantic hypotheses that dominant streams of lexical semantics offer go well beyond aspects of meaning that are reasonably constitutive of linguistic competence. Very little about how

---

<sup>24</sup> It should also be noted that the sources of the data encoded in WordNet are highly diverse and unsystematic, putting into question how much of the information captured in WordNet corresponds to things that speakers know by virtue of their linguistic competence, as opposed to things that surface in actual language processing and use.

<sup>25</sup> This does not mean that the discipline cannot mature into a field of study that would not warrant such a deference.

lexical semantics is done would help us distinguish those aspects of word meaning that have to be known or grasped to count as linguistically competent from those that do not.

To sum up: On the one hand, the prospects for grounding a conception of understanding in lexical semantics seem dim. On the other hand, grounding a conception of understanding in compositional semanticists does seem to yield a perfectly acceptable notion of epistemic analyticity, and one that is far from empty. But it is one that is unlikely to do more than account for the a priority of trivial structural entailment patterns. If we are to capture the kinds of substantive a priori knowledge that we were originally hoping to explain by reference to epistemic analyticity, we will need to turn to a thicker conception of understanding.

### 3.2. Thick Understanding

The thick conception links understanding to our ability to reason with concepts, or to the ability to make an interesting set of rational judgments by relying on our concepts. This general idea will need to be developed bit by bit.

There are three motivations for a thick conception of understanding. First is the one just indicated, that a thick conception of understanding might do better than a thin conception with respect to the goal of explaining substantial a priori knowledge. Second, there is a natural connection between understanding and rational reasoning, but rational reasoning of any kind requires more than mere linguistic competence. And third, the thick conception captures some intuitions about understanding that demand that we see understanding as involving additional complex cognitive capacities. For example, consider Momo, who is perfectly capable of verbally identifying cars when being presented with them in the actual world, but who balks at the suggestion to call something a ‘car’ when it is presented to him in a fictional story, or in a scenario that he visually imagines. Intuitively, Momo does not really understand the concept ‘car’. Or consider Magali, who can correctly verbally identify every animal species in the world, but who cannot answer a single question about what it is about the creature in front of her that leads her to identify it as a zebra or a sloth or a wonderpus photogenicus<sup>26</sup>.

---

<sup>26</sup> No need to google unless you want a picture: A wonderpus photogenicus is a rare octopus species with a dramatic color pattern of spots and bars.

Intuitively, Magali does not really understand the concepts ‘zebra’, ‘sloth’ or ‘wonderpus photogenicus’. Granted, intuitions of this kind can be debated. Someone who is happy to characterize Williamson’s Stephen as a subject who fully understands ‘every’, even though he rejects ‘every vixen is a vixen’ might want to resist the intuitive judgments about Momo and Magali. But remember that the goal is not to decide which conception of understanding is correct one, but rather to explore the consequences of different conceptions for epistemic analyticity.

Appeals to a thick notion of understanding can be found in the work of Paul Boghossian, Christopher Peacocke, David Chalmers and Frank Jackson, among others. For Boghossian, “grasping” a non-defective concept involves “being disposed to engage in some sort of inferences involving them and not in others”<sup>27</sup>. For Peacocke a thinker’s “understanding” of at least some concepts involves her “possession of an implicit conception” which is manifested by her “reflection” on the truth-value of various claims involving the concept.<sup>28</sup> For Chalmers and Jackson “possession of a concept (...) bestows a conditional ability to identify a concept’s extension under a hypothetical epistemic possibility.”<sup>29</sup> These views have in common the idea that understanding a concept is at least partly constituted by an ability to make rational judgments of certain sorts involving the concept. This will need refinement, and not all of these philosophers are engaged in the project of trying to explain a priority in terms of epistemic of analyticity. But the thick conception of understanding that they have in common does support a characterization of epistemic analyticity, one that says that a true thought is epistemically analytic if and only if having the ability to make rational judgments of a certain sort involving the relevant concepts puts one in a position to know it.

What sorts of rational judgments are the ones that matter for understanding? For Peacocke and for Chalmers and Jackson, they are the sorts of judgments that we make when we engage in the process of conceptual analysis. It is precisely when engaged in this process that put to use the ability that constitutes (thick) understanding of concepts. This is not to say that the notion of ‘conceptual analysis’ is taken as explanatorily prior to the notion of ‘understanding’. But we can use our grasp of the process of conceptual

---

<sup>27</sup> Boghossian (2003), 240. See also Boghossian (1996).

<sup>28</sup> Peacocke (1998), 44-46.

<sup>29</sup> Chalmers/Jackson (2002), page (online 9)

analysis to extract, in a reverse-engineering sort of way, information about the cognitive ability involved in this process.

In his paper “Implicit Conceptions, Understanding and Rationality” Peacocke describes a rational way for a subject to come to accept the simple inference rule: ‘From A, the conclusion A or B can be inferred’:

“The reflection involves a simulation exercise. The thinker imagines (...) that A is true and B is false. His aim is to address the question of whether the alternation ‘A or B’ should be regarded as true or false in the imagined circumstances. (...) This capacity is the very same, understanding-based capacity he would be exercising in a real case in which he had the information that A is true and B is false, and has to evaluate the alternation ‘A or B’. As in the corresponding real case, in the imaginative exercise he goes on to hold that ‘A or B’ will be true in the imagined circumstances. In coming to hold that ‘A or B’ is true in the simulated circumstances, our thinker employs only the information about the truth-values, within the simulation of A and of B, together with his understanding of alternation. He does not draw on any other resources. Next our thinker proceeds to consider imaginatively another case, say in which A is true and B is true...”<sup>30</sup>

A quite similar way of coming to know a true thought – in this case the thought that one can justifiably and truthfully believe a proposition and yet not know it – is described by Chalmers and Jackson in their paper “Conceptual Analysis and Reductive Explanation”:

“Let G be the conjunction of the statements in the following passage: ‘Smith believes with justification that Jones owns a Ford. Smith initially has no beliefs about Brown’s whereabouts. Smith forms a belief that Jones owns a Ford or Brown is in Barcelona, based solely on a valid inference from his belief that Jones owns a Ford. Jones does not own a Ford, but as it happens, Brown is in Barcelona.’ Let K be state statement ‘John does not know that Jones owns a Ford or Brown is in Barcelona.’ (...) The conditional  $G \supset K$  (...) plays an essential role in Gettier’s argument for the conclusion that knowledge is not

---

<sup>30</sup> Peacocke (1998), 45.

justified true belief. (...) The argument proceeds by presenting the possibility that G holds and appealing to the reader's concept of knowledge to make the case that if G holds, K holds (and J holds where J is a corresponding positive claim about John's justified true belief). Empirical information plays no essential role in justifying this conditional. (...) The model of conceptual analysis that emerges is something like the following. When given sufficient information about a hypothetical scenario, subjects are frequently in a position to identify the extension of a given concept, on reflection, under the hypothesis that the scenario obtains. Analysis of a concept proceeds at least in part through consideration of a concept's extension within hypothetical scenarios and noting regularities that emerge."<sup>31</sup>

The resulting picture of conceptual analysis – or of the process of making an implicit conception explicit, to speak in Peacocke's terms – is the following. A subject is looking for an answer to a certain question about X. Is a certain rule of inference involving disjunction valid? Is knowledge justified true belief? She imagines or supposes that a certain scenario obtains. She then makes a certain judgment involving X about the scenario under the supposition, or 'under the imagining'. She judges that in this case 'A or B' is true. Or she judges that in this case Smith does not know. Finally, she decides how the judgment she made under the supposition or under the imagining bears on the question under discussion. Sometimes the judgment might straightforwardly logically entail an answer to it. If Smith does not know, then knowledge cannot be justified true belief. In other cases, it will only do so in conjunction with other judgments about other cases. The validity of the inference rule 'From A, the conclusion A or B can be inferred' can be accepted based on a series of judgments about hypothetical cases that represent all possible truth-value assignment for A and for B. In still other cases, the judgment will only combine with other judgments about other cases to form an data set from which an answer to the question under discussion can be abductively inferred.

Let us assume for the sake of this paper, that we can rationally engage in this type of analysis. According to a proponent of the thick conception, when we engage in this type of analysis we manifest our understanding. The complexity of the process of analysis has to be matched by a complex and rich conception of understanding. What does

---

<sup>31</sup> Chalmers/Jackson (2001), 320-322.

understanding have to involve to enable us to engage in conceptual analysis thus understood? It plausibly has to include linguistic competence with words and semantic structure. But it also has to include the ability to imagine or to suppose possible hypothetical scenarios. It has to include basic logical abilities. And it has to include the ability to reason according to inference to the best explanation.<sup>32</sup> Peacocke acknowledges that the method of conceptual analysis requires general cognitive abilities that we use in other epistemic contexts:

“Some of the intellectual skills required to succeed in making an implicit conception explicit will be skills useful in any enterprise of building explanations from instances. (...) So, even in trying to articulate what is influencing one in making judgments involving the concept in particular cases, the skills and methodology involved are those pertinent to any abductive investigation.”<sup>33</sup>

And when Chalmers and Jackson express their ideas carefully, they explicitly note that a mere possession of concepts without functioning reasoning abilities will not be enough to enable us to perform conceptual analysis: “If a subject possesses a concept and has unimpaired rational processes, then sufficient information about the actual world puts a subject in a position to identify the concept’s extension.”

The challenge for a thick conception of understanding is that not all evaluations of hypothetical cases amount to instances of conceptual analysis. And not all judgments we make about hypothetical scenarios are plausibly judgments based solely in understanding, even in understanding in the thick sense. Consider a scenario like the following. Suppose that you imagine a city skyline at random. Reflecting on the content of your imagination, you ask yourself whether the skyline is asymmetrical and rationally judge that it is. As it happens, the imagined skyline you have come up with is a very good representation of the actual skyline of your home town, Chicago. Noticing this, you also rationally judge that the imagined skyline is the skyline of Chicago, and by a simple logical inference you conclude that the skyline of Chicago is asymmetrical. Even though you arrive at this conclusion via the same kind of process as the process employed in conceptual analysis,

---

<sup>32</sup> For more on how we come to learn new fact via conceptual analysis see Balcerak Jackson, M. (2013). A thick notion of understanding is also discussed in Balcerak Jackson, M. / Balcerak Jackson, B. (2012).

<sup>33</sup> Peacocke (1998), 51.

intuitively you did not do conceptual analysis in this case. Your judgment is not just a manifestation of your understanding alone, but also of a recognitional capacity based on perceptual background knowledge. The question for the thick conception of understanding, then, is which ways or instances of evaluating hypothetical scenarios as described above count as manifestations of understanding.

We could refuse to answer the question, and simply allow that any way of evaluating hypothetical scenarios counts. But to do this would be to bring into life a notion of understanding that surely has no interesting relationship to a priority. If we want to avoid this then it is not enough just to describe the structure of the reflection that takes place during conceptual analysis. We need to decide, which are the basic abilities the exercise of which counts as part of understanding, and which basic abilities are not? But, to compile the relevant list in non-arbitrary ways seems to be a very difficult task. How do we decide what is in and what is out?

One way of making a non-arbitrary decision is provided by Chalmers. For him the aspect of meaning that reflects our understanding of the concepts is constituted by those judgments about full descriptions of hypothetical scenarios that manifest ideal rational reflection, which he understands as ideal a priori reflection.<sup>34</sup> Many thoughts will come out as epistemically analytic on this account. As normal reasoners we are not capable of ideal a priori reflection. But we can often approximate it, and when we do then the judgments we make about hypothetical scenarios will plausibly give us knowledge (or at least reasonable belief) of many interesting propositions, including propositions that have philosophical significance. We might be able to discover good theories about knowledge, for example, or about personal identity or the nature of belief by relying on these a priori judgments and our a priori reasoning abilities.<sup>35</sup>

But notice that if we go down this path, we take the notion of a priority to be more fundamental than the notion of understanding; our conception of understanding is itself explained in terms of what we can judge on the basis of a priori reflection. To do this is to effectively give up on the project that motivated us in the first place, the project of

---

<sup>34</sup> See Chalmers (2006).

<sup>35</sup> How we can make epistemic progress on such a picture has been described in Balcerak Jackson, M. (2013).

using epistemic analyticity to help explain and ground a priori knowledge. The epistemically analytic truths *just are* certain truths that we are in a position to know a priori.<sup>36</sup>

#### 4. Moving Forward

Where do we go from here? We have explored various possibilities to explain a priority by relying on a notion of epistemic analyticity that is grounded in a conception of understanding. We have followed paths that construe understanding thinly as linguistic competence, and that look at linguistic research for help in understanding this cognitive ability. We have followed paths that construe understanding thickly and looked at the method of conceptual analysis to elucidate what such a thick conception might entail. But none of the paths we have walked have led us to a conception of understanding and a notion of epistemic analyticity that could meet our expectations. The most solid thin conceptions of understanding leave us with trivial a priori thoughts. The most robust thick conceptions of understanding leave us with an account that puts the cart in front of the horse by taking a priority as explanatorily prior to understanding.

One could object that we have not explored all the possible paths. The space of possible conceptions of understanding is not exhausted by the thin and thick conceptions considered here. And we surely have not explored all the ways in which linguistic practice could lend support to a thin conception of understanding, or all the ways in which we could non-arbitrarily delineate the cognitive abilities that make up a thick conception of understanding. Maybe someone can come up with a moderately chubby conception of understanding that is well-founded and has great explanatory power. But so far, such a conception is not on the horizon. The options we have discussed seem like the most motivated and eligible options that we have.

But more importantly, the discussion of *why* the thin conception of understanding fails in explaining *substantial* a priori knowledge and *why* the thick conception fails in *explaining* substantial a priori knowledge, both tell us something about how we should move forward. I think that the focus on, or sometimes even the obsession with a conception of

---

<sup>36</sup> This is not a problem for Chalmers, who is not in the business of explaining a priority. But it is a problem for the project we are pursuing here.

understanding and the related concept of epistemic analyticity is misguided. It hinders our progress in understanding a priori knowledge rather than helping it. We have seen that there are cognitive capacities that we know, use and study – linguistic competence, supposition, imagination, deductive inference, abductive reasoning etc. – that plausibly play a role in our acquisition of a priori knowledge. We could continue to ask which of these capacities helps constitute understanding, in some good sense of the term, and what sorts of epistemic analyticities might issue from understanding in this sense. But perhaps we should put understanding and epistemic analyticity to the side, and focus on these more basic cognitive capacities themselves, and on their interplay in methods such as conceptual analysis. Perhaps in this way we can arrive at better explanations of how substantial a priori knowledge is not merely possible, but even achievable.

## **Bibliography**

Balcerak Jackson, B. (2009): Understanding Semantic Structure: A Reply to Timothy Williamson, *Proceedings of the Aristotelian Society* 109, 1: 337-343.

Balcerak Jackson, B. (2017): Structural Entailment and Structural Natural Kinds, *Linguistics and Philosophy* 40, 3: 207-237.

Balcerak Jackson, M. (2013): Conceptual Analysis and Epistemic Progress, *Synthese*, 190: 3053-3074.

Balcerak Jackson, M. / Balcerak Jackson B. (2012): Understanding and philosophical methodology, *Philosophical Studies* 161, 2: 185-205.

Boghossian, P. (1996): Analyticity Reconsidered, *Nous* 30, 3: 360-391.

Boghossian, P. (2003): Blind Reasoning, *Aristotelian Society Supplementary Volume* 77: 225-248.

Chalmers, D. (2006): Two-Dimensional Semantics, in: *Oxford Handbook of the Philosophy of Language* (eds. Lepore, E. / Smith, B.), Oxford University Press.

Chalmers, D. / Jackson, F. (2002): Conceptual Analysis and Reductive Explanation, *Philosophical Review* 110, 3: 315-360.

Davidson D. (1967): *The Logical Form of Action Sentences, The Logic of Decision and Action* (ed. Rescher, N.), University of Pittsburgh Press.

Fellbaum, C. (2005): WordNet and wordnets. in: *Encyclopedia of Language and Linguistics* (eds. Brown, K. et al) Second Edition, Oxford: Elsevier, 665-670.

Fodor, J. / Katz, J. (1963): The Structure of a Semantic Theory, *Language* 39, 2: 170-210.

- Fodor, J. / Lepore, E. (1996) *Analyticity Again*, *The Blackwell Guide to the Philosophy of Language* (eds. Devitt, M./Hanley, R.), Oxford University Press.
- Geeraerts, D. (2010): *Theories of Lexical Semantics*, Oxford University Press.
- Goddart, C. / Wierzbicka, A. (2014): *Words and Meanings. Lexical Semantics Across Domains, Languages and Cultures*, Oxford University Press.
- Higginbotham (1985): *On Semantics*, *Linguistic Enquiry* 16, 4: 547-593.
- Parsons, T. (1994): *Events in the Semantics of English*, MIT Press.
- Peacocke, C. (1998): *Implicit Conceptions, Understanding and Rationality*, *Philosophical Issues* 9: 43-88.
- Pietroski, P. (2000): *The Undeclared Domain of Semantics*, *Sats: The Nordic Journal of Philosophy* 1: 161-176.
- Pietroski, P. (2008): *Small Verbs, Complex Events: Analyticity Without Synonymy*, in: *Chomsky and His Critics* (eds. Anthony, L. / Hornstein, N.), Blackwell.
- Rey, G. (2003/2013): *The Analytic/Synthetic Distinction*, *Stanford Encyclopedia of Philosophy*, <https://plato.stanford.edu/entries/analytic-synthetic/>
- Russell, G. (2008): *Truth in Virtue of Meaning. A Defense of the Analytic/Synthetic Distinction*, Oxford University Press.
- Schein, B. (1994): *Plurals and Events*, MIT Press.
- Wierzbicka, A. (1992): *Defining Emotion Concepts*, *Cognitive Science* 16: 531-589.
- Williamson, T. (2007): *The Philosophy of Philosophy*, Oxford University Press.