Lexical Innovation and the Periphery of Language

Luca Gasparri, IJN Paris / FU Berlin | 1st draft, October 2018

Abstract

Lexical innovations (e.g., nonce words and zero-derivations coined on the fly by a speaker) seem to bear semantic content. Yet, such words cannot be attributed semantic content based on the lexical conventions available to the language, since they are not part of its lexicon. This is in tension with the commonplace view that lexical semantic content is constituted by lexical conventions. In recent work, Josh Armstrong has argued that this tension can be alleviated in two ways. The first is to stick to the conventionalist assumption and deny that lexical innovations bear semantic content. The second is to dynamicize the conventionalist assumption, i.e., argue that lexical innovations trigger a rapid update of the lexical conventions of the language and receive semantic content via the updated lexical conventions. Armstrong lays out and defends the second option. I propose a third way: the view that the interpretation of lexical innovations relies on an algorithm which generates a hypothesis about the semantic properties projected by the unfamiliar occurrence and feeds them into sentence meaning with no prior update of the lexical conventions of the language.

1 The Problem

Consider the following sentence.

1) [*Context: the speaker is describing a basketball play*]
John was ten feet from the basket and tried to shaq his way to the rim.

The word "shaq" is a novel zero-derivation created by turning the proper name 'Shaq', one of the several nicknames of NBA superstar Shaquille O'Neal, into a verb.[1] Even with no prior exposure to the

1 I am not aware of previous uses of this verbification, and apologies if anyone else should be credited for the invention. That said, I can reasonably predict that most readers, even those acquainted with the recent history of NBA basketball,

novel verb, basketball aficionados can recover the intended meaning of (1) easily: John established a low post position and attempted to exploit his physical force to bump the defender, approach the rim, and score. Practically speaking, the interpretation of "shaq"$_{(1)}$ is thus a simple task to accomplish, provided the example is presented to a rational listener endowed with the relevant background knowledge about the game of basketball and how Shaquille O'Neal used to play it. Theoretically speaking, however, examples like (1) raise important questions. Some of such questions pertain to cognitive processing: borrowing from the terminology of machine learning, examples like (1) call for an account of the psychological mechanisms that sustain "zero-shot" lexical interpretation in humans.[2] But they also raise foundational questions, as they generate a tension between two independently plausible assumptions about the nature of linguistic meaning. Call them, for brevity, Semanticity and Conventionalism (cf. King and Stanley 2005).

Semanticity The interpretation of natural language sentences rests on the computation of the semantic content of their word-level constituents.

will be encountering the word for the first time. Also, notational caveat: I will use expressions of form *word*$_{CL}$ to notate vocabulary items (lemmas) and their grammatical class, and expressions of form "word"$_{S}$ to notate occurrences of vocabulary items within sentences. Accordingly, I will use *Shaq*$_{N}$ and *shaq*$_{V}$ to refer to the proper name 'Shaq' and the verb 'shaq', respectively, and "shaq"$_{(1)}$ to refer to the occurrence of *shaq*$_{V}$ in (1). This toy notation will be sufficient for the purposes of this paper.

2 In machine learning (esp. computer vision), "zero-shot learning" refers to the processes by which a machine learns how to recognize objects in an image without any labeled training data to help in the classification. Approximating a bit, "zero-shot learning" deals thus with the problem of building machines capable of emulating human-level abilities in the classification of objects belonging to a class they have never encountered before. See, e.g., Xian et al. (2018). Since the interpretation of "shaq"$_{(1)}$ can be aptly characterized as a "zero-shot" task in this sense, I will borrow the label as a quick reminder to the fact that in cases such as (1) listeners are faced with the task of interpreting a word which is not part of their lexicon and is encountered for the first time.

Conventionalism The semantic content of word-level linguistic units is determined by the lexical conventions of the language.

Semanticity suggests that the interpretation of (1) should rely on the computation of the semantic content of its lexical constituents. Conventionalism adds that the lexical constituents of (1) should be assigned semantic content relative to the lexical conventions of English. The conjunction of these two assumptions is unproblematic if measured against sentences containing familiar words ("Mary ate the apple"). However, *shaq*$_{V}$ is not in the lexicon of English. So, contra Conventionalism, "shaq"$_{(1)}$ cannot be attributed semantic content based on the lexical conventions available to the listener. As a consequence, Semanticity and Conventionalism cannot be embraced simultaneously in the analysis of (1). Either "shaq"$_{(1)}$ projects semantic content, as per Semanticity, and the case speaks against Conventionalism, or word-level semantic content can be determined exclusively by the form-to-meaning mappings set by the lexical conventions of the language, as per Conventionalism, and (1) speaks against Semanticity.

We can provide a generalized formulation of the problem. Language users produce lexical innovations and manifest an ability to interpret them.[3] However, it appears that the combination of Semanticity and Conventionalism cannot capture such an ability. Either the zero-shot interpretation of lexical innovations rests on the recovery of their semantic content and lexical innovations provide

3 In what follows, I will use expressions such as "lexical innovations", "lexical novelties", and "unfamiliar words" interchangeably, but always to talk about a specific phenomenon: the spontaneous production of words which are not part of the lexicon of the language. I am making this explicit because labels like "lexical innovation" have (at least) two possible readings: one referring to the production of brand new words (like "shaq"$_{(1)}$), one referring to atypical uses of existing words (e.g., cases where a familiar word is employed with a non-standard or "shifted" meaning). As I will reiterate below, these two senses individuate two distinct objects of analysis (cf. Blank 1999), and this paper is about the former.

counterevidence to Conventionalism, or Conventionalism sets the appropriate boundaries for the determinants of semantic content, and we should withdraw from the Semanticity-inspired assumption that the zero-shot interpretation of lexical innovations rests on the apprehension of semantic content.

The goal of this paper is twofold. The first, programmatic goal is to interface the philosophical discussion on lexical conventions and language change with recent advances in linguistic theory on the semantics of "peripheral" devices of communication like iconic pseudowords and gestures (e.g., Kendon 2004; Lascarides and Stone 2009; Ebert and Ebert 2014; Schlenker 2018), and suggest that a combined investigation of these two domains can lead to fruitful foundational insights on the nature of lexical innovation. The second, specific goal is to use the newly drawn parallel to provide (what I believe is) a simple and attractive non-conventionalist solution to the tension between Semanticity and Conventionalism illustrated above.

The plan is as follows. Section 2 will describe Armstrong's (2016) take on the problem: the view that lexical innovations like "shaq"$_{(1)}$ trigger a rapid update of the lexical conventions of the language and receive semantic content via the updated lexical conventions. Section 3 will draw a parallel between (1) and the zero-shot interpretation of novel onomatopoeic words, iconic pseudowords and iconic gestures, and use the comparison to argue that Armstrong's treatment of the tension can be challenged by an alternative account. Section 4 will lay out the alternative account: the view that the zero-shot interpretation of lexical innovations relies on an algorithm which generates a hypothesis about the semantic features projected by the unfamiliar occurrence and feeds them into sentence meaning with no intermediate update of the lexical conventions of the language. Section 5 will recapitulate and offer some concluding remarks.

2 Solution: Conventions on the Fly

Armstrong (2016) provides a nice formulation of the tension surfacing from the combination of Semanticity and Conventionalism in cases like (1), and maintains that there are two basic ways out. The first way out is to keep Conventionalism in force, suspend Semanticity, and argue that the content recovered upon first exposure to "shaq"$_{(1)}$ is pragmatic. On this view, (1) lacks semantic content and our ability to assign "shaq"$_{(1)}$ its perceived meaning does not run through the formulation of a hypothesis about the semantic features projected by the unfamiliar occurrence. The second way out is to keep Semanticity and revise Conventionalism. On this view, "shaq"$_{(1)}$ does bear semantic content, and the correct diagnosis of the tension cropping up from (1) is that we should improve the current formulation of Conventionalism.

After detailing the two alternatives (keep Conventionalism and drop Semanticity vs. keep Semanticity and revise Conventionalism) Armstrong argues in favor of the second option. He observes that cases like (1) challenge the conventionalist doctrine only if the body of conventions underlying the calculation of lexical semantic content is understood as a *static* inventory. However, there is no need to presuppose that the assignation of semantic content to "shaq"$_{(1)}$ can rely only on the state the lexical conventions available to the listener are in at the time of exposure to (1). He thereby argues that examples like (1) motivate a *dynamic* reinterpretation of Conventionalism. Let us call it, for brevity, Conventionalism*.

Conventionalism* The semantic content of word-level linguistic units is determined by the lexical conventions *encoded or imported on the fly* in the language.

Conventionalism* agrees with its static ancestor on the proposition that word-sized constituents have their semantic content determined relative to the conventions encoded in the lexicon of the language. But it adds that the body of lexical conventions available to a language is not an immovable inventory. Instead, it is flexible, ever changing, and can be updated on the fly to secure fluent communicative interaction with our peers. Once Semanticity and Conventionalism* are combined, the argument continues, the zero-shot interpretive algorithm assigning semantic content to "shaq"$_{(1)}$ can be characterized as follows.[4]

Lexical Expansion

i. The listener hears (1) and encounters "shaq"$_{(1)}$.

ii. The listener infers from the grammatical environment of the word that "shaq"$_{(1)}$ is a verb, searches the lexicon for an entry indicating what semantic content should be assigned to "shaq"$_{(1)}$, but finds no *shaq*$_{V}$ in the lexicon.

iii. The listener retrieves the world knowledge associated to the denotation of *Shaq*$_{N}$, isolates the chunk of world knowledge evaluated to be relevant to the interpretation of "shaq"$_{(1)}$, and updates their lexicon with a new entry specifying the presumed conventional semantic features of *shaq*$_{V}$.

iv. The attribution of semantic content to "shaq"$_{(1)}$ is performed.[5]

4 The label "Lexical Expansion" is mine, but what comes below it is a charitable summary of Armstrong's position. An approach to lexical innovations pressing an even more radical defense of Conventionalism is laid out by Lepore and Stone (2017).

5 A clarification. Although I have appealed and will appeal, for pedagogical convenience, to the received notion that lexical semantic conventions are stored in the lexicon, Lexical Expansion is not necessarily committed to a grammar contemplating an autonomous lexical component. What the account appears to be pressing is the simple claim that on first encounter with a novel word W recognized to be of a grammatical class CL, the listener builds a new semantic

Now, if you are a fan of semantic conventions, Lexical Expansion appears a very good view to have.[6] The thesis improves on static approaches when it comes to giving a promising convention-based explanation of facts about language change in general (Bybee 2015) and lexical semantic change in particular (Ludlow 2014). For example, Lexical Expansion provides an attractive descriptive framework for the processes triggered on first exposure to a semantically unfamiliar use of a familiar word: listeners retrieve the conventional interpretation of the word specified by their lexicon, and then modulate the convention on the basis of what they can infer from the contextual peculiarities of the use at stake. The thesis also reconciles, at least in intent, well-known objections to Conventionalism based on language variation (see, most notably, Davidson 1984; 1986; 1991) with the received picture that

convention *Sem*(W_{CL}) encoding the semantic properties of W_{CL} and assigns semantic features to the occurrence of W by reasoning relative to a model where the language has *Sem*(W_{CL}). Hence, what I take Lexical Expansion to be committed to (and what I will discuss in what follows) is the high-level proposition that the algorithm in charge of the zero-shot interpretation of (1) runs on a (however implemented) prior expansion of the lexical semantic resources of the language. On some approaches, these resources will be encoded in a dedicated storehouse: the lexicon. On other approaches, these resources will be distributed through other components of the grammar. Hence, the discussion below should be meaningful and coherent both for models on which lexical semantic conventions are stored in an autonomous lexicon, and for models of the architecture of the grammar rejecting the notion of an autonomous lexicon, such as Elman (2009) or the syntax-all-the-way-down approach of Distributed Morphology (Halle and Marantz 1993; 1994).

6 It is also the only way I think we can be serious about the idea that episodes of zero-shot lexical interpretation can be given a conventionalist treatment. I would not suffice to insist, alongside Lewis (1969), that there are grounds to argue that conventions can be established even after a single episode of coordination on the meaning of a new word, for this begs the question of how we can press a conventionalist story about the interpretive mechanisms at work on the *first* encounter with that word. Even if the single episode of coordination on the meaning of "shaq"$_{(1)}$ did suffice to establish a convention for *shaq*$_V$, *that* inaugural episode of semantic coordination would remain an exception to the conventionalist principle.

communication is essentially a matter of coordination on shared linguistic conventions.[7] The idea that the zero-shot interpretation of new words is ultimately a matter of rapidly updating one's lexical conventions on the basis of context seems to give us a framework where the processing of linguistic conventions and non-conventional reasoning (e.g., reasoning about intentions and so-called "social triangulation") can peacefully coexist without any either/or lurking in the background (Armstrong 2015).

Notwithstanding these advantages, I believe there are reasons to cast doubt the claim that Lexical Expansion produces the best account of the processes involved in interpreting sentences which, like (1), feature an occurrence of a brand new word. In other words, the high-level view of the dynamic features of the lexical system delineated by Conventionalism* may as well hit the mark and track a real mechanism of language change. Even so, Lexical Expansion is not the only, and probably not the best story we can tell about how a listener with no prior knowledge of $shaq_{V}$ can interpret "shaq"$_{(1)}$, even if we assume that the content recovered for the zero-derivation does not boil down to a purely pragmatic meaning. I will maintain, in particular, that Armstrong's initial aut-aut (either keep Conventionalism and drop Semanticity, or keep Semanticity and revise Conventionalism) overlooks an attractive third way: agree that on the zero-shot interpretation of (1) "shaq"$_{(1)}$ is interpreted semantically, and drop the requirement that the attribution of semantic features to the unfamiliar occurrence should be mediated by the rapid construction of a new lexical convention. On this view, the lesson we should draw from the

7 It should perhaps be mentioned that Davidson took cases of full-blooded neologisms like (1) to be as challenging for Conventionalism as cases of inadvertent malapropisms and slips of the tongue. This was probably a mistake: most cases of inadvertent production errors (e.g., "The police apprehended two *auspicious* individuals"), can be readily reconciled with Conventionalism by hypothesizing that listeners perform a plausibility inference to reconstruct the word the speaker was attempting to pronounce ("suspicious"), reverse engineer the intended sentence ("The police apprehended two *suspicious* individuals"), and interpret the reverse-engineered sentence by purely conventional means. A theory of malapropisms in this spirit can be found in Predelli (2010).

dilemma raised by (1) is not that we should dynamicize Conventionalism, but that Conventionalism should be suspended in the analysis of the zero-shot interpretation of lexical innovations.

3 Widening the Landscape

To drive the point home, we need to reexamine Armstrong's case for Lexical Expansion, and work our way to the alternative account of the zero-shot interpretation of (1) just hinted at. Let me start with two observations.

The first observation (actually a brief programmatic reminder) is that a meaningful analysis of the tension surfacing from (1) should not lose sight of the richness and the productivity of the lexical system. It is widely accepted that lexical entries are informationally complex entities which impose various requirements on the structures in which they find themselves. Think of argument structure and of the distributional phenomena grouped under the notions of S-selection and C-selection in textbook linguistic theory (e.g., Fromkin 2000). But most importantly, the lexicon is a generative engine (a.o., Bauer 2001; Jackendoff 2002). Suppose the lexicon of a speaker K contains an entry for the prefix *out-*, an entry for the verb *populate*, and the inflectional suffix *-s*, but does not feature an entry for *outpopulate*. Suppose also K encounters for the first time the simple sentence "India outpopulates Japan". Even in the absence of an entry specifying the conventional meaning of $outpopulate_{V}$, it seems safe to assume that K's lexicon has all the resources it takes to determine that the sentence features a complex inflected verb made of three familiar constituents (the prefix *out-*, a root borrowed from *populate*, the inflectional suffix *-s*), combine compositionally the semantic features of these constituents, and handle the innovation with no need to compute any new lexical convention. To be sure, these mechanisms have a limited scope of application and do not help account for cases where the

listener faces the task of interpreting semantically "primitive" lexical innovations.[8] But it is important to keep in mind that the lexicon is considerably richer and more productive than a mere list of form-to-denotation pairs.

The second, more at-issue observation is that the interpretive algorithm posited by Lexical Expansion is stipulative, and potentially redundant. The main issue is the following. Once the grammatical class and the inflectional characteristics of "shaq"$_{(1)}$ have been identified, once the interpretive system has isolated the bits of world knowledge that are relevant to produce a theory of the action designated by "shaq"$_{(1)}$, and once such a hypothesis is formulated, why would the interpretive system integrate it (together with grammatical class) in a novel lexical entry *shaq*$_{V}$ and then retrieve it again to interpret (1) via conventional reasoning, instead of feeding it directly into the content recovered for the sentence? While we could not conceive of an interpretive algorithm for (1) not contemplating the inspection of the world knowledge possessed by the listener about Shaquille O'Neal,[9] the formulation of a hypothesis about the conventional semantic features of a putative lemma *shaq*$_{V}$ seems stipulative in an explanation of how a listener might zero in on a theory of the semantic features projected by "shaq"$_{(1)}$. To be sure, the move offers a handy way of generalizing Conventionalism to the realm of sentences featuring novel words. But at this point we cannot take it as antecedently evident that our resolution of the tension raised by (1) should be a chapter of Conventionalism. The key question, thus, is whether there are *independent* reasons to believe that the

8 E.g., a new verbal element that, like "shaq"$_{(1)}$, cannot be analyzed as the combination of known parent lexical parts via semantically productive mechanisms of morphological synthesis. More on this below.

9 As a matter of fact, cooperative lexical innovation tends to rely on an implicit theory of the informational resources required to interpret the new word, and on an assumption that such resources are available to the listener. For example, the production of (1) seems to be governed by an underlying assumption that the addressee is familiar with Shaquille O'Neal and associates him to the properties that are relevant to the interpretation of "shaq"$_{(1)}$ (being physically imposing and bumping defenders under the rim) in an almost paradigmatic or stereotypical fashion.

zero-shot interpretation of "shaq"$_{(1)}$ rests on the rapid formulation of a hypothesis about the conventional semantic features of *shaq*$_{V}$, absent clear indications that the informational resources fueling the process should contemplate this step.

Suppose one replies that the rapid formulation of a hypothesis about the conventional semantic features of a putative lemma *shaq*$_{V}$ occurs as a result of the productive role played by the morphology in the interpretation (and the production) of the zero-derivation. On first exposure to (1), the interpretive algorithm creates a new lexical convention for the lemma *shaq*$_{V}$ because of the particular reasoning involved in grasping that the novel word is the verbal counterpart of a known vocabulary item (*Shaq*$_{N}$): switching grammatical class while keeping word form constant is clearly a function which takes in a lexical entry and outputs another lexical entry. Once, the reply continues, the system understands that "shaq"$_{(1)}$ is designed to be interpreted as an occurrence of a verb derived from a preexisting name though a process of change in grammatical class without change in form, it promptly postulates a lemma *shaq*$_{V}$ for a verbal element with known morphosyntactic characteristics and semantic features to be defined, defines the semantic features of the lemma, and then moves on to the interpretation of the unfamiliar occurrence.

This line of argument is an open option for (1), but faces an immediate issue of generality. Lexical Expansion is supposed to give us a *general* model of how we come to terms with lexical innovation. However, words coined on the basis of morphological conversion are just one type of lexical innovation, and the rejoinder does not apply to cases where the novel word is not generated based on familiar lexical material. Lexical innovations that have no vocabulary-level parent and are introduced by inference on properties in the domain of sensory cognition, such as new onomatopoeic verbs, offer an instructive comparison. Consider (2).

2) Lea saw the girls jump on their motorbikes and *schwoom* through the alley.

Listeners exposed to (2) appear able to identify a plausible meaning for the occurrence of the novel verb simply by analyzing its shape and the sentential environment where it occurs. Since linguistic context and reasoning about the shape of the word do all the required semantic work (modulo the identification of grammatical class and inflectional characteristics), it is highly plausible that the zero-shot interpretation of "schwoom"$_{(2)}$ is performed by an algorithm that generates a candidate meaning for the occurrence and plugs it into the content recovered for (2) without first updating the lexical semantic conventions of the language. The example, thus, raises concerns parallel to (1): it is unclear why one should postulate an interpretive routine where the zero-shot interpretation of "schwoom"$_{(2)}$ is achieved via the formulation of a hypothesis about the conventional semantic features of *schwoom*$_{V}$.[10] But unlike (1), the introduction of the new verb of (2) does not rely on the conversion of any preexisting lexical material, so the formulation of a hypothesis about the conventional semantic features of *schwoom*$_{V}$ cannot be justified on the basis of morphological considerations like those

10 This is not to say that the two examples are not remarkably different. While the reconstruction of the action designated by "schwoom"$_{(2)}$ rests on context plus an iconic analysis of the shape of the verb, the interpretation of "shaq"$_{(1)}$ requires context and a more complex plausibility inference. The listener has to determine what properties, among those they associate to Shaquille O'Neal, provide the correct blueprint for the interpretation of the verb (physical dominance and style of play as opposed to having an eccentric media personality). But this difference is orthogonal to the claim that the mechanisms in charge of identifying a plausible meaning for "shaq"$_{(1)}$ and "schwoom"$_{(2)}$ share the macrostructural property of zeroing in on a hypothesis about the meaning of the occurrence without first fixing a new convention. Besides, recall that the properties of the shape [ʃwum] are insufficient to fix any univocal meaning and that, much like in (1), reasoning about context has to pull its weight. In the sentence "The stocks were going up yesterday but today they'll most certainly *schwoom*", the verb would be readily understood to mean something very different from the event pictured by "schwoom"$_{(2)}$ (a fall instead of a horizontal movement). So (1) and (2) are not as distant as too much emphasis on iconicity might lead to believe.

illustrated above.[11] Bottom line: we still need to be given a specific reason why the zero-shot interpretation of sentences like (2) and, retrospectively, (1) should run through the rapid creation of a new semantic convention for the unfamiliar word.

In one key turn of his argument, Armstrong (2016) does offer a recipe to answer the issue. As we have seen, the argument starts off with the claim that the tension between Semanticity and Conventionalism raised by (1) can be alleviated in one of the following ways: classify "shaq"$_{(1)}$ as a constituent that receives only a pragmatic meaning (which would amount to claiming that lexical innovations are not in the domain of Semanticity), or allow it to project semantic features on the basis of a rapid update of the lexical resources of the language, as per Conventionalism* plus Lexical Expansion. However, Armstrong (2016: 95-96) observes, the idea that the content projected by "shaq"$_{(1)}$ is not in the domain of Semanticity suffers from difficulties. In particular:

a) due to the absence of *shaq*$_{V}$ in the lexicon, "shaq"$_{(1)}$ has no literal meaning, so the content it contributes to (1) cannot be derived pragmatically from any prior literal meaning;

b) (1) does not trigger judgments of ungrammaticality;

11 Another option would be to argue that the identification of the syntactic and inflectional features of "schwoom"$_{(2)}$ suffice to prove that the listener is performing implicit reasoning about the properties of a new lemma *schwoom*$_{V}$, since it is impossible, e.g., to compute inflectional features without *ipso facto* computing the inflectional root they modify. But the matter is more complex. The fact that listeners assign "schwoom"$_{(2)}$ the morphosyntactic features of an infinitive might as well run through the implicit inference of a new inflectional root *schwoom*$_{V}$, but this is insufficient to conclude that the *meaning* of "schwoom"$_{(2)}$ is individuated thanks to an algorithm which enriches the lexicon with a hypothesis of the conventional semantic features of the inflectional root *schwoom*$_{V}$. Once again, one can tell a simpler story: the algorithm posits an inflectional root *schwoom*$_{V}$ to fix the morphosyntax of "schwoom"$_{(2)}$, but then, instead of generating a candidate meaning for "schwoom"$_{(2)}$ via reasoning about the context-insensitive semantic features of the vocabulary item *schwoom*$_{V}$, it focuses directly on the formulation of a hypothesis about the semantic features projected by the occurrence "schwoom"$_{(2)}$ in the specific context of (2).

c) “shaq”$_{(1)}$ interacts compositionally with the rest of (1).

Because of (a-c), the argument continues, the contribution made by “shaq”$_{(1)}$ to the overall meaning of (1) cannot be understood as purely pragmatic in kind. Hence, “shaq”$_{(1)}$ projects semantic content. But the assignment of semantic content to the lexical parts of a sentence is normally governed by their status in the lexical conventions of the language. Hence, we should embrace Lexical Expansion. This is the key turn: the heuristics in (a)-(c) suggests that “shaq”$_{(1)}$ projects semantic content. And since we have antecedent evidence that Conventionalism provides explanatory purchase on the projection of lexical semantic content, we have grounds to generalize the premise to lexical innovations and claim that “shaq”$_{(1)}$ projects semantics features relative to the rapid formulation of a hypothesis about the conventional meaning of *shaq*$_{V}$.

The reasoning is clear, and certainly gives *prima facie* reasons to think that Lexical Expansion is on the right track. But now consider (3) and (4).

3) Great. Now Mark’s mood will [*iconic pseudoword*: [u]-SOUND DESCENDING IN PITCH] and stay low for hours.
4) Great. Now Mark’s mood will [*pro-speech gesture*: HAND WITH PALM UP + ARM DESCENDING DIAGONALLY] and stay low for hours.[12]

We mentioned that (2) cast doubt on Lexical Expansion because it presented a case where the attribution of semantic features to a novel word seemed possible with no need of any rapid update of

12 A note on the terminology: pro-speech gestures are gestures that replace entire words. They are distinguished from co-speech gestures, which are produced simultaneously with the spoken words they modify, and from post-speech gestures, which follow the expressions they modify.

the lexical conventions of the language, and led us to the hypothesis that the interpretation of (1) could be understood to follow the same macrostructural path. (3) and (4) present a strengthened variant of the same problem. Once again, the interpretation of the pro-speech gesture and of the pseudoword rely, plausibly, on mechanisms outside the purview of a theory of the lexical semantic conventions of the language, due to their iconic nature. The meaning they contribute to the sentences they interact with is not calculated based the inspection of the appropriate set of conventions, but rather recovered from the combination of sentential context and the information afforded by the physical makeup of the sign. So, as it happened with (2), it is unclear why the zero-shot interpretation of (3) and (4) should contemplate a stage where listeners update their lexicon with a hypothesis about the conventional semantic features of a new lemma-like unit. The additional peculiarity of these examples is that there is simply no way the interpretation of (3) and (4) may be mediated by an update of the lexical conventions presupposed by the listener, due to the non-lexical nature of [u]-SOUND DESCENDING IN PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$.[13] While (2) could, in principle, be interpreted via the

13 One might object that we are ignoring the possibility that the zero-shot interpretation of [u]-SOUND DESCENDING IN PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ is in fact a form of covert lexical reasoning because the pseudoword and the pro-speech gesture are implicitly treated as codes for familiar words. On this view, [u]-SOUND DESCENDING IN PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ are implicitly construed as placeholders for "descend" or "drop down", and their interpretation is taken over by normal linguistic computation. However, the hypothesis conflicts with evidence that pro-speech gestures and iconic pseudowords convey gradient iconic information that cannot be emulated with standard lexical items. For example, the modulation of the angle and the amplitude of the movement of the arm in (4) can be used to trigger fine-grained inferences about the speed and the magnitude of the change in Mark's mood (Mark will slowly go from neutral to sad vs. Mark will rapidly go from euphoric to utterly depressed). The fact that modulations of this sort cannot be reproduced with standard lexical items militates strongly against the claim that SOUND DESCENDING IN PITCH"$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ can be analyzed as modes of non-phonological externalization of lexical words. See Schlenker (to appear) for further discussion on this point.

rapid formulation of a hypothesis about the conventional semantic features of the lemma *schwoom*$_{V}$ (the issue was that the move seemed stipulative and could be challenged by an alternative story), Lexical Expansion looks like a nonstarter when it comes to capturing (3) and (4).

But wait, one might complain: why bring in gestures and pseudowords? Precisely because (1) and (2) contain a novel lexical occurrence, whereas the salient constituents of (3) and (4) are non-lexical, it would be unfair to imply that because Lexical Expansion has a hard time illuminating what happens on first encounter with (3) and (4), then it should be ruled out as a palatable explanans for the zero-shot interpretation of "shaq"$_{(1)}$ and "schwoom"$_{(2)}$. (3) and (4) would cast a shadow on Lexical Expansion only on the premise that the interpretive algorithm it illustrates can be justifiably asked to apply, besides to (1), also to the way interpreters come to terms with "peripheral" constituents such as iconic gestures and iconic pseudowords. But the story did not aim to apply to gestures and pseudowords in the first place: it was motivated by a desire to illuminate "core" instances of linguistic innovation like those found in (1) and (2), while remaining agnostic on whether non-linguistic innovations like those observed in (3) and (4) triggered the same interpretive algorithm.

The reaction would be perfectly sound. And yet, the parallel with (3) and (4) may have something to teach us. Recall the diagnostics in (a)-(c). The argument was that because "shaq"$_{(1)}$ satisfies (a)-(c), the novel zero-derivation projects semantic features, and since we have antecedent motivations to believe that the projection of semantic features for word-sized constituents relies on the lexical conventions of the language, we should dynamicize Conventionalism and appeal to Lexical Expansion. However, if we test (3) and (4) against the same heuristics, it turns out that [u]-SOUND DESCENDING IN PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ *do* satisfy (a-c). Both [u]-SOUND DESCENDING IN PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ contribute content which cannot be calculated based on a prior literal meaning (a), they do not trigger judgments of ungrammaticality (b), and they seem to interact with the grammatical structure of the respective

sentences under the principles of compositionality (c). For comparison, see (5), which manages to introduce an antecedent for modal anaphora, (6), which successfully embeds the whistle under negation, and (7), where quantification is perfectly in order.

5) Had Mark kept allowing his mood to [*iconic pseudoword*: LONG DESCENDING WHISTLE], Bea would have left him.
6) Bill's mood did not [*iconic pseudoword*: LONG DESCENDING WHISTLE]. But he wasn't in his best shape either.
7) Whenever somebody wears Lucy's hat, everybody else's mood [*iconic pseudoword*: LONG DESCENDING WHISTLE].

In short, if we test (3) and (4) against the same criteria set by Armstrong to argue that the zero-derivation of (1) is in the scope of Semanticity (and thereby should be characterized as recommended by Lexical Expansion), the pseudowords and the gesture of (3) and (4) turn out to be in the domain of Semanticity as well.[14] However, it *remains* clear that the mechanisms whereby these constituents

14 Once again, one might object that this conclusion is a figment of (a)-(c), and that a more careful heuristics might do the job of disentangling (1) from (3) and (4) relative to the issue of the semantic status of the relevant constituents. However, there is growing consensus that iconic constituents like [u]-SOUND DESCENDING IN PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ do trigger sophisticated semantic effects, even on first exposure. Besides standard entailments, they seem to trigger scalar implicatures, presuppositions and associated phenomena (e.g., anti-presuppositions), homogeneity inferences that are characteristic of definite plurals, as well as some expressive inferences normally found in pejorative terms (see, a.o., Abush 2012; Greenberg 2013; Ebert and Ebert 2014; Schlenker, to appear). Notice that I am not describing a homogeneous front, as research on these subjects is its infancy. For example, Hunter (to appear) takes issue with Schlenker's (2018) view that co-verbal gestures trigger conditional presuppositions and argues that we are not yet in a position to make any general claim about their at-issue status. But these disputes seems to be pursued in the context of a broader agreement: constituents like [u]-SOUND DESCENDING IN

receive a semantic interpretation are unlikely to rely on any prior update of the semantic conventions of the language. (3) and (4) are thus counterexamples to the thesis that the inspection of the body of semantic conventions available to a listener at the time of interpretation of an utterance is the only possible source of semantic meaning for its word-sized parts. Absent this intermediate premise, there is no logical consequence between the assumption that the content projected by the occurrence of a new word-level constituent is in the domain of Semanticity and the view that such a content should be analyzed within a framework appealing to the rapid formulation of a hypothesis about conventional meaning (as per Lexical Expansion). This is the relevant lesson we can learn from (3) and (4): word-level constituents can project complex semantic properties in zero-shot settings even if the body of semantic conventions available to the listener remain silent about them throughout the interpretive process. Pace Armstrong's initial aut-aut, it is not necessarily the case that either the novel zero-derivation of (1) gets semantic features relative to the rapid formulation of a hypothesis about the conventional semantic features of $shaq_{V}$, or it receives only a pragmatically derived meaning. The comparison with (3) and (4) reveals a third option in logical space: converge on the proposition that "shaq"$_{(1)}$ satisfies Semanticity and put Conventionalism on hold with a theory on which the possible sources of word-level semantic content extend beyond the processing of lexical conventions.

4 A Third Way

Now that we have all the required ingredients, we are ready to formulate our alternative take on the problem of lexical innovation. The recipe is as follows. "Shaq"$_{(1)}$ contributes semantic features to (1) and the sentence is not semantically gappy at the position occupied by the novel zero-derivation, despite the absence of any entry specifying the conventional semantic features of $shaq_{V}$ in the lexicon

PITCH$_{(3)}$ and HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$ project complex semantic effects whose analysis deserves the full might of formal semantics.

of the listener. However, the zero-shot interpretation of (1) does rely on any prior expansion of the lexical resources of the language, and does not run through the rapid construction of an *ad hoc* lexical convention for the lemma *shaq*$_{V}$. Instead, we can hypothesize that the zero-shot interpretive mechanisms assigning semantic features to "shaq"$_{(1)}$ are built on the same macrostructural blueprint as the mechanisms underlying our ability to assign complex semantic features to novel onomatopoeic words (e.g., "schwoom"$_{(2)}$), novel iconic pseudowords (e.g., [u]-SOUND DESCENDING IN PITCH$_{(3)}$), and novel pro-speech gestures (e.g., HAND WITH PALM UP + ARM DESCENDING DIAGONALLY$_{(4)}$). The hypothesis, in particular, is that on first exposure to (1), the attribution semantic features to "shaq"$_{(1)}$ is performed by an algorithm which searches the world knowledge associated to *Shaq*$_{N}$, generates a hypothesis about the semantic features projected by "shaq"$_{(1)}$ in the specific context of (1) and feeds it into the semantics of the host sentence without the intermediate formulation of a theory about the conventional semantic features of a new lemma *shaq*$_{V}$. The hypothesis is, effectively, that Lexical Expansion should be amended as follows (notice the difference at step (iii)).

Productive Periphery

i. The listener hears (1) and encounters "shaq"$_{(1)}$.

ii. The listener infers from the grammatical environment that "shaq"$_{(1)}$ is a verb, searches the lexicon to check if it contains an entry indicating what semantic content should be assigned to "shaq"$_{(1)}$, but finds no *shaq*$_{V}$ in the lexicon.

iii. The listener retrieves the world knowledge associated to the denotation of *Shaq*$_{N}$, isolates the chunk of world knowledge evaluated to be most relevant to the interpretation of "shaq"$_{(1)}$, and builds a hypothesis about the situated semantic features projected by "shaq"$_{(1)}$ in (1) and maps them onto the constituent without first encoding them into a lexical convention for *shaq*$_{V}$.

iv. The attribution of semantic content to "shaq"$_{(1)}$ is performed.

The departure from Lexical Expansion is clear enough. On Lexical Expansion, the exposure to (1) triggers the construction of a novel lexical entry *shaq*$_{V}$ specifying the semantic features of the lemma and the attribution of semantic content to "shaq"$_{(1)}$ is performed via *shaq*$_{V}$. The procedure involves, crucially, the rapid formulation of a theory about the semantic features of the vocabulary item *shaq*$_{V}$, and the assignation of semantic properties to "shaq"$_{(1)}$ relative to such a theory. On the account I am encouraging to consider, the attribution of semantic features to "shaq"$_{(1)}$ is not mediated by any vocabulary-level middle man. Once the morphosyntactic features of the constituent have been extracted from the linguistic environment, and the interpretive machinery has identified what properties, among those known about the referent of *Shaq*$_{N}$, provide the correct blueprint for the identification of the semantic features projected by the unfamiliar occurrence in the specific environment of (1), the machinery maps them onto "shaq"$_{(1)}$ and produces a model of the situated semantic properties of the constituent without first updating the lexical resources of the language.[15]

Of course, the third way I am encouraging to consider is fairly abstract (just as abstract as Lexical Expansion itself, as far as I can tell), and leaves open important questions. What are the specifics of the

15 Nota bene: I am not trying to argue that the interpretation of novel words cannot, as a matter of principle, follow Lexical Expansion. Prior reasoning about lexical conventions *can* happen, especially in cases where context, world knowledge and the concurrent sentential material are insufficient to suggest a strong candidate meaning for the unfamiliar word. If a speaker produces out of the blue the sentence "Tom was late, he had to *koot*", a perfectly natural reaction on the listener's part would be to pause, ask "what does *koot* mean?", gather more information about what the unfamiliar word might mean conventionally, and then proceed to the interpretation of the sentence. The proposal I am making by introducing Productive Periphery is that Lexical Expansion should be rejected as a *universal* and *generic* model of zero-shot lexical interpretation, though it may successfully capture the interpretive dynamics triggers by some instances of lexical innovation (and thus hold *existentially*).

inferential steps allowing to move from the body of world knowledge possessed by the listener about Shaquille O'Neal to an hypothesis about the situated semantic properties "shaq"$_{(1)}$ is understood to project?[16] Does the narrative do justice to the intuition that the first exposure to "shaq"$_{(1)}$ establishes a lexical precedent which can undergo full-fledged conventionalization? What is the interaction between Productive Periphery and known constraints on lexical innovation, such as lexical competition and the resistance to new function words? These questions are certainly worth investigating, but nothing in the account we are considering seems to stand in the way of giving them good answers. The second question, for example, might be answered in the affirmative by highlighting that nothing in the proposed account makes it especially problematic to produce a story about conventionalization, and that it is far from evident that a good theory of how full-fledged conventional coordination on the meaning of a new content word is achieved within a group of speakers should rely on the analysis of (1) given by Lexical Expansion. Lexical Expansion argues that the creation of a proto-convention for *shaq*$_{V}$ is a *precondition* for the zero-shot interpretation of "shaq"$_{(1)}$. On Productive Periphery the acquisition of the proto-convention is not a required stage of the zero-shot interpretation of "shaq"$_{(1)}$, but occurs as a *result* of the episode of coordination involved in the interpretation of the unfamiliar occurrence: the listener interprets "shaq"$_{(1)}$ and the generated interpretation becomes (modulo the external feedback received on the interpretive attempt) the default semantic blueprint for the interpretation of any subsequent occurrences of the verb they might encounter.

Naturally, the claim that the mechanisms which plausibly govern the interpretation of new iconic pseudowords and pro-speech gestures help us shed light on the mechanisms whereby language users cope with non-iconic lexical novelties like "shaq"$_{(1)}$ should be interpreted with due nuance. I am

16 In Rayo's (2013) terminology: how do we move from the "grab bag" of mental items (memories, mental images, pieces of encyclopedic information, pieces of anecdotal information, and so forth) we can associate to Shaquille O'Neal to a hypothesis about the meaning of "shaq"$_{(1)}$?

certainly not trying to imply that the interpretive algorithms at work in the two domains are identical twins and obliterate the received distinction between iconic and arbitrary signs.[17] Quite the contrary: there are considerable differences. For example, the interpretation of iconic gestures relies on an integration of information from the visual modality that is absent in cases of non-iconic lexical novelty (the same goes for pitch with the pseudoword of (3)). Furthermore, the interpretation of pro-speech gestures relies on mechanisms that often make it possible to have a precise and considerably fine-grained understanding of the meaning of the gestures at hand. Listeners exposed to cases like (1), by contrast, seem to match the newly encountered words with a rather coarse-grained meaning. Even when elements of sound symbolism are available, as in the case of "schwoom"$_{(2)}$, the meaning we are able to recover on first exposure to the new word is admittedly loose: a fast horizontal movement possibly producing a certain type of sound.[18] The proposal I am making accepts these differences, but invites to consider one major point of potential commonality across the iconic/non-iconic divide: in essence, the property of producing a hypothesis about the semantic features projected by the occurrence of an unfamiliar sign without running a prior update of the conventional resources of the language.

17 Though it should be mentioned that over the last decade or so linguists have been steadily shifting away from the received orthodoxy that while iconicity can sometimes be found in spoken languages, phenomena such as onomatopoeia and sound symbolism are "asterisks to the far more important principle of the arbitrary sign" (Pinker 1999: 2). A flurry of findings suggests that iconic form-to-meaning mappings are much more pervasive (in signed and spoken lexicons alike) than the dogma of arbitrariness suggests, and that elements of iconicity tend to be present, in various degrees and at various levels (phonemic, prosodic, morphological), even in the linguistic forms we have learned to characterize as "arbitrary". See, e.g., Perniss, Thompson and Vigliocco (2010).

18 Recall, however, that semantic coarse-grainedness seems to be the norm even in our communicative transactions employing conventionalized linguistic material. Hear, e.g., Yalcin (2014: 24): "even allowing that linguistic communication does centrally involve some degree of coordination on items of content [...], it remains open that coordination on items of content is a highly approximate, more-or-less affair, with perfect coordination on content not being especially important, and rarely or never happening".

Notice that the account, besides offering a viable approach to (1), can be extended smoothly to more stereotypical cases of lexical acquisition. Consider (8).

8) [*Context: Louis is French and has an elementary knowledge of the English vocabulary. He is visiting his American friend Sue. Sue and Louis have baked a cake and have left it to cool down in the kitchen. They are sitting at a table in the living room. On the table there are three objects: a blue napkin, a long knife, and a wooden tray. Sue says:*]
Time to cut the cake: grab the *slicer* and follow me.

Suppose Louis is completely unfamiliar with the word *slicer*$_{N}$ but is competent enough to parse the rest of the sentence, form correct beliefs about the grammatical class of "slicer"$_{(8)}$, understand that the odd occurrence is intended to pick out Sue's knife, and interpret the sentence accordingly. Save for the systematic desideratum of generalizing Conventionalism, I see no reason why the understanding that "slicer"$_{(8)}$ picks out Sue's knife should feature a prior update of Louis's theory of the lexical conventions of English (or: why Louis would need to generate the conclusion that "slicer"$_{(8)}$ picks out Sue's knife via implicit reasoning that *slicer*$_{N}$ is a term denoting knives). The interpretation of "slicer"$_{(8)}$ by Louis will, in all likelihood, *produce* such an update: following the interpretation of (8), a temporary (and presumably underspecified) form-to-meaning pair *slicer*$_{N}$ will be added to the arsenal of semantic conventions available to Louis and become the default blueprint for the interpretation of any subsequent occurrences of the noun he might encounter. But nothing in the recovery of a complete semantic template for (8) seems to *presuppose* the update. Whatever prediction Louis might form about the semantic features of "slicer"$_{(8)}$ based on the grammatical environment of the occurrence, the candidate referents available in the context, their salience in Louis's attentional state, the context set and the question(s) under discussion, on the current hypothesis such prediction will be fed directly into

the meaning recovered for the sentence without being first encoded in a dedicated semantic convention for the novel word.

5 Conclusion

The discussion has proceeded as follows. Section 1 introduced Semanticity and Conventionalism, and illustrated the tension surfacing when their combination is tested against sentences featuring lexical innovations. Section 2 described Armstrong's (2016) solution to the tension: the view that zero-shot lexical innovations are interpreted semantically on the basis of a rapid update of the lexical conventions of the language. Section 3 reviewed the case for Lexical Expansion and, building on a comparison with the interpretation of novel iconic verbs, iconic pseudowords and pro-speech gestures, drew attention to the possibility of an alternative analysis of the tension. Section 4 laid out the alternative analysis: the view that the zero-shot interpretation of lexical innovations relies on an algorithm which formulates a hypothesis about the situated semantic features projected by the unfamiliar occurrence without first updating the lexical conventions of the language.

Let me conclude with two final comments. First, a key component of Productive Periphery is the idea Conventionalism should be suspended in the analysis of the semantic effects of lexical innovation. But it is important to emphasize that the suspension of Conventionalism encouraged by Productive Periphery is no way enemy to the proposition that in cases *not* involving lexical innovation ("Mary ate the apple") the recovery of semantic content is successfully characterized by Conventionalism and, accordingly, that the principle captures the "normal" *modus operandi* of our linguistic transactions. Davidson (1984; 1986; 1991) and a line of subsequent commentators have certainly assumed that the fate of Conventionalism hinged on its possibility to generalize to cases of lexical innovation. If the principle turned out to fail to illuminate what happens when we encounter novel words, then the

conclusion to draw was that Conventionalism is false: lexical conventions do not provide any explanatory purchase on an account of linguistic meaning and communication. But I believe that this assumption is mistaken. There is no obvious inconsistency between the claim that Productive Periphery offers a sensible approach to the semantic aspects of zero-shot lexical innovation, and the claim that Conventionalism is a sound principle to hold for sentences and constituents built on familiar linguistic material. As long as the line we seem able to draw between "standard" and "innovative" linguistic production tracks two distinct facets or spheres of language use (each responding to distinct principles and requiring a distinct analysis), our inability to generalize Conventionalism to the realm of lexical innovation should pose no harm to its validity for the run-of-the-mill aspects of language it was originally designed to capture.

Second, my goal in the discussion above has been to argue the anti-conventionalist policy of Productive Periphery is a contender to Lexical Expansion when it comes to capturing the zero-shot interpretation of (1) assuming that the process should feature the attribution of situated semantic features to "shaq"$_{(1)}$. The move was fueled, a.o., by considerations of charity for Armstrong's argument, by the growing consensus around the idea that the innovations found in (3) and (4) warrant a distinctively semantic analysis, and by the intention to show that Lexical Expansion is not a necessary consequence of the assumption that "shaq"$_{(1)}$ bears semantic content. However, in thinking about what happens upon first exposure to (1) outside this theoretical setup, there is room for skepticism that Semanticity is such a solid assumption to hold with respect to constituents like "shaq"$_{(1)}$ and "slicer"$_{(8)}$. Take (8): even if Sue's utterance of "slicer"$_{(8)}$ were covered by a sudden burst of noise in the room or replaced by an unintentional cough (or by anything else we would resist giving any semantic analysis), Louis would probably still be able to infer that Sue wants him to pick up the knife. The same seems to hold in (9).

9) After yesterday's class Rebecca was exhausted, but today she's feeling much less *sloomey*.

Thanks to the richness of the cues set by the prior sentential context, there is a strong possibility that at the point of encounter with "sloomey"$_{(9)}$ the listener has already formulated a robust hypothesis about what the upcoming word might communicate, and simply sticks to such a prediction without venturing in any additional reasoning about the semantic features projected by the unfamiliar adjective.[19] So, to put the point in Kaplanese: what if the zero-shot interpretation of "slicer"$_{(8)}$ and affiliates goes straight to occasion-specific content (i.e., that "slicer"$_{(8)}$ picks out the knife on Sue's table) without the intermediate formulation of an implicit hypothesis about projected character (i.e., that "slicer"$_{(8)}$ denotes a type of knife)?

One might answer that since the recovery of intended meanings is generally the goal of the interpretation of natural language sentences, it stands to reason to assume that the interpretation of unfamiliar words should as well involve the implicit entertainment of a character-like meaning. However, this is at best a *prima facie* justification. The observation that lexical innovations are not part of the communal lexicon can be used to resist the premise, and argue that while Semanticity does capture the ordinary regimen of lexical interpretation, occurrences of words which are not represented in the linguistic lexicon are handed over to general-intelligence processes which, based on context, world knowledge and the like, pair them to an occasion-specific content without making any implicit hypothesis about their semantic status. Another possibility would be to refrain from any generalization and argue that examples are going to differ in this regard on a case by case basis: while some instances of lexical innovation are processed via an hypothesis about projected semantic character (like (1) and possibly (2)), others are not (like (9) and possibly (8)); it all depends on context and on the richness of

19 The mechanisms involved in the use of information from the sentence context to predict the meanings of upcoming words have been studied extensively in psycholinguistics. See, e.g., Kamide, Altmann and Haywood (2003).

the cues set by the concurrent sentential material. In any case, caution is advised. I do not think that, at the moment, we can definitively rule out that the zero-shot interpretation of "shaq"$_{(1)}$ and "slicer"$_{(8)}$ consists of the direct recovery of a fine-grained situated meaning which does not run through anything semantic in kind.[20] Were that the case, we would have both another reason to resist the narrative produced by Lexical Expansion (since we would no longer have to explain how new words might pull the magic of projecting semantic features even if they are not in the lexicon), an immediate dissolution of the tension that gave rise to our discussion (since Semanticity would be suspended), *and* a reason to consider an even more pragmatically oriented version of Productive Periphery on which the zero-shot exposure to lexical innovations trigger nothing but the direct recovery of occasion-specific content.

20 Criterion (a) in Armstrong's (2016: 95-96) heuristics, for example, seems to press another dispensable aut-aut. It is true that the absence of a literal meaning for "shaq"$_{(1)}$ in English makes it hard to claim that its contributed content is recovered via the processes familiar to linguistic pragmatics. But this is insufficient to conclude that we should made do with semantics. We can agree with the part of the argument that tells us that the interpretation of "shaq"$_{(1)}$ cannot rest on linguistic pragmatics, and argue that interpreters fix an occasion-specific meaning for "shaq"$_{(1)}$ entirely by exercising sensitivity to context and common sense (despite the linguistic nature of the constituent, and precisely because the lexicon of the language is silent about it).

References

Abusch, D. 2012. Applying discourse semantics and pragmatics to co-reference in picture sequences. In *Proceedings of Sinn und Bedeutung* 17. https://semanticsarchive.net/sub2012/Abusch.pdf

Armstrong, J. 2015. Coordination, triangulation, and language use. *Inquiry* 59: 80-112.

Armstrong, J. 2016. The problem of lexical innovation. *Linguistics and Philosophy* 39: 87-118.

Bauer, L. 2001. *Morphological Productivity*. Cambridge: Cambridge Unviersity Press.

Blank, A. 1999. Why do new meanings occur? A cognitive typology of the motivations for lexical semantic change. In *Historical Semantics and Cognition*, ed. A. Blank and P. Koch. 61-90. New York, NY: Mouton de Gruyter.

Bybee, J. 2015. *Language Change*. Cambridge: Cambridge University Press.

Davidson, D. 1984. Convention and communication. *Synthese* 59: 3-17.

Davidson, D. 1986. A nice derangement of epitaphs. In *The Philosophical Grounds of Rationality*, ed. R. Grandy and R. Warner. 443-446. Oxford: Oxford University Press.

Davidson, D. 1991. James Joyce and Humpty Dumpty. *Midwest Studies In Philosophy* 16: 1-12.

Ebert, Cornelia, and Christian Ebert. 2014. Gestures, demonstratives, and the attributive/referential distinction. Handout of a talk given at *Semantics and Philosophy in Europe (SPE 7)*, Berlin, June 28, 2014. https://semanticsarchive.net/Archive/GJjYzkwN/EbertEbert-SPE-2014-slides.pdf

Elman, J. L. 2009. On the meaning of words and dinosaur bones: lexical knowledge without a lexicon. *Cognitive Science* 33: 547-582.

Fromkin, V. 2000. *Linguistics: An Introduction to Linguistic Theory*. Malden, MA: Blackwell.

Greenberg, G. 2013. Beyond resemblance. *The Philosophical Review* 122: 215-287.

Halle, M., and A. Marantz. 1993. Distributed Morphology and the pieces of inflection. In *The View from Building 20*, ed. Kenneth Hale and S. Jay Keyser. 111-176. Cambridge, MA: MIT Press.

Halle, M., and A. Marantz. 1994. Some key features of Distributed Morphology. *Papers on Phonology and Morphology, MITWPL 21*. 275-288. Cambridge, MA: MIT Working Papers in Linguistics.

Hunter, J. to appear. Relating gesture to speech: reflections on the role of conditional presuppositions. *Linguistics and Philosophy.* http://www.juliejhunter.com/uploads/3/9/6/1/39617901/relating_gesture_to_speech.pdf

Jackendoff, R. 2002. *Foundations of Language*. Oxford: Oxford University Press.

Kamide, Y., G. Altmann, and S. Haywood. 2003. The time-course of prediction in incremental sentence processing: evidence from anticipatory eye movements. *Journal of Memory and Language* 49: 133-159.

Kendon, A. 2004. *Gesture: Visible Action as Utterance*. Cambridge: Cambridge University Press.

King, J., and J. Stanley. 2005. Semantics, pragmatics, and the role of semantic content. In *Semantics versus Pragmatics*, ed. Z. Szabo. 111-164. Oxford: Clarendon Press.

Lascarides, A., and M. Stone. 2009. A formal semantic analysis of gesture. *Journal of Semantics* 26: 393-449.

Lepore, E., and M. Stone. 2017. Convention before communication. *Philosophical Perspectives* 31: 245-265.

Lewis, D. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8: 339-359.

Ludlow, P. 2014. *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. New York, NY: Oxford University Press.

Perniss, P., E. L. Thompson, and G. Vigliocco. 2010. Iconicity as a general property of language: Evidence from spoken and signed languages. *Frontiers in Psychology* 1: Article 227. https://doi.org/10.3389/fpsyg.2010.00227

Pinker, S. 1999. *Words and Rules: The Ingredients of Language*. London: Weidenfeld & Nicolson.

Predelli, S. 2010. Malapropisms and the simple picture of communication. *Mind & Language* 25: 329-345.

Rayo, A. 2013. A plea for semantic localism. *Nous* 47: 647-679.

Schlenker, P. 2018. Gesture projection and cosuppositions. *Linguistics and Philosophy* 41: 295-365.

Schlenker, P. to appear. Gestural semantics. *Natural Language & Linguistic Theory*. http://ling.auf.net/lingbuzz/003485

Xian, Y., C. H. Lampert, B. Schiele, and Z. Akata. 2018. Zero-shot learning: a comprehensive evaluation of the good, the bad and the ugly. https://arxiv.org/abs/1707.00600v3

Yalcin, S. 2014. Semantics and metasemantics in the context of generative grammar. In *Metasemantics*, ed. A. Burgess and B. Sherman. 17-54. New York, NY: Oxford University Press.