

One Argument, Many Interpretations: Polysemy and the Argument from Copredication

Christopher A. Vogel

Vogel, C.A. (2025) One Argument, Many Interpretations: Polysemy and the Argument from Copredication. *Inquiry*, pp. 1–28
(please cite original source: <https://doi.org/10.1080/0020174X.2025.2493058>)

Abstract

Cases of copredication like *The book was long but boring* are marshaled against externalist theories of meaning, alleging that the rampant polysemy of natural language expressions resists externalist explanations. Various externalist responses deny the claimed polysemy, appealing to the role of speaker intuitions about reference in arguments that mobilize copredication data against externalism. I argue that these denialist strategies fail because they miss the import of arguments from copredication, which point to natural language speakers' inferences, not intuitions about reference.

keywords: externalism, internalism, polysemy, truth-conditions

disclosure: The authors report there are no competing interests to declare regarding this research.

1 Introduction

Natural language expressions exhibit a pronounced semantic flexibility. Recently this flexibility has been proffered as evidence against externalist conceptions of linguistic meaning (Chomsky 2000; Pietroski 2005; Collins 2009; Glanzberg 2014; Yalcin 2014; Azzouni 2017; Murphy 2023), often citing copredication as revealing the *polysemy* of natural language expressions (Pietroski 2025; Quilty-Dunn 2021; Collins 2017). Sentences like

- (1) The bank went out of business after burning down.
- (2) Lunch was delicious but took forever.
- (3) The book is a compilation, and it is too heavy to lift.

instantiate the relevant properties, where a single use of an expression (e.g., *book*) solicits multiple, distinct ways of representing the sentence’s subject. Such expressions are arguably polysemous, expressing multiple senses/concepts with a single use. An utterance of (3), with the single use of *book*, invites one to think about books both as concrete objects that move scales, and as abstract contents with narrative types. Standard externalist accounts of such copredication seemingly require a single bizarre hybrid object to satisfy the conflicting predicates.

This critical attention to copredicational sentences, unsurprisingly, motivates externalist accounts of the relevant phenomena (cf. Ortega-Andrés 2023; Liebesman and Magidor 2017; Vicente 2019; Ludlow 2011). But one kind of response *denies* that these data reflect natural language polysemy, by denying either that the relevant expressions are ambiguous, or that polysemy is a unique form of ambiguity (King 2018; Devitt 2021). More recently, a strategy in this vein explicitly denies that polysemy exists (Brody and Feiman 2023). I argue that these *denialist* strategies rely on a fundamental misunderstanding of the challenge copredication poses for externalists. Below I focus on the novel denialist variant in Brody and Feiman (2023) to highlight the mistaken assumption common to denialists, namely that arguments from copredication rely on speakers’ intuitions about reference. I’ll show that the challenge posed by copredication pertains to the sentences speakers are apt to infer from sentences that utilize flexible expressions.

I begin in §2 by rehearsing the copredication data, the ontological concerns these data have raised, and the denialist response to them. In §2.1 I

detail the considerations Brody and Feiman (2023) raise regarding the relevance of speakers' intuitions about reference for what they call the *argument from polysemy* (or AFP). In §2.2 I show that denialist responses to the AFP assume that copredicational nominal reference, and related speaker intuitions, are the data externalists must explain. In §3 I motivate what I call (for clarity's sake) the *argument from copredication* (or AFC). This argument obligates a semantic theory to explain the *inferences* speakers endorse regarding sentences like (1)–(3). In §3.1 I show that the referential assumption underwriting the denialist response to the AFP misses the core data from semantic flexibility. I then indicate the inferential pattern cited by the AFC. In §3.2, I offer evidence that this pattern of inference generalizes to *prima facie* polysemes. While externalist commitments suggest a referential interpretation of the problem copredication poses, the argument from copredication can be made without invoking a notion of reference. I conclude by entertaining three objections, and relate them to future directions that responses to the AFC might take.

2 Copredication and Reference

Sentences like (1)–(3) are the basis for what Brody and Feiman (2023) label the *argument from polysemy* (or AFP). These are cases of *copredication* since multiple predicates are applied in tandem to a single grammatical subject. Data related to uses of these sentences arguably undermine externalist semantics (see Murphy 2023 for a helpful review).

In §2.1 I sketch the AFP, indicating how anti-externalists take examples like (1)–(3) to undermine externalist theses. In §2.2 I review a class of externalist responses to the AFP, which identify the problem stemming from copredication as a problem of reference (or the related speaker intuitions). These *denialist* responses deny that expressions are genuinely polysemous, arguing that sentences like (1)–(3) fail to exhibit a unique kind of ambiguity. In §3 I argue that these denialist responses fail to address the core challenge that sentences like (1)–(3) pose, which pertains to the inference of speakers, not their intuitions about reference.

2.1 AFP: The Argument from Polysemy

An externalist accepts the following hypothesis:

- (E) The meaning of a declarative utterance of a sentence (in a context) determines its truth-conditions (in that context).¹

The familiar idea is that sentence meanings express truth-conditions, indicating what the world is like if the sentence is true. Relative to a context, using a sentence expresses a determinate truth-condition, and natural language speakers understand the meanings of expressions insofar as they understand those truth-conditions.

AFP arguments cite copredication as a class of cases that externalism struggles to explain. But copredication *simpliciter* is not troubling for externalists.² Sentence (4) poses no obvious problems for externalism:

- (4) The dog is a terrier, but it is bipedal.

If meanings determine truth-conditions, (4) requires that the worldly domain contains things that could jointly satisfy both predicates. That is, there must exist a kind of thing such that some single thing of that kind could satisfy the two predicates by having three properties: it is a dog, of a particular breed, and has exactly two legs. Put more plainly, for externalists the felicity of (4) requires that a two-legged dog (can possibly) exist. That ontological commitment is unproblematic.

What distinguishes the copredicational sentences in (1)–(3) from (4) is that they illustrate the apparent *flexibility* of natural language. Put roughly,

¹See Pietroski (2005, 254, 268–9, 276, 287, 289) for a similar characterization of externalism, which itself comports with what externalists themselves identify as the central tenet of externalism. See for example Stanley (2007, 7) and King (2018, 2).

²The use of *copredication* in the polysemy literature tends to highlight cases where the pair of predicates are in some sort of tension, insofar as ‘copredication is the phenomenon whereby a single, felicitous occurrence of an expression appears to take several conflicting interpretations’ (Viebahn 2022, 1066). For reasons that will be clear below, it will make more sense to talk about the inferences that speakers are apt to make across various cases where two predicates are applied to a single grammatical subject *irrespective* of whether the predicates generate ‘conflicting interpretations’ of a single expression. In that sense, my use of *copredication* (and its cognates) is somewhat novel/idiosyncratic—though see Collins (2021) who suggests a similar usage with the phrase ‘polysemous copredication’ and Asher (2011) who defines copredication as ‘a grammatical construction in which two predicates jointly apply to the same argument’ (11). My slightly broader usage more succinctly captures the import of copredication arguments: offer a uniform account of *all* kinds of copredication inference, not merely those that exhibit the intuitive tension in examples like (1)–(3).

while *the dog* in (4) is about a single pet, the grammatical subjects in (1)–(3) seem to recruit multiple distinct ways of thinking about objects that no *single* object can satisfy. Call sentences like (1)–(3) instances of felicitous *flexible* copredication, which contrast with sentences like (4) in terms of the indicated flexibility.

Instances of felicitous flexible copredication appear to reveal the *polysemy* of expressions like *book*, having multiple ‘different senses’ whereby speakers invoke in ‘a process of “modulation” [that] picks out the contextually relevant one [i.e., sense], putting the right sense together with the right predicate’ to explain the felicity of flexible copredicational sentences (Brody and Feiman 2023, 4). With *book*, each sense³ is needed to provide the kind of object each predicate demands: one for concrete objects with mass, and another for abstract contents that minds entertain. Otherwise, monosemous *book* would (according to (E)) require that some *single* thing can be at once both abstract and concrete.

Brody and Feiman (2023) contend however that the relevant expressions (i.e., *lunch*, *bank*, and *book*) are not, simply in virtue of such felicitous flexible copredication, thereby polysemous. Rather, characterizing them as polysemes requires an assumption linking ‘intuitions of ontology’ (and reference) with claims about meaning. Specifically, attributing polysemy to these expressions assumes the following, which I label (R):

- (R) If people attest that a word is being used to refer to distinct entities, then these entities correspond to different senses of the word (Brody and Feiman 2023, 5).

So according to (R), if a single expression (e.g., *the book*), as used in a context, is seen by speakers as referring to distinct objects (e.g., a concrete artifact, and an abstract content), then that expression must have multiple

³Brody and Feiman (2023) use the term *sense* to indicate the content a declarative use of an expression contributes to a sentence’s interpretation. This terminology can be misleading in various ways, but attempts at clarification invariably beg many a question. The following is an approximate, though not completely unobjectionable, gloss on what is at stake regarding polysemy: whether the use of a single expression in a declarative utterance, relative to a single context, can provide either *one* or *more than one* satisfaction condition to the proposition that that use expresses, irrespective of which proposition a speaker intends to convey. Taking ‘senses’ as such, Brody and Feiman (2023) contend that polysemy does not exist insofar as no expression can provide more than one such satisfaction condition per context.

senses, each capable of contributing to the utterance’s interpretation. The AFP, with (R) in hand, purports to undermine externalist theories of meaning (and correlative views about conceptual contents). The argument seems straightforward. The data in (1)–(3) indicate that speakers interpret expressions like *the book* as *referring* to both concrete objects (with leaves) and abstract contents (with narratives). Given (R), words like *book* express at least two senses, contradicting externalist theories of meaning, which require monosemous meanings to determine a single referent.

2.2 Denialist Responses to the AFP

One class of replies to the AFP *denies* that the relevant flexible expressions are polysemous. The *ontological* version of the denial strategy insists that expressions like *the book* are semantically no different than expressions like *the dog*, accepting that a use of *the book* in sentences like (3) has some single referent. There is (at least possibly) a *single* thing that is both too heavy to lift and a compilation. Devitt (2021) offers this kind of ontological response, treating the threat of copredication to externalism as depending on referential considerations.⁴ Expressions like *the book* work as *the dog* does by denoting a single referent, and are thereby *not* polysemous. Consider the difference suggested between (polysemous) *lunch* and (monosemous) *book*:

“Lunch” ...can refer to *objects* eaten or *events* of eating. These are two very different categories of entities. They are not “aspects” or “conceptualizations” of the one entity. ... “Books” are quite unlike “lunches” in this respect. (Devitt 2021, 232)

The contention is that *book* is *not* polysemous. For Devitt, (apt) uses of *the book* refer to a *kind* of object with a singular essence:

⁴Devitt (2021) discusses at length the role of intuition data from speakers and its importance in linguistic theorizing. He often talks as if such judgments are irrelevant to theorizing (21–27), preferring to identify those conventions around the use of words as the primary source of data. But surely part of those conventions regards what speakers take other speakers to be talking about when using expressions like *the book* to communicate. Indeed, his way of addressing the challenge from copredication only makes sense if framed within the need to take such judgments seriously, otherwise statements like ‘ “This book” (semantically) refers to that one token physical object and the two predications are thus of *one and the same object*, the one picked off my desk’ would be a *non sequitur* (233).

For an object to be a book is for it to have *both* that form *and* that content; that’s the nature/essence of the kind *book*. . . . In virtue of its somewhat complex nature, a book, *the one object*, can be both heavy and real fun. (Devitt 2021, 232)

On this strategy addressing the problem posed by flexible copredication requires metaphysical ingenuity to identify referents for the offending terms—in this case speculation about books given uses of *the book*.⁵ With such a metaphysics of *book* in hand, the problem posed by copredication is resolved because apt uses of *the book* denote a single object instantiating both ‘predicated *properties of the book, that very object picked off my desk*. So these predications do not demonstrate any polysemy at all’ (Devitt 2021, 233). Thus *the book* is no more polysemous than *the dog*, denoting a single (hybrid) object. This *denies* that *book* is polysemous.

As indicated above, Devitt distinguishes *book* from *lunch*, whereby *lunch* admits to ‘regular polysemy’, able to refer to both *events* of eating, and *objects* being eaten: ‘it can refer to an event or to food’ (Devitt 2021, 260).

(2) Lunch was delicious but took forever.

The ‘regular polysemy’ exhibited in (2) is purportedly explained by the second version of the denial strategy, which treats polysemy as *homophony/homonymy*:

We can say simply that the displayed sentences containing them show that the one token of such a word can indeed express both of its conventional senses in the one sentence. So, they illustrate

⁵Devitt (2021) contends that such an account of book-ontology is not motivated by semantics, but by broader appeals to parsimony: positing such ‘complex’ books is more explanatory (233). However, one might want to know what precisely is being explained by such posits. If the *explananda* pertain to our treatment of books, imbuing pulverized dead tree matter with both concrete and abstract properties is not required to explain why humans that have acquired an I-language might treat objects with orthographs differently than a wad of paper and ink (232). Further, an appeal to this prescriptive book-ontology ignores the fact that English speakers abide by the clear convention of using *book* in discussing, for example, *notebooks* one has not yet taken notes in, or *books* as aspects of a digital environment that a video-game character might collect, or plastic pieces in a doll house that can be arranged on a small *bookshelf*, none of which *must* contain orthographs to be aptly described using *book*.

a phenomenon akin to pronouns of laziness and the “sloppy reading” of verb phrase anaphora . . . We might say that (2) is elliptical for:

[(2')] The lunch-qua-food was delicious but the lunch-qua-event took forever.

But that is a clumsy way to get the message across. So it is not surprising that we have come up with a syntactic convention that enables us to get the message across by (2), exploiting both meanings of ‘lunch’ at once. (Devitt 2021, 260)

Sentence (2) is just a ‘lazy’ way of saying sentence (2’), which uses two distinct expressions for lunches. Thus what speakers understand ‘with a syntactic convention’ when they understand (2) is the sentence in (2’). As Devitt indicates, ‘on this view, the two senses of the polysemous “lunch” yield two entries in the lexicon, just as does the homonymous “bank” . . .’ (Devitt 2021, 260, *my emphasis*).

This second kind of denial approach treats expressions like *lunch* as instances of a different ambiguity, homonymy/homophony. Two meanings are associated with *lunch*: *lunch-qua-food* and *lunch-qua-event*, where the second permits reference to eating events and the first permits reference to objects eaten. Much like the homonymous/homophonous *bank* whereby a single pronunciation /bæŋk/ gets assigned two meanings (one for financial institutions and another for waterway boundaries), /lʌntʃ/ is paired with two meanings—there are two lexical entries for *lunch* (Devitt 2021, 260). This response simply *denies* that polysemy is a distinct form of ambiguity, treating ‘regular polysemy’ cases as instances of homophony/homonymy.⁶

The ontological and homophonous replies share in common the contention that polysemy does not exist, by either modeling it as homophony, or by

⁶King (2018) arguably pursues the homophonous version of the denial strategy in considering cases like (3). He claims that sentences like (3) do not ‘force us to reject an externalist semantics any more than the phenomenon of ambiguity does’ (King 2018, 9). The difference between ‘ambiguous’ homonyms like *bank* and polysemes like *book*, he contends, is that the meanings associated with homonym pronunciations diverge considerably. So sentences like *The bank is flooded and it is an investor* cannot be interpreted as being about a overflowing river and a financial institution. While ‘in cases where [meanings] are “sufficiently close” meanings’, like *book* in (3), an analogous interpretation is available (King 2018, 9). But as Collins (2023) points out, this assumes that the lack of a unified meaning for *book* renders it as ‘ambiguous’ in the way *bank* is.

contending that the offending expression is not actually ambiguous, readily providing the single needed referent. These denial approaches take the data utilized by the AFP in sentences like (1)–(3) as pertaining to reference. Brody and Feiman (2023) makes this assumed appeal to referential intuitions and the denial of polysemy explicit by contending that the AFP argument first, *assumes* (R), and, second, that (R) is *false* (5).

The problem with (R), as Brody and Feiman (2023) underscore, is that it over-generates, rendering (uses of) expressions as polysemous that seem otherwise. Consider one class of counterexamples that leverages ontological composition, with copredicational sentences (purportedly) expressing multiple properties that apply to an object’s *parts*.

- (5) The cat got shaved and vaccinated.

Applying (R) to such cases is problematic because it ‘will diagnose every word with multiple senses’ (Brody and Feiman 2023, 5). For the sentence in (5) principle (R) mandates that *the cat* is polysemous along super- and sub-cutaneous lines. The idea is that *the cat* references ‘distinct things’ given the copredication: ‘shaving the cat only affects the *outside* of the cat, while vaccinating it only affects the *inside* of the cat. Intuitively these are quite different entities...’ (5). They argue that (R) mistakenly renders *the cat* in (5) polysemous because of these two distinct referents. But to insist that this use of *the cat* is polysemous is implausible: ‘are we therefore to believe that *cat* is polysemous between $CAT_{insides}$ and $CAT_{outsides}$?’ (5).

The claim in (R) is clearly about reference, pinning the import of the AFP to speakers’ referential intuitions. For Brody and Feiman (2023), since (R) is central to the AFP and is false, the AFP fails to undermine externalist theses. Given this failure, externalists are entitled to simply *deny* that polysemy exists. If the relevance of copredication for semantic theorizing was simply a function of explaining referential intuitions, the denialist’s approach seems initially attractive. After all, preserving the idea that names are directly referential, *Ulysses* (the book) works just like *Max* (the dog), referring to a single thing. Books are plausibly hybrid things, at once both abstract and concrete. And after considering the naturalness of treating humans as hybrid objects (see §4), accepting a hybrid ontology might seem not so bizarre after all. Further, if the referential principle in (R) is required to press the critical point of copredication data against externalism, then noting that (R) miscategorizes monosemes as polysemes offers further motivation for the denialist approach.

In the next section I argue however that the direct argument appealing to lexical flexibility—the *argument from copredication* (AFC)—makes *no* such assumption about referential intuitions, including (R). So, the arguments Brody and Feiman (2023) offer against (R) do not address the problem that sentences like (1)–(3) present for externalism. The ontological and homophonous denial strategies share this mistake, ignoring the import of the *inferential* data pregnant in (1)–(3).

3 Copredication and Inference

In the previous section I reviewed the denialist response to the AFP, which leverages the felicity of copredicational sentences as evidence against externalist theses. As we saw Brody and Feiman (2023) claim that AFP arguments fail because the purported problem that copredicational sentences pose for externalism stems from ‘intuitions of ontology [that] certain kinds of things cannot have certain kinds of properties’ (4). Such an argument from *polysemy* simply assumes that the felicity/acceptability of a copredicational sentence indicates that the relevant (use of the) noun is polysemous, having multiple senses expressible by a single use of the expression. This, they claim, is mistakenly thought to be a consequence of the fact that such nouns (like *France*) ‘[appear] to describe different referents, or even different types of referents’ that speakers ‘can still some how pick . . . out using a single word’ (4). So, the AFP leverages referential judgments, and speaker assumptions about the intuited referents, working backwards from them and insisting that nouns like *France* must be polysemous. More simply, the AFP wrongly assumes (R).

In this section I outline the *argument from copredication* (AFC) which avoids this referential commitment, citing a collection of speaker inferences that externalism struggles to explain. I begin in §3.1 by detailing the AFC, which is an invitation to explain the inferences that native speakers are apt to make in cases of felicitous copredication. For felicitous *flexible* copredicational sentences, speakers are disposed to make different inferences compared to other cases of copredication. The AFC suggests that externalist commitments forestall attempts to explain this difference. In §3.2 I show that this inferential pattern generalizes. There I introduce the technical notion of a *flexeme*, which applies to any expression that admits to the felicitous flexible copredication pattern of inference. The upshot is that a broad class of

natural language expressions are flexemes, admitting to a pattern of inference that is best explained by positing polysemy, and for which the most straight-forward externalist explanations fail.

3.1 AFC: The Argument from Copredication

The seminal argument from copredication in Pietroski (2005) takes inspiration from an example in Austin (1962), and considerations raised by Chomsky (2000). Consider:

- (6) France is hexagonal and it is a republic.
- (7) France is hexagonal.
- (8) France is a republic.

The flexibility of *France* exhibited in these sentences poses a problem for externalist theories of meaning because of the inferences speakers draw between the sentences. Roughly, speakers are apt to infer (7)–(8) from (6) while felicitously denying that some single thing satisfies both predicates. Externalist commitments frustrate plausible explanations for this pattern of inference. The argument from copredication regards the *inferences* speakers are apt to make in cases of flexible copredication, remaining silent on matters of reference.

Since the AFC directly appeals to the inferences of speakers (as I detail below), it's worth highlighting that reverence for inferential data is not foreign to externalism. Because individual acts of inferring one sentence from another depend on the interpretations speakers apply to sentences, patterns of speaker inferences offer insight into the nature of meaning, and constrain plausible theories of meaning. Historically, one central source of argument favoring (E) appeals to the inferences supported by adverbial modification (Davidson 1967). Modeling speakers' inferences as abiding entailment⁷ relations between fundamentally eventitive logical forms offers an arguably straight-forward explanation for adverbial inferential patterns. How-

⁷Importantly, the *data* in question here cannot be described as exhibiting entailment relations, *pace* King (2018, 13). Insisting that sentences have entailment properties (as opposed to claiming that speakers infer some sentences from others), and further demanding that a theory of meaning must explain them, presupposes what many internalists deny: that sentences are truth-evaluable, even relative to a context. After all, entailment is a relation between two propositions and truth. So describing the natural language inferences of speakers as illustrating entailment relations between sentences already assumes

ever, precisely because such an explanation is taken to count as evidence for externalist theses, one should not conflate the speaker data regarding acts of inferring with the theoretical explanation of those data that appeals to entailment relations. Speakers make the inferences they make. The Davidsonian insight is that those actions can be rendered intelligible by a theory of meaning that takes these acts of inferring to be nothing more than acts of deriving entailments.

The AFC likewise takes speaker inferences as central data for good semantic theorizing. In particular, speaker inferences from flexible copredicational sentences like (6) are paradigm data that require explanation. Since English speakers that accept (6) will thereby also accept (7) and (8), a theory of meaning must explain such facts. For theories that endorse (E), standard explanations for speaker inferences in terms of entailment must take onboard additional implausible assumptions regarding flexible copredication inferential patterns.

To see the point, consider how an externalist explains non-flexible copredicational inferences:

- (9) Max is quadrupedal and it[*/she*] is a terrier.
- (10) Max is quadrupedal.
- (11) Max is a terrier.

English speakers that accept (9) will also accept (10) and (11). An externalist readily explains this inferential pattern: the truth-conditions determined by the meaning of (9) entail the truth-conditions for (10) and (11). For the externalist, a speaker's semantic competence is (at least partly) constituted by their ability to derive the truth-conditions/propositions determined by the meaning of a sentence. So acts of inference are explained as instances of speakers deriving (and recognizing) the entailment relations between the relevant truth-conditions/propositions. Here the entailment is a matter of conjunction reduction, as made clear by explicitly indicating their (purported) truth-conditions (in a single context):

- (9') $Q(m) \ \& \ T(m)$
- (10') $Q(m)$

that sentences (as opposed to the propositions speakers express by the use of sentences) are routinely either true or false.

(11') $T(m)$

If we take the anaphoric *it* in (9) as copying the semantic value of *Max*—for externalism its determined referent—then the truth-condition expressed by an assertive use of (9) is likewise expressed by the formal sentence (9'). Some single referent, indicated by m , has a pair of properties. Sentence (9) is true just in case some single thing is both four-legged and a terrier. Plainly, if (9) expresses that truth-condition, it entails the truth-conditions indicated in (10') and (11'). Thus the inferences of English speakers *according to the externalist* are explained because apt uses of *Max* refer to a single individual which can satisfy both of the predicates in (9). If that individual satisfies those predicates by having the properties of *being-four-legged* and *being-a-terrier*, it likewise satisfies the predicates in (10') and (11').

The AFC indicates that an analogous explanation for the inferential patterns in instances of *flexible* copredication requires taking on implausible assumptions. Direct analogs to the example just rehearsed utilizing *France* indicate truth-conditions of the following sort:

(6) France is hexagonal and it is a republic.

(7) France is hexagonal.

(8) France is a republic.

(6') $H(f) \ \& \ R(f)$

(7') $H(f)$

(8') $R(f)$

An analogous *externalist* explanation of the inferences from (6) to (7) and (8) posits a single thing that apt uses of *France* can refer to, the thing indicated by f in (6')–(8'). As with *Max*, whereby speaker inferences are captured by the possibility of a single (mind-external) domain object that is both a terrier and quadrupedal, for the externalist *France* must follow suit: there are single things capable of being both hexagonal and a republic. Such hybrid objects are able to satisfy the predicates in (6') by having the properties of *being-hexagonal* and *being-a-republic*. This is the ontological version of the denialist strategy, applied to countries instead of books. There are things at once both abstract and concrete such that both predicates in (6)/(6')

could apply to them.⁸ So, *externalism*'s assumptions regarding the nature of reference carry an apparent ontological commitment to abstract-concretea. This commitment is required *by the externalist* to explain the inferences from (6) to (7) and (8).

Pietroski (2005) and others take such a commitment to be implausible, but not (simply) because there are no hybrid abstract-concretea:

For even if one grants that the terrain of France is hexagonal, one might deny that the French terrain is the French republic; and one might think that republics, whatever they are, cannot be hexagonal. Moreover, even if there are such [hybrid] things, one might be suspicious of a *semantic* theory according to which it *follows* from [(6)] that there are hexagonal republics. Such a theory apparently *mischaracterizes* the meaning of [(6)]. For a competent speaker who asserts [(6)] might deny that at least one thing is both hexagonal and a republic. (Pietroski 2005, 267)

The point here is that if, as (E) suggests, understanding the meaning of a sentence is to understand what the world must be like according to a use of it, then accepting (6) as apt and felicitous requires a single object as a predicate satisfier. After all, the felicity of an expression on externalist theories stems from its (speaker-perceived) truth. But given that (6') directly entails

(12') $\exists x.H(x)\&R(x)$

⁸An externalist might balk at this conclusion, insisting that just because a thing has the properties of (say) *being-France*, *being-hexagonal*, and *being-a-republic*, it does not follow that such a thing is at once both abstract and concrete. Liebesman and Magidor (2017) make a similar claim regarding books:

the key challenge raised by copredication arises from the supposition that the properties ascribed in copredicational sentences [for example, about books] cannot be instantiated by both [a physical book and an informational book]. . . . Our view is that this is bad metaphysics. Informational books are distinct from physical books, but there are many properties that both can instantiate (137).

Whether such suppositions are misplaced on metaphysical grounds or not, permitting the existence of informative physical books (or terrain with political power structures) does not respond to the AFC. Whatever books are like, one wants to know why a speaker that accepts *The book is verbose and it is a collectible* need *not* thereby infer that *Something is a verbose collectible*. See §3.2 below. If the meaning of *the book* determines a single informative physical referent, this inference should be pristine.

the externalist fails to explain why an English speaker that endorses (6) (or in externalist mode, takes it to be true) need *not* infer (12) from (6):

(12) Something is a hexagonal republic.

And that inferential fact needs to be explained *even if* (12) is true because there are abstract-concretea to satisfy the formal sentence in (12') as the denialist contends.

The collection of inferences that speakers are apt to make from copredicational expressions requires an explanation. The AFC offers a challenge: explain these inference. The problem for theses like (E) is that straightforward applications of them cannot explain the relevant inferences without accepting additional implausible assumptions.⁹ But of note, nothing in the AFC just presented asserts or assumes (R), *pace* Brody and Feiman (2023). The argument draws out an apparent implication of (E), which is bolstered by common externalist commitments about names as directly referring expressions.¹⁰ The AFC need not even assume that speakers interpret uses of expressions as referring to anything, ever. The argument is stated in terms of speaker inferences from copredicational sentences.

3.2 Generalizing the Inferential Pattern

The AFC demands an explanation for the inferential patterns endorsed in cases of copredication.¹¹

⁹There are many ways in which externalists have responded to the problem copredication poses for a theory of meaning that accepts (E) (cf. Ortega-Andrés 2023; Luo 2012; Gotham 2017; Gotham 2022; Vicente 2018; Vicente 2019; Löhr 2024). See Vicente (2021) and Falkum and Vicente (2015) for a helpful discussion. The purpose of this paper is not to show that all of these responses fail, but rather to show that the problem copredication poses is not a problem of reference. Rather, the point is that *externalism* makes copredication a problem about reference.

¹⁰See for example King (2018) who, in defending externalism, states ‘externalist semantic theories posit a semantic relation (perhaps relative to a context or other parameters) between expressions of the language and entities in the world. Such relations have gone by many names: “__ having __ as its semantic value (relative to context *c*)”; “__ refers to __ (relative to *c*)”; “||__||_{*c*} = __”; etc. I think it is safe to say that much recent semantic theorizing is externalist in this sense’ (1).

¹¹A point of clarification about notation. In the inferential patterns exemplified in (13)–(15), the inferential data of interest is indicated by the (in)felicity of prepending *Therefore* before each sentence, where the important data point pertains to the existential sentence

- (13) The dog is quadrupedal and it is a terrier.
- a. [Therefore] The dog is quadrupedal.
 - b. [Therefore] The dog is a terrier.
 - c. [Therefore] Something is a quadrupedal terrier.
- (14) # The dog is quadrupedal and it is a biped.
- a. [?? Therefore] The dog is quadrupedal.
 - b. [?? Therefore] The dog is a biped.
 - c. [# Therefore] # Something is a quadrupedal biped.
- (15) The country is hexagonal and it is a republic.
- a. [Therefore] The country is hexagonal.
 - b. [Therefore] The country is a republic.
 - c. [# Therefore] Something is a hexagonal republic.

Since the inferences from (13) to each of (13a–c) are pristine, one wants to know why structurally identical inferences in (14) and (15) fail. While externalist theories of meaning can offer plausible explanations for the differing patterns in (13) and (14) by appealing to a notion of (in)commensurability, such an appeal is unhelpful for explaining the inferences in (15). That is, the infelicity of inferring (14a–c) from (14), indicated by marking a pretended *Therefore* as either infelicitous (#) or confusing (??), is plausibly explained by the incommensurability of the two predicates in (14). The predicates are incommensurable insofar as no single object can instantiate the pair of properties the two predicates express (i.e., no single thing can have exactly four legs and exactly two legs). But a similar proposal fails for the pattern in (15), whereby the inference from (15) to (15a–b) is mandated, but not the inference to (15c). If (14) and the related inferences are infelicitous because the two predicates are incommensurable as such, then one wants to know why *only* the inference from (15) to (15c) is infelicitous for the (presumably) incommensurable predicates in (15). Alternatively, if the two predicates in

Something is a . . . Put concretely, the sentence *Something is a hexagonal republic* might be felicitous (or not). But inferring that sentence from *France is hexagonal and it is a republic* is not mandated—a fact reflected in the use of the pound sign in ‘[# *Therefore*]’. This hopefully avoids confusion about the purpose of the AFC, which is not directly about the felicity of copredicational sentences, but about which inferences speakers are apt to make or deny. But, see the discussion on page 18 about enhancing the felicity of copredicational sentences.

(15) *are* commensurable (say, because there is a hybrid object that can satisfy both predicates), one wants to know why the inference to (15c) isn't mandated in the way that that analogous inference to (13c) is. Such an incomplete externalist solution to the copredication problem might appeal to a notion of incommensurability, apparently leveraging referential considerations. But the problem posed by the AFC of explaining inferential patterns like those in (13)–(15) does not rely on any assumptions about the nature of reference in the manner denialists assume.

The goal of this section is to show that the pattern in (15) generalizes to paradigm polysemes, solidifying the robustness of the phenomenon and grounding the criticism of externalist theses. However, because many externalist approaches deny that polysemes exist, stating that generalization in terms of polysemy unfairly begs the question against those externalist proposals. For the sake of precision, let's call copredication cases using *flexible* nouns (e.g., *France*, *book*, *country*, *lunch*, etc.) *flexemes* just in case they fit with the inferential pattern in (15). As an approximate definition, a *flexeme* is any expression whereby 1) speakers will infer from the felicitous copredicational use of that expression two single-predicate sentences, each predicating one of the copredicates to that expression, and 2) speakers can felicitously deny the inference to a sentence that applies those copredicates as a stacked existential. So the felicity of a copredicational use of a flexeme (like *country*) does not mandate the inference to a stacked predicate sentences like (15c) that make use of the initial copredicates. This contrasts with a use of non-flexemes like *dog*, which does mandate the stacked predicate inference, as in (13c).

Introducing this terminology will enable us to plainly state proposals for explaining why inferential (flexeme) patterns emerge, without begging questions against the denialist.¹² Since denialists reject the idea that words

¹²It's natural to think the notion of a flexeme is related to that of a polyseme. After all, the cases of flexemes thus far mentioned and invoked later in this section are all instances of fairly canonical polysemes. One might assume that flexemes are a special class of polysemes. But recall, the denialist claims that there are no polysemes—talk of polysemes, for the denialist, is like talk of phlogiston. So, to describe the inferential data to be explained as pertaining to the semantic or inferential behavior of polysemes, or some sub-class of them, appears to beg the question against the denialist. Importantly, while the denialist can argue on theoretical grounds that polysemy does not exist, they cannot likewise argue (at least not on theoretical grounds) that speakers do not make the inferences they in fact make. Invoking the notion of a flexeme is simply a means to identify that critical inferential data in a non-question-begging way.

express multiple concepts/sense, and so would deny that words like *country* are polysemes, we do better by avoiding stating the problem posed by copredication in terms of polysemy. So if the AFC is a demand to explain a collection of inferences, describing these patterns as the ones exhibited by ‘polysemes’ like *country*, *book*, etc. begs the question against the denialist. We can now state what the denialist wants to say about the relevant inferences more precisely, insofar as they contend that flexemes are not polysemous. The challenge offered by the AFC then is to explain the flexeme pattern in light of that denial.

With the notion of a flexeme in hand, in the remainder of this section I show that the most discussed (presumed) polyseme types in the literature abide by the flexeme pattern of inference. That is, canonical polysemes, but not monosemes, permit speakers to deny the inference to a stacked existential sentence from a felicitous copredicational one. Since Pustejovsky (1995), discussions of polysemous expressions often appeal to the types of senses/-concepts a (proposed) polyseme recruits. As discussed in §2.2, the expression *book* is a paradigm case of a proposed polyseme, recruiting two different ways of thinking about books: one as a PHYSICAL object, and another as a body of INFORMATION (or literary content).¹³ In contrast, *lunch* exhibits a duality with respect to the types FOOD (*Lunch was tapas*) and EVENT (*Lunch ended early*). In terms of these types, the most commonly discussed (presumed) polysemes regard some combination of the types PHYSICAL, INFORMATION, EVENT, and INSTITUTION (Xue and Luo 2012; Murphy 2024). Expressions that recruit these sense/concept types follow the flexeme pattern of inference.

The expression *book* is a flexeme, exhibiting the type duality of PHYSICAL and INFORMATION. In a discussion about the relationship between the aptitude of authors and the rarity of their books, the sentence in (16) is perfectly felicitous. Yet the pertinent uses of *book* do not mandate an inference to the stacked predicate sentence.

- (16) The book is verbose and it is a collectible.
- a. [Therefore] The book is verbose.
 - b. [Therefore] The book is a collectible.
 - c. [# Therefore] Something is a verbose collectible.

¹³On Pustejovsky (1995)’s Generative Lexicon approach, the meanings of polysemes encode multiple ‘qualia roles’, creating complex ‘dot-types’ as a combination of more basic types. See also Asher (2011). For the present task, the invocation of these types is helpful merely for taxonomic purposes.

A speaker can coherently deny that any collectibles are verbose, even upon accepting (16). This contrasts with monosemous cases, like the use of *dog* in (13), since someone that accepts (13) cannot likewise deny that at least one four-legged terrier exists.

Consider a second case involving *school*. Like *country*, *book*, and *France*, uses of *school* exhibit a concrete-abstract duality, expressing both a PHYSICAL type and an INSTITUTION type. In a conversation about the relationship between an institution’s popularity and the features of its campus, the use of *school* in (17) is felicitous and fits with the flexeme inference pattern:

- (17) The school is overenrolled and it is an eyesore.
 - a. [Therefore] The school is overenrolled.
 - b. [Therefore] The school is an eyesore.
 - c. [# Therefore] Something is an overenrolled eyesore.

Again, we see a paradigm polyseme fitting with the attested inferential pattern, whereby the inference to the stacked predicate is not licensed by felicitous copredication.

Finally, cases of copredication involving *lunch* similarly present the flexeme inferential pattern, with *lunch* exhibiting both FOOD and EVENT types. In a discussion about the most apt kind of meal to serve for a particular kind of gatherings, (18) is felicitous.

- (18) The lunch is sandwiches and it is a meeting.
 - a. [Therefore] The lunch is sandwiches.
 - b. [Therefore] The lunch is a meeting.
 - c. [# Therefore] Something is a sandwiches meeting.

Again, we see that the inference to the stacked predicate sentence is not licensed by the felicitous use of *lunch* in (18).

Admittedly, (18) does seem somewhat zeugmatic, making the sentence sub-optimal as a data point about ambiguity (Viebahn 2018; Liu 2023). However, we can adjust the example slightly to eliminate the zeugmatic judgment without altering the critical inferential pattern. While the contours of copredication judgments are not well understood, one feature that seems to matter for their semantic assent is coherence (Murphy 2024). When speakers are more readily able to decipher how the predicates in a copredicational sentence are related this facilitates acceptability. To the degree that (18) is semantically marked, (19) seems (at least comparatively) unobjectionable:

- (18) ?? The lunch is sandwiches and it is a meeting.
 (19) The lunch was sandwiches because it was a meeting.

The use of *because*, and the shift in tense, seems to implicate a contrast between the two predicates for the purpose of explaining the food choice *based on* the event's purpose. Nonetheless, the conjunctive aspect of *because* licenses the same pattern of inference relevant for (13)–(15) (Carston 1993). That is, replacing the conjunct *and* with *because* preserves the generalization that monosemes license the inference to the stacked predicate sentence while *prima facie* polysems do not. Consider a variant of (13) above, in contrast with (19):

- (19) The lunch was sandwiches because it was a meeting.
 a. [Therefore] The lunch was sandwiches.
 b. [Therefore] The lunch was a meeting.
 c. [# Therefore] Something was a sandwiches meeting.
- (20) The dog was quadrupedal [and/]because it was a terrier.
 a. [Therefore] The dog was quadrupedal.
 b. [Therefore] The dog was a terrier.
 c. [Therefore] Something was a quadrupedal terrier.

The conjunctive use of *because*, when replacing *and*, preserves the inferential patterns relevant to (non)flexemes. Despite the substitution, the paradigm polyseme *lunch* exhibits that same flexeme pattern pertaining to the stacked existential, while the monosemous (use of) *dog* retains the mandatory inference to the stacked existential. Why does this matter? The AFC demands an explanation for the inferences that speakers are apt to make from copredicational sentences. This demand is orthogonal to the more direct task of accounting for the acceptability of copredicational sentences on which much of the polysemy literature focuses. So, to illustrate the generality of the inferential pattern in need of theoretical capture, sentences like (18) may prove misleading to the degree that they are zeugmatic, plausibly because of a coherence constraint on admissible (flexible) copredication.

Prising apart the felicity conditions of copredicate expressions from the inferential data introduces two further considerations regarding the trouble copredication cases pose for externalist semantics. First, substituting *and* for *because* improves the acceptability of (18), and the typifying flexeme

inferential pattern emerges. But this shift in acceptability is particularly curious for externalist proposals, since the (would be) truth-conditions for (18) are entailed by the (would be) truth-conditions for (19). For externalists, the acceptability of an expression is tied to the truth-conditions its meaning determines. As such, if there is something semantically odd about (18), it's unclear why the *more* demanding claim in (19) rectifies that oddity. And note, that while the shift from *and* to *because* in (18) compared to (19) improves speaker acceptability, an analogous amendment to (14) does not help.

(14) # The dog is quadrupedal and it is a biped.

(21) # The dog is quadrupedal because it is a biped.

If anything, (21) is even worse than (14), or at least more confusing. This underscores the contrast between cases like (14) for which appealing to an externalist notion of incommensurability is plausible, and flexible copredication cases like those involving *country*, *book*, and *lunch* for which such an explanation seems implausible. So, whatever the explanation for these data might be, an appeal to an externalist notion of incommensurability is unhelpful.

Second, by appealing to the flexeme pattern of inference we can aptly capture the significance of purported counterexamples to copredication arguments, like (5) above. If one takes the inferential pattern in (15) as the paradigm for flexemes, (5) fails to generate that pattern, aligning instead with the monosemous (use of) *dog* in (13):

(5) The cat got shaved and [it got] vaccinated.

a. [Therefore] The cat got shaved.

b. [Therefore] The cat got vaccinated.

c. [Therefore] Something got shaved and vaccinated.

Unlike paradigm (uses of) flexemes like *France* or *country* that do not license the stacked existential inference, this use of *cat* does. So while an appeal to (R) miscategorizes monosemous (uses of) expressions as polysemous (as we saw in the previous section), an appeal to inferential patterns as the relevant diagnostic does not. Sentences like (5) may well undermine AFP arguments that appeal to referential principles like (R), but they do not serve as an undermining counterexample to the AFC. Rather, the failure of *cat* in (5) to pattern with flexemes makes clear why (at least on this usage) *cat* should not be a candidate for being polysemous.

Of course, none of these data prove that *country*, *school*, *book*, or *lunch* are polysemous. Copredication does not mandate that a particular use of an expression is polysemous (Liebesman and Magidor 2024). But, the AFC does not assert that certain expressions are polysemous, and then proceed to admonish externalist proposals for failing to explain that polysemy. Rather, the AFC offers a direct challenge to *any* theory of meaning: explain why speakers are apt to make the cited inferences in structurally (i.e., grammatically) similar cases. The collection of expressions in (13)–(15) serve as a kind of minimal triplet, indicating that an externalist appeal to entailment relations and incommensurability is not up to the task of explaining the behavior of flexemes. That demand is distinct from the further general proposal that the differential pattern of inference in cases like (13)–(15) is explained by an appeal to polysemy. That is, polysemous (uses of) expressions will betray the flexeme pattern of inference displayed in (15), while monosemous (uses of) expressions will not. This general proposal seems quite plausible given the data rehearsed in this section. But we should separate the challenge posed by the AFC and the general (under specified) positive polysemy proposal which purports to explain these inferential patterns.

Before turning to objections in the closing section, a word about the examples deserves attention. The cases involving paradigm polysemes above make use of a particular anaphoric structure, where an anaphoric clause and its antecedent containing clause are joined by a conjunct. The purpose of using these structures is to present, in the simplest form possible, a syntactically uniform set of examples that seem to differ only with regard to the target noun position (and the appropriate verbs). One reason for the simplicity is to avoid the introduction of complex grammatical structures whose proper analysis is still contentious. Another reason for this structure, and the indefinite in the anaphora clause in particular, is to mandate an individual reading of the target noun. This makes particularly vivid the trouble for an externalist solution to the AFC challenge, since on the most salient externalist proposals the indefinite requires a single referent to satisfy the joint (and stacked) predicates.

But, similar inferential patterns emerge in more complex syntactic frames. To offer a sampling, consider the following:

- (22) The translation took over a year. It is on the table.
 - a. [Therefore] The translation took over a year.
 - b. [Therefore] The translation is on the table.

- c. [# Therefore] ?? Something took over a year and is on the table.
- (23) The school that gave me a position closed after it burned down.
- a. [Therefore] The school gave me a position.
 - b. [Therefore] The school burned down.
 - c. [# Therefore] ?? Something gave me a position and burned down.
- (24) The dog ran over a mile. It is on the couch.
- a. [Therefore] The dog ran over a mile.
 - b. [Therefore] The dog is on the couch.
 - c. [Therefore] Something ran over a mile and is on the couch.
- (25) The dog that gave me a wound barked after it ran away.
- a. [Therefore] The dog gave me a wound.
 - b. [Therefore] The dog ran away.
 - c. [Therefore] Something gave me a wound and ran away.

Again, we see the (*prima facie*) polysemes *translation* and *school* pattern in the expected way, despite the different syntactic frames—in one case involving multiple sentences, and another involving a relative clause. In contrast, the monosemous use of *dog* in identical syntactic frames does not abide that inferential pattern. The challenge offered by the AFC could extend in a variety of ways not limited to the particular structures in (13)–(15).

Taking stock, we’ve seen that the denialist explanation of copredication is motivated by a particular gloss on the threat posed by flexible copredication cases. Because the perceived problem of accounting for the felicity of copredicational sentences regards the nature of the purported referents of flexible nouns, denialist solutions target referential intuitions in various ways. But as we saw in §3.1, the AFC pertains to the inferences that speakers are apt to make from a copredicational sentence. The demand to explain those data, and the related criticism of externalist proposals, makes no assumption about referential intuitions. Nor does that demand presuppose any view about the ontology of the claimed referents. Further, in this section (§3.2) we’ve seen that this inferential pattern generalizes, distinguishing the behavior of paradigm polysemous (uses of) nouns from monosemous ones. These data support the broad proposal that such flexibility is best explained as polysemy, wherein the flexible behavior of these (uses of) nouns is possible

because these nouns can recruit multiple distinct senses/concepts with a single use. The challenge then for the externalist (and denialist) is *not* to explain why some nouns beget referential intuitions that seem to imply the existence of bizarre hybrid objects. Rather, the challenge is to explain the flexeme inferential generalization, irrespective of one's ontological allegiances.

4 Objections and Future Developments

I conclude by considering three objections, drawing out the implications for further research on copredication and polysemy.

One might object that the flexeme inferential generalization is not universal. After all, some paradigmatic flexible expressions seem to pattern with monosemous (uses of) nouns. The following inference seems pristine:

- (26) Esteban is two meters tall and he is a platonist.
 - a. [Therefore] Esteban is two meters tall.
 - b. [Therefore] Esteban is a platonist.
 - c. [Therefore] Something[/Someone] is a two meter tall platonist.

The proper name *Esteban*, like *France*, is (*prima facie*) polysemous, capable of being used to talk about both a person (with goals, friends, ontological commitments, etc.) and a physical organism extended in space. As evidenced by (26), a copredicational sentence can successfully recruit both of these senses/concepts in a single use. Yet unlike *France* the stacked predicate sentence *does* seem to be inferentially mandated. So the claim that all polysemes will pattern with paradigm flexemes is mistaken.

While the data point seems apt, it fails to serve as a criticism of the AFC. The pattern of inference exhibited in (26) does require an explanation, and that complicates the general strategy of addressing the AFC by appealing to polysemy. But precisely because it seems to be exceptional, undermining the apparent universality of the generalization which indicates that flexemes are polysemous, one wants to know why that generalization seems to hold. Maybe pronominals used to refer to human persons are unique, which might prove unsurprising given the messiness of dualist intuitions. Maybe, if pressed about the nature of human mind/brains, speakers might relinquish the inference to the stacked predicate in (26). Notice, for example that not all predicates regarding human person pronominals pattern with (26).

- (27) Esteban is mostly carbon and he is a genius.
- a. [Therefore] Esteban is mostly carbon.
 - b. [Therefore] Esteban is a genius.
 - c. [# Therefore] ?? Something[/Someone] is a mostly carbon genius.

These data are interesting and beg for an explanation. But, the AFC, as a challenge to explain copredication inferences, is not undermined by citing such cases. Maybe human persons serve as exemplar hybrid objects that are less palpably bizarre, leaving room to argue for a denialist ontological account for all flexemes. Importantly, that would require a response to the AFC beyond merely insisting that the most parsimonious ontology posits hybrid objects. One still wants to know why, given such objects, speakers seem to balk at the idea of treating inferences involving hybrid France differently than those regarding non-hybrid dogs (and possible hybrid Esteban).

A second objection pertains to the veracity of the claimed inferential patterns. If the AFC claim is that natural language speakers are apt to make certain inferences from felicitous copredicational sentence but not others, and that these patterns must be explained by any theory of meaning, then that argument plainly rests on an empirical claim about native speakers. One might question whether speakers actually exhibit the inferential patterns. Since our only source of data are the professed intuitions of trained academics, the objection contends, native speaker linguistic competence might diverge from the select population of speakers that research polysemy.

As far as I am aware no formal study has examined the robustness of the inferential generalization claimed by the AFC. So, the support of that empirical claim depends on the acceptability judgments of trained linguists and philosophers. The plausibility of this empirical claim indeed rests on the weight of those expert judgments. Somewhat recent concerns over linguistic methods have called for an examination of the linguistic judgments of experts, which might diverge from native speakers (Gibson and Fedorenko 2010; Gibson and Fedorenko 2013). While some evidence suggests that expert judgments/intuitions are representative of native speaker competencies (Culicover and Jackendoff 2010; Sprouse and Almeida 2013), others are less convinced (see Sprouse, Schütze, and Almeida 2013 for a helpful review). But given the evidence indicated above, there is a case for the general pattern of inference that underwrites the AFC. So dismissing the AFC simply based on

the select group of native speakers that inform the empirical generalization is unwarranted.

What this criticism does point to however, is one possible response to the AFC: undermine the claimed empirical generalization. A fruitful area of research could involve studying the relevant patterns of inference, both to ensure that the claimed generalization holds, and to detail the character of that generalization. Paired with work on the features of copredicational sentences that render them (in)felicitous, a detailed accounting of what speakers are willing to infer from a copredicational sentence, and how willing they are to make those inferences, would likely shed light on copredication and polysemy. But absent findings to the contrary, the inferential pattern indicated above serves as compelling evidence for the AFC's central empirical claim, despite being derived from a privileged English-speaking population.

A final objection concerns the relevance of the copredication inferential patterns to externalism. If the AFC is a challenge to any theory of meaning, one might wonder whether externalism is uniquely troubled by the inferential data. Presumably, an internalist must tell some story as to why (say) the inference from (15) to (15c) is bad, and the inference from (13) to (13c) is not. Surely, whatever that story is, an externalist could appeal to similar considerations, making the data not uniquely troubling for externalism.

As the objection notes, the AFC is a challenge to an theory of meaning. The central claim I've defended is that externalist arguments minimizing the relevance of copredication data to externalism by citing the role of referential intuitions are mistaken. That narrow point does not impugn externalism uniquely with respect to explaining the inferential patterns rehearsed above. However, we saw that the most straight-forward application of externalist theses like (E) fail to capture the inferential patterns in (13)–(15). That failure relates to the constraints that externalism places on the referential behavior of determiner phrases (and nominals). Given those constraints, it's difficult to see how an externalist can provide a semantic solution for the AFC since externalist accounts of anaphora limit available solutions—i.e., the anaphoric *it* and the appropriate antecedent must share that antecedent's meaning-determined referent. If the externalist denies that the issue is semantic, a pragmatic solution to the problem needs to address the differing pattern of judgments in (13), (14) and (15) without abandoning (E). Whatever pragmatic mechanism one appeals to in explaining why (15c) need not be inferred from (15), one wants to know why that mechanism isn't at work in (13) and (14) in the same way. The most salient difference between the

cases, namely the incommensurability of the predicates, does not seem to help the externalist. If the copredicates in (15c) are somehow responsible for the pragmatic-induced infelicity of inferring it from (15) because they cannot be jointly applied to a single object, that same mechanism should be in play for (14). But that's inconsistent with the differing patterns of inference in (14) and (15). Citing these obstacles, of course, does not mean that no externalist solution of the AFC is possible. But they indicate why externalist theses present a special challenge for explaining these data.

5 Conclusion

The relevance of copredication to externalist hypotheses about linguistic meaning is often construed as being a matter of reference. But the import of such semantic flexibility can be pressed without making assumptions about reference or the related speaker intuitions. Denialist responses either implicitly assume or, in the case of Brody and Feiman (2023), explicitly state that addressing the threat of polysemous copredication depends on referential intuitions. The argument from copredication shows that this assumption is mistaken. The inferences that speakers are apt to make with respect to polysemous copredication are the primary data to be explained. Naturally, externalist attempts to capture these data require substantive proposals about reference. But this is a consequence of externalist commitments about what meanings are, not a feature of the data that underwrite the AFC's challenge.

References

- Asher, Nicholas. 2011. *Lexical Meaning in Context: A Web of Words*. Cambridge University Press.
- Austin, J. 1962. *How to Do Things with Words*. Oxford University Press.
- Azzouni, J. 2017. *Ontology without Borders*. Oxford University Press.
- Brody, G., and R. Feiman. 2023. "Polysemy Does Not Exist, at Least Not in the Relevant Sense." *Mind & Language* 39:179–200.
- Carston, Robyn. 1993. "Conjunction, Explanation and Relevance." *Lingua* 90 (1–2): 27–48.

- Chomsky, Noam. 2000. *New Horizons in the Study of Language and Mind*. Cambridge University Press.
- Collins, John. 2009. “Methodology, not Metaphysics: Against Semantic Externalism.” *Proceedings of the Aristotelian Society* Supplementary Volume LXXXIII:53–69.
- . 2017. “The copredication argument.” *Inquiry* 60 (7): 675–702.
- . 2021. “Internalist Perspectives on Language.” In *The Cambridge Handbook of the Philosophy of Language*, edited by Piotr Stalmaszczyk, 157–173. Cambridge University Press.
- . 2023. “Copredication as Illusion.” *Journal of Semantics*, pp. 1–31.
- Culicover, Peter W, and Ray Jackendoff. 2010. “Quantitative methods alone are not enough: Response to Gibson and Fedorenko.” *Trends in Cognitive Sciences* 14 (6): 234–235.
- Davidson, Donald. 1967. “The Logical Form of Action Sentences.” In *The Logic of Decision and Action*, edited by Nicholas Rescher, 81–95. University of Pittsburgh Press.
- Devitt, Michael. 2021. *Overlooking Conventions*. Springer.
- Falkum, Ingrid Lossius, and Agustín Vicente. 2015. “Polysemy: Current Perspectives and Approaches.” *Lingua* 100 (157): 1–16.
- Gibson, Edward, and Evelina Fedorenko. 2013. “The need for quantitative methods in syntax and semantics research.” *Language and Cognitive Processes* 28 (1-2): 88–124.
- Gibson, Edward A, and Evelina G Fedorenko. 2010. “Weak quantitative standards in linguistics research.” *Trends in Cognitive Science* 14 (6): 233–234.
- Glanzberg, Michael. 2014. “Explanation and partiality in semantic theory.” In *Metasemantics: New Essays on the Foundations of Meaning*, edited by A. Burgess and B. Sherman, 259–292. Oxford University Press Oxford.
- Gotham, Matthew. 2017. “Composing criteria of individuation in copredication.” *Journal of Semantics* 34 (2): 333–371.
- . 2022. “Property Inheritance, Deferred Reference and Copredication.” *Journal of Semantics* 39:87–116.

- King, J.C. 2018. “W(h)ither Semantics!(?).” *Nous* 52:772–795.
- Liebman, David, and Ofra Magidor. 2017. “Copredication and property inheritance.” *Philosophical Issues* 27:131–166.
- . 2024. “Ambiguity Tests, Polysemy, and Copredication.” *Australasian Journal of Philosophy* 103 (2): 551–60.
- Liu, Michelle. 2023. “How to Think about Zeugmatic Oddness.” *Review of Philosophy and Psychology*, pp. 1–24.
- Löhr, Guido. 2024. “Does polysemy support radical contextualism? On the relation between minimalism, contextualism and polysemy.” *Inquiry* 67 (1): 68–92.
- Ludlow, Peter. 2011. *The Philosophy of Generative Linguistics*. Oxford University Press.
- Luo, Zhaohui. 2012. “Formal semantics in modern type theories with coercive subtyping.” *Linguistics and Philosophy* 35:491–513.
- Murphy, Elliot. 2023. “The Citadel Itself: Defending Semantic Internalism.” *Global Philosophy* 33, no. 7.
- . 2024. “Predicate Order and Coherence in Copredication.” *Inquiry* 67 (6): 1744–1780.
- Ortega-Andrés, Marina. 2023. “The Denotation of Copredicative Nouns.” *Erkenntnis* 88 (7): 3113–3143.
- Pietroski, Paul M. 2025. “One Word, Many Concepts: Endorsing Polysemous Meanings.” In *Oxford Handbook to Contemporary Philosophy of Language*, edited by E Lepore and U Stojnić, 258–286. Oxford University Press.
- Pietroski, Paul M. 2005. “Meaning Before Truth.” In *Contextualism in Philosophy*, edited by Georg Peter Gerhard Preyer, 255–302. Oxford University Press.
- Pustejovsky, J. 1995. *The Generative Lexicon*. MIT Press.
- Quilty-Dunn, Jake. 2021. “Polysemy and Thought: Toward a Generative Theory of Concepts.” *Mind & Language* 36 (1): 158–185.
- Sprouse, Jon, and Diogo Almeida. 2013. “The empirical status of data in syntax: A reply to Gibson and Fedorenko.” *Language and Cognitive Processes* 28 (3): 222–228.

- Sprouse, Jon, Carson T Schütze, and Diogo Almeida. 2013. “A comparison of informal and formal acceptability judgments using a random sample from Linguistic Inquiry 2001–2010.” *Lingua* 134:219–248.
- Stanley, Jason. 2007. *Language in Context: Selected Essays*. Oxford University Press.
- Vicente, Agustín. 2018. “Polysemy and Word Meaning: An Account of Lexical Meaning for Different Kinds of Content Words.” *Philosophical Studies* 175 (4): 947–968.
- . 2019. “Chomskyan Arguments Against Truth-Conditional Semantics Based on Variability and Co-predication.” *Erkenntnis* 86 (4): 919–40.
- . 2021. “Approaches to Co-predication: Inherent Polysemy and Metaphysical Relations.” *Journal of Pragmatics* 182:348–357.
- Viebahn, Emanuel. 2018. “Ambiguity and Zeugma.” *Pacific Philosophical Quarterly* 99 (4): 749–762.
- . 2022. “Copredication, Polysemy and Context-Sensitivity.” *Inquiry* 65 (8): 1066–1082.
- Xue, Tao, and Zhaohui Luo. 2012. “Dot-Types and Their Implementation.” *Logical Aspects of Computational Linguistics: 7th International Conference*, Volume 7. Springer, 234–249.
- Yalcin, Seth. 2014. “Semantics and metasemantics in the context of generative grammar.” In *Metasemantics: New Essays on the Foundations of Meaning*, Volume 17, 17–54. Oxford University Press Oxford.