

## ***One Word, Many Concepts: Endorsing Polysemous Meanings (28Jan2023)***

**Paul M. Pietroski, Rutgers University**

for the *Oxford Handbook to Contemporary Philosophy of Language* (E. Lepore and U. Stojnić, eds.)

The phenomenon of polysemy, often illustrated with words like ‘book’ or ‘window’, tells against the idea that human linguistic meanings determine extensions; see, e.g., Chomsky (1995, 2000) and Collins (2017). If ‘book’ has one sense that corresponds to certain abstract *contents* and another sense that corresponds to certain concrete *vehicles* for contents, then as discussed below, dubious maneuvers are required to maintain that some set of things is the extension of ‘book’. Likewise, if ‘window’ has a sense that corresponds to certain *apertures* and a distinct sense that corresponds to certain *panes* (perhaps framed for installation into apertures of the relevant sort), then dubious maneuvers are required to maintain that ‘window’ has an extension.<sup>1</sup>

Pietroski (2005, 2010, 2014, 2017a, 2018, 2021) connects this point with other reasons for denying that words have extensions. The proposed alternative is that meanings are recipes for how to generate distinctive *concepts*—mental symbols of a special sort—from a stock of basic concepts that can be accessed via lexical addresses; where a single address can be shared by a family of concepts that are not coextensive. On this view, a word that is often used to fetch a concept C may be used on other occasions to fetch a related but distinct concept C\*.

Sentences like (1) and (2) can help motivate this idea.

(1) The window they had cut in the wall was nicer than the one they installed.

(2) The book she brought weighed two pounds, but it had been well reviewed.

It seems that in an utterance of (1), a speaker can use ‘window’ and the anaphoric pronoun ‘one’ to express—and invite a hearer to access—distinct concepts from her ‘window’-family.

Similarly, it seems that in an utterance of (2), ‘The book she brought’ and ‘it’ can be used to express concepts whose constituents include different concepts from the ‘book’-family.

Still, one might hope to salvage the hypothesis that meanings determine extensions by saying that extensions include mappings from contexts to senses that determine more classical extensions. Perhaps in contexts where (1) is used to make a coherent claim, ‘window’ has its aperture-sense, and ‘one’ has the pane-sense of its antecedent. Appealing to contexts in this way risks the charge of saving a dogma by trivializing it. But in my view, the problem is worse: the meaning of a polysemous expression does not determine a mapping from contexts to senses. Correlatively, I claim, sentence meanings do not map contexts to propositions (or truth values).

In this paper, my central argument will focus on (3), which can be used in multiple ways.

(3) The window that Click preferred was nicer than the one that Clack preferred.

The key point is that a speaker of English can use (3) to make an assertion, and thereby exclude some possibilities that had been left open, without using ‘window’ to express *any* sense. A speaker can use (3) without knowing or indicating whether the windows in question were apertures, fillers, or one of each. Expressions can be used in conceptually *neutral* ways. Such uses are, I contend, unsurprising if expression meanings are recipes for how to build concepts. Sections two through four provide the details. In section five, I suggest that the conclusion applies even to words like ‘water’ and other so-called kind terms. Section one addresses some preliminary matters, including my uses of expressions like ‘word’, ‘language’, and ‘meaning’.

---

<sup>1</sup> If it helps, think of a sense as a “use” of an expression that has a single meaning. I return to the implausible idea that polysemy is a species of homophony, and that each expression has a single sense. Vincente and Falkum (2017) provide a helpful introduction to the topic of polysemy and previous literature; see also Falkum and Vincente (2015), and for discussion of relevant psycholinguistic experiments, Quilty-Dunn (2021) and Murphy (2021).

## 1. Words, Kinds, and Human Languages

Like most words, ‘word’ is polysemous. A person who promised to tell others about a party might keep her word by spreading the word, and so do by using words. In this respect, ‘word’ is like ‘book’, which can be used to talk about contents that can be downloaded and plagiarized, or about bound volumes that can be kicked and defaced. But let’s focus on words of the human languages, spoken or signed, that children regularly acquire.

These words are grammatically diverse. Perhaps ‘mud’ and the untensed verb ‘tell’ are atomic. But consider ‘told’, ‘broken’, ‘windows’, ‘mailbox’, etc. A polysynthetic language like Mohawk allows for sentences like (4), which means that Shawatis gave Uwari a basket.

(4) Uwari washakotheru' ne Shawatis.

The second word in (4) has a complex meaning that speakers of English might transliterate with ‘he/her-basket-give’. More dramatically, the one-word sentence displayed as (5) means roughly that he made the thing that you put on your body ugly for her; see Baker (2002).

(5) Washakoty'a'tawitsherahetkvhtv'

Few if any such words get memorized. Their meanings can be worked out, compositionally, as needed. But if only for this reason, words seem more like a hodge-podge than a theoretically interesting kind. Perhaps ‘told’ and ‘mailbox’ are memorized and in that sense *lexical* items; though if the tense and plural morphemes in ‘broke windows’ are lexical items, not all such items are words. Similarly, the count noun ‘rabbit’ may comprise a count morpheme and a simpler noun that can be used to talk about rabbit. In this sense, the polysemy of a word can reflect grammatical structure.<sup>2</sup> But for present purposes, let’s focus on lexical items that are plausibly sources of the polysemy exhibited by more complex expressions.

The count noun ‘language’ is also polysemous. But in this case, I think there is an interesting kind in the vicinity. In acquiring a spoken or signed language, a human child acquires lexical items that can be combined to form complex expressions, most of which will never be formed. Yet however many expressions the child memorizes, she will construct far more in speaking, thinking, and understanding others.<sup>3</sup> This suggests that the languages in question are *procedures that generate* expressions, and that these procedures can be biologically implemented in humans. To be sure, one can also use ‘language’ to talk about sets of expressions, while insisting that a person can acquire a set *by* acquiring a procedure that generates its elements. But if humans regularly acquire expression-generating procedures of a special sort, that is potentially relevant for purposes of characterizing polysemy in a theoretically perspicuous way.

In acquiring English, a child comes to have a mind that implements a procedure that generates expressions that are meaningful and pronounceable. But a child who acquires English in Brooklyn doesn’t acquire the very same procedure as a child who acquires English in Sydney, Glasgow, or another part of Brooklyn. So let’s not assume that ‘English’ is a name for a shared procedure that generates expressions whose meanings are determined by tacit conventions. We can use ‘English’ to talk about certain pubs, the history of The English Pub, or a family of

---

<sup>2</sup> See Bromberger (2011) for related discussion in reply to Kaplan (1990), Hawthorne and Lepore (2011). We can introduce a notion of *phonological-word* that covers ‘a’ and (5), or more restricted notions that exclude compounds of nouns and verbs; a notion of *semantic-word* that includes lexical items and combinations of open-class items with functional elements, but excludes the pleonastic subject of ‘It rained’; a notion of *orthographic-word* that treats ‘run’ as a single word with many meanings, or more liberal notions that allow for homographs; and so on.

<sup>3</sup> *Prima facie*, most episodes of constructing linguistic expressions occur in intrapersonal cognition, and not just because of silent soliloquy. In understanding someone else as having uttered the sentence ‘Language is essentially for communication’, I form that complex expression, even if I would never use it to make a sincere assertion.

languages that turn out to be procedures implemented in human heads.<sup>4</sup> These procedures are similar enough to support communication, to the degree they do, among good-willed users. But even among good-willed members of the same family, speaking is no guarantee that hearers will form the thoughts expressed. Theorists should not assume otherwise when studying the expression-generating procedures that humans naturally acquire. Two speakers of English may connect a given pronunciation with different meanings—think of ‘lift’, ‘biscuit’, and ‘robin’—just as they may connect the same meaning with different pronunciations.

In order to talk about expression-generators without relying exclusively on ordinary locutions, Chomsky (1986) introduced the technical term ‘I-language’. He used the contrasting term ‘E-language’ to cover any other things that might be called languages—e.g., sets of expressions or normatively governed dispositions to speak (see Quine 1960, Lewis 1975). Church (1941) had contrasted functions *in intension* with functions *in extension*; the point was that distinct intensions, corresponding to different algorithms that could be implemented on a Turing Machine, may determine the same set of input-output pairs. I-languages are intensional in this sense, though the prefix ‘I-’ also connotes ‘internal’ and ‘individualistic’.

Given this terminology, one can say that acquiring a spoken-or-signed human language is a matter of coming to implement an I-language that generates certain expressions and thereby connects some meanings with pronunciations in a particular way. This doesn’t yet show that the I-languages in question exhibit deep commonalties. But a long line of investigation—initiated by Chomsky (1957, 1959, 1964, 1965) and extended by many others—confirms the old idea that for humans, grammatical variation reflects different ways of filling in an innately determined Universal Grammar that corresponds to a distinctively human kind of generative procedure.

The central discovery was that human I-languages respect substantive *constraints* on how they connect meanings with pronunciations. This matters. For while ‘meaning’ is polysemous with a vengeance (see, e.g., Grice 1989), the meanings that human I-languages connect with pronunciations conform to principles that are far from obvious. So at the risk of making an introductory section too long, let me offer a few examples of how human I-languages *don’t* link meanings to pronunciations, since this will help distinguish the meanings I am talking about from other things (e.g., concepts or extensions) that might get called meanings in other contexts.

The words listed in (6) can be combined to form the declarative sentences (7) or (8).

(6) the, who, was, in, hiker, lost, kept, walking, circles

(7) The hiker who was lost kept walking in circles.

(8) The hiker who lost was kept walking in circles.

The words in (6) can also be combined to form the polar (yes/no) interrogative sentence (9).

(9) Was the hiker who lost kept walking in circles?

But while (7) is more expected than (8), given the list of constituent words, (9) can only be understood as the yes/no question corresponding to (8). In this respect, (9) is interestingly *unambiguous*; it fails to have the polar question meaning that corresponds to (7).

Similarly, (10) must be understood as the polar interrogative correlate of (11).

(10) Was the guest who fed waffles fed the parking meter?

(11) The guest who fed waffles was fed the parking meter.

---

<sup>4</sup> Used in this taxonomic way, ‘English’ is like ‘Germanic’, which can be used to talk about a family of languages that include versions of Norwegian. Dummett (1986) suggests that every speaker of English speaks it imperfectly, without thereby speaking Japanese very badly. This strikes me as a misleading way of saying that for purposes of communication, there are normative constraints on how those who have acquired one or more English languages/procedures should use them. But perhaps there are social abstracta that deserve to be called languages.

Despite what we know about guests, feeding, waffles, and parking meters, (10) cannot be understood as the perfectly reasonable yes/no question corresponding to (12).

(12) The guest who was fed waffles fed the parking meter.

This makes the unambiguity of (10) quite interesting, since human I-languages permit ambiguity.

The pronunciation indicated with ‘an artistic sheriff drew a gun near a bank’ is shared by at least eight English sentences, each with its own meaning. The pronunciation indicated with ‘bank’ (a.k.a. /bæŋk/) is shared by at least two relevant nouns, and the pronunciation indicated with ‘drew’ (a.k.a. /dru/) is shared by at least two relevant verbs. We can mark such distinctions by adding superscripts to standard orthography as in ‘bank<sup>S</sup>’, ‘bank<sup>V</sup>’, ‘drew<sup>✓</sup>’, and ‘drew<sup>✗</sup>’. But however these *lexical* homophonies are resolved, the prepositional phrase ‘near a bank’ can modify the noun ‘gun’ or the verb phrase ‘drew a gun’. We can mark this *structural* homophony with brackets as in [drew<sup>✓</sup> [a [gun [near [a bank<sup>S</sup>]]]]] vs. [[drew<sup>✓</sup> [a gun]] [near [a bank<sup>S</sup>]]].<sup>5</sup> So a speaker of English who knows the relevant words has acquired an I-language that generates at least eight sentences that share the pronunciation of ‘an artistic sheriff drew a gun near a bank’, thereby connecting this pronunciation with at least eight sentential meanings.

It is striking that such an I-language, which imposes no upper bound on how many ways a pronunciation can be understood, won’t let the pronunciation of (10) have the meaning of (13).

(13) The guest who was fed waffles fed the parking meter?

And this is not a parochial feature of a few human I-languages. The relevant constraint, which precludes displacing an auxiliary verb from a relative clause, seems to reflect the mental systems that determine the class of expression-generating procedures that humans can naturally acquire.<sup>6</sup>

Linguists have discovered many such constraints, thus bolstering the case that humans acquire I-languages of an interesting kind. The meanings that these languages connect with pronunciations may also exhibit deep psychological, computational, and even biological commonalities; see Pietroski (2018). But whatever human linguistic meanings are, they allow for homophony and polysemy: distinct meanings can share a pronunciation, subject to constraints; and a single meaning can be associated with distinct senses or concepts. Moreover, lexical polysemy can be interestingly related to structural homophony.

For example, the string of words shown in (14) can be understood in two ways.

(14) the duck was ready to eat

(14a) The duck was prepared to be one that *does some eating*.

(14b) The duck was fit to be one that *gets eaten*.

But (15) has only the duck-as-eater meaning paraphrased with (15a).

(15) the duck was eager to eat

(15a) The duck was eager to be one that does some eating.

(15b) #The duck was eager to be one that gets eaten.

Likewise, if ‘eager’ is replaced with ‘reluctant’. By contrast, (16) has only the duck-as-eaten meaning paraphrased with (16b); cp. Chomsky (1964).

(16) the duck was easy to eat

(16a) #It was easy for the duck to eat (something).

(16b) It was easy (for others) to eat the duck.

---

<sup>5</sup> The ambiguity of ‘drew’ is lexical, even though in each case, the verb is tensed. But ‘a name’ and ‘an aim’ share a pronunciation, for structural reasons, without having homophonous lexical constituents; see Chomsky (1957).

<sup>6</sup> See Ross (1967) and much work since. Berwick et. al. (2011) revisit the corresponding “poverty of stimulus” arguments, and reasons for maintaining that meanings correspond to ways of deriving strings from grammars, in light of developments since Chomsky (1957, 1965, 1980); see also Pietroski (2017b, 2018).

So unlike (16) or (15), (14) is structurally homophonous: ‘the duck’ can be understood as linked to the subject position of the embedded verb ‘eat’, or as linked to the object position of that verb.

This formal ambiguity highlights the fact that ‘ready’ is polysemous in a subtle way; cp. Recanati (2004). It can be used to express a concept of psychological readiness—being mentally prepared for what is to come—or a concept of being fit for use; ‘ready for battle’ may reflect a mixture. Used in the first way, ‘ready’ is like ‘eager’ in implying a mind with an attitude. Used in the second way, ‘ready’ is like ‘easy’, which carries no such implication. Correlatively, the grammatical subjects of (14a) and (15) are linked to the subject (or “agent”) position of ‘eat’, but the grammatical subjects of (14b) and (16) are linked to the object (or “patient”) position.<sup>7</sup>

In describing (14) as a case of structural homophony, I assumed—quite standardly—that ‘ready’ is a constituent of both (14a) and (14b). The two readings of (14) correspond to two ways of arranging those words: one with ‘the duck’ linked to the subject/agent position of ‘eat’; the other with ‘the duck’ linked to the object/patient position of ‘eat’. But both ways of generating an expression with the pronunciation of (14) involve ‘ready’. It’s not that there are distinct words, ‘ready<sup>1</sup>’ and ‘ready<sup>2</sup>’, that share a pronunciation but have distinct meanings.

Similarly, it’s not that there is one word ‘window’ that gets used to access concepts of apertures and a distinct but homophonous word that gets used to access concepts of fillers. This would imply that in (1), the pronoun ‘one’ can have an antecedent that does not appear in (1).

(1) The window they had cut in the wall was nicer than the one they installed.

But that’s not how antecedence works in human I-languages. Consider (17).

(17) The bank we swam towards was wetter than the one where the teller worked.

This sentence cannot be construed as implying that some bank<sup>∇</sup> was wetter than some bank<sup>§</sup>. To insist that the ‘window’-homophones are special homophones, in a way that supports anaphora in (1) but not (17), is to use ‘special homophone’ as a device for talking about senses of expressions—as opposed to expressions, each of which has a single meaning that allows for multiple senses. But even given this misleading way of talking about human I-languages, ‘bank<sup>§</sup>’ would remain ambiguous: certain buildings are banks<sup>§1</sup>; and institutions like Chase Manhattan are banks<sup>§2</sup>. Moreover, it seems that ‘bank<sup>§1</sup>’ would remain ambiguous: one can walk into a bank<sup>§1a</sup> by walking through its front door; and one can walk into a bank<sup>§1b</sup> by bumping into it.<sup>8</sup>

In general, polysemy outruns homophony, given any independently plausible way of individuating lexical items and their meanings; see Caplan et. al. (2020), Srinivasan and Berner (2019). This isn’t surprising, given that each child is regularly faced with the task of deciding whether other speakers are using lexical items that the child has already acquired—perhaps using them in unexpected ways—or new lexical items that happen to be homophonous with old ones.

---

<sup>7</sup> The subject of (15/15a) is said to “control” the unpronounced subject/agent position of ‘eat’, while the pleonastic subject of (16b) leaves that unpronounced position free to be construed as the relevant eaters. Relatedly, (15) can be paraphrased with ‘The duck was eager to be one who eats something’, with ‘eats something’ understood as implying ‘eats something suitable for it’. But if a duck ate a tack, and in that sense ate/ingested something, it didn’t thereby *eat* (dine, or fuel up) unless it was an odd duck that can digest tacks. So it seems that ‘eat’ can be used to express distinct concepts, one of which is subtly *normative*; see Chomsky (1986) who credits the point to Howard Lasnik.

<sup>8</sup> Repeat for ‘door’ and ‘walk through’. Similarly, in ‘France is hexagonal, and it is a republic’, the antecedent for ‘it’ does not fail to appear in the sentence; see Pietroski (2005). So *pace* King (2018) and Devitt (2020, 2021), it wouldn’t help to say that the pronunciation of ‘France’ is shared by homophonous proper nouns, one of which names a certain chunk of terrain, while the other names a certain polis. To my ear, even if the conjunction with ‘it’ is slightly odd, it is less odd than ‘France is a hexagonal republic’ or ‘France is a republic and a hexagonal one’. This suggests that ‘France’ and the linked pronoun ‘it’ can be used to access distinct concepts: one that applies to certain terrain but not the corresponding nation, and one that applies to the nation but not the corresponding terrain.

The child's default assumption cannot be that new uses of words are uses of new words. Resorting to homophony is presumably an option only given evidence that outweighs the presumption that foreign uses of familiar lexical pronunciations are uses of familiar lexical items.

With these points in mind, let's return to examples like (1) and (2).

(1) The window they had cut in the wall was nicer than the one they installed.

(2) The book she brought weighed two pounds, but it had been well reviewed.

Words like 'book' and 'window' suggest that a lexical item can be linked to two or more concepts that somehow "belong together." Indeed, children may use lexical items in the process of acquiring and connecting concepts that dovetail given human minds and purposes.

At least for these purposes, let's take concepts to be mental symbols with which one can think about things in certain ways; see Fodor (1998). In thinking about something *as* a book that can be downloaded and plagiarized, a thinker employs a concept (of books) that we can call a content-concept; and in thinking about something *as* a book that can be kicked and defaced, a thinker employs a concept (of books) that we can call a vehicle-concept. The content-concept presumably differs from the vehicle-concept, even if the books that can be downloaded can also be kicked; cp. concepts of people and their bodies. But at least for human thinkers, it seems natural to have both the content-concept and the vehicle-concept if one has any concept of books.

It seems even clearer that an aperture-concept of certain holes (that we can call windows) can differ from a filler-concept of certain panes (that we can call windows). It's hard to see how an aperture-concept could let one think about holes-*or-fillers*, much less think about them in the same way that some filler-concept lets one think about *holes-or-fillers*. And presumably, a child might initially link 'window' to an aperture-concept that doesn't apply to the fillers, or initially link 'window' to a filler-concept that doesn't apply to the holes. Though at least typically, both concepts will eventually be acquired and brought together under a shared lexical heading, in a way that may promote extended uses of these concepts and/or acquisition of new concepts.<sup>9</sup>

It might be that we couldn't, or least wouldn't, think about a sizable hole in a large rock as a window if we didn't already think about certain holes in walls as windows. And we might not think about panes of glass as windows if we didn't think about them as potential fillers of relevant holes. Other uses of 'window'—e.g., in describing certain storefront display spaces or certain gaps in counters for transactions with bank tellers—seem to reflect more extended uses of polysemous words to express other concepts that get treated as members of the same family. (A speaker might even form a permissive concept that applies equally to windowpanes and the holes they can fill.) One can say that such uses are secondary, in the sense of depending on hole/filler uses already being in place, without saying that secondary uses are metaphorical or non-literal in any further sense; see Carston (2019, 2021) and section five below. Relatedly, uses that sound poetical to adults can be grist for the next generation's extended polysemy.

As one more example, consider 'line', which can be used to talk about perceptible things drawn on a board (or in sand) or imperceptible things that conform to geometric theorems. The perceptible lines have width, and the imperceptible ones don't. This suggests that 'line' can be used to access one or more concepts that apply to certain perceptible things—perhaps including lines in faces, waiting lines, and fishing lines—but also used to access one or more concepts that

---

<sup>9</sup> Quilty-Dunn (2021) suggests that concepts *point to* complex contents, and that 'window' corresponds to a single concept. But this is misleading if the pointers are meanings, and the contents are represented with families of mental symbols. If one takes a concept to be a lexical item as used to access a mental representation, the issue will be whether the accessed representations can be used without the lexical item to think about things in systematic ways.

apply to certain imperceptible things and let us think about such things in ways that let us *deduce consequences* (e.g., that doubling the sides of a square quadruples the area), at least when prompted by Socratic chat and perceptible diagrams.

By contrast, it's hard to see how 'line' could have a single extension-determining sense that covers all the things we use that word to talk about. So while the word's *meaning* is neutral, in the sense of permitting varied uses, we shouldn't infer that this meaning determines a neutral extension. On the contrary, whatever sense of 'line' one uses on a given occasion of uttering 'the lines' or inscribing '{x: x is a line}', it seems that any set specified in this way would be too parochial to be the alleged extension of 'line'. Parallel remarks apply to 'window' and 'book'. As noted at the outset, one might hope to salvage the idea that meanings determine extensions by saying that extensions include mappings from contexts to senses that determine more classical extensions. But if it's granted that words are polysemous, then blocking the salvage operation isn't hard. As argued in the next sections, contexts don't map words to senses.

## 2. Installing Windows

Suppose that Click and Clack bought a framed pane of glass at a hardware store, and then installed it into a hole they had cut in the west wall of a cottage they were building. In the end, Click and Clack disagreed about whether the window they cut in the wall was nicer than the one they installed. One of them preferred the original hole to the framed pane of glass they put in it; the other one preferred the framed pane of glass to the hole. (If it helps, imagine that the hole had been finished off and pleasantly shuttered for months before the rainy season arrived.) We don't know who preferred what. We never saw the original hole or the final result. But we know that whenever Click and Clack disagree, we agree with Click. So we endorse (3).

(3) The window that Click preferred was nicer than the one that Clack preferred. But we're agnostic about whether Click preferred the original hole or the framed pane of glass.

In this context, we are not using 'window' to access one of its senses and then accessing another sense via the anaphoric pronoun 'one'. Our use of 'window' is more neutral. Still, by endorsing (3), we're making an assertion and thereby rejecting both (a) the possibility that Click preferred the original hole, but the framed pane of glass was nicer, and (b) the possibility that Click preferred the framed pane of glass, but the original hole was nicer. We're not, however, rejecting either (a') the possibility that Click preferred the original hole, and it was nicer than the framed pane of glass, or (b') the possibility that Click preferred the framed pane of glass, and it was nicer than the original hole. On the contrary, even given our belief that Click was right and Clack was wrong, both (a') and (b') remain as "live epistemic possibilities" for us.

It's hard to see how any proposition plausibly *expressed* by (3), an expression of a human I-language, could carve the relevant space of possibilities in this way. That isn't a problem if (3) doesn't express a proposition. But one wants to know how a speaker can use a sentence assertively and thereby reject some possibilities—along lines suggested Stalnaker (1978) and many others—if the sentence doesn't express a proposition that determines a set of possibilities.

If only to help clarify the issues, let 'Holy' and 'Philly' be names for the relevant hole and filler. Let 'Clickbait' and 'Clackbait' be names for the window that Click preferred and the window that Clack preferred. In our scenario, either Holy is Clickbait or Holy is Clackbait; but we don't know whether Holy is Clickbait or Clackbait. Likewise, we don't know whether Philly is Clickbait or Clackbait. But let's grant that the invented sentences (18-23) express propositions.

- |                                  |                                  |
|----------------------------------|----------------------------------|
| (18) Nicer(Holy, Philly)         | (19) Nicer(Philly, Holy)         |
| (20) Nicer(Clickbait, Philly)    | (21) Nicer(Clickbait, Holy)      |
| (22) Nicer(Clickbait, Clackbait) | (23) Nicer(Clackbait, Clickbait) |

Let's also assume that either Holy was nicer than Philly, or vice versa. Then (18) or (19) is correct, but we don't know which. *Given* that Click was right, (20) or (21) is correct; and (22) is correct. But (3) does not express any of these propositions. Even if it is a priori that (22) is correct if and only if (3) is, (22) has a different modal profile than (3)—assuming that the names in (22) designate rigidly in Kripke's (1980) sense, while the definite descriptions in (3) do not.

Let's also grant that (24-27) express propositions.

(24) Identical(Holy, Clickbait) & Nicer(Holy, Philly)

(25) Identical(Philly, Clickbait) & Nicer(Philly, Holy)

(26) Identical(Philly, Clackbait) & Nicer(Clickbait, Philly)

(27) Identical(Holy, Clackbait) & Nicer(Clickbait, Holy)

The possibilities left open by endorsing (3), in the imagined context, can be described with the disjunction of (24) and (25). But this doesn't show that relative to the imagined context, (3) expresses what (28) expresses, even though (28) is the disjunction of (24) and (25).

(28) [Identical(Holy, Clickbait) & Nicer(Holy, Philly)] v  
[Identical(Philly, Clickbait) & Nicer(Philly, Holy)]

The possibilities left open can also be described with (29), the disjunction of (26) and (27).

(29) [Identical(Philly, Clackbait) & Nicer(Clickbait, Philly)] v  
[Identical(Holy, Clackbait) & Nicer(Clickbait, Holy)]

But if the conjecture that (3) expresses a proposition leads to the conclusion that (3) expresses what (29) expresses, we should reject the conjecture. Likewise if the conjecture leads to the conclusion that (3) expresses what (28) expresses.

Put another way, we can use (30) to specify the possibilities in which Click was wrong.

(30) [Identical(Holy, Clickbait) & Nicer(Philly, Holy)] v  
[Identical(Philly, Clickbait) & Nicer(Holy, Philly)]

So we can specify the four relevant possibilities with the disjunction of (28) and (30). This makes it easy to see how *our belief that Click was right* excludes the possibilities specified with (30), leaving the possibilities specified with (28). But it's very hard to see how the human sentence (3) could do this possibility-excluding work, even in the imagined situation.

(3) The window that Click preferred was nicer than the one that Clack preferred.

Whatever the logical form of (3) is, it's not a disjunction, much less a disjunction of propositions about Holy and Philly (or holes and fillers).

One might think that (31) is a better model for the proposition expressed with (3);

(31) Nicer{ $\lambda x$ [Windo(x) & ClickPreferred(x)],  $\lambda x$ [Windo(x) & ClackPreferred(x)]}

where by stipulation, 'Windo' applies to *both* the unsmashable holes that count as windows *and* the smashable fillers that count as windows (cp. 'dog', poodles, and terriers). But even if the definite descriptions in (3) are quantificational in the way that (31) suggests, this does not yet reflect the polysemy of 'window' or the anaphoric use of 'one' in (3). And *prima facie*, 'Windo' is importantly unlike 'window'. For example, imagine Philly lying on the floor near Holy, before Philly was installed. It seems wrong to say that there were two windows in that room, one on the floor and one in the wall. We typically count windows relative to a single (de-polysemized) use of 'window', as if the requisite sortal is provided by a specific concept/sense and not the word itself. Similarly, if someone is carrying a single volume that includes two novels by Jane Austen, it seems wrong to conclude that the person is carrying three books. But if 'Windo' is a good model for 'window', then other things equal, (32) should be as true as (33).<sup>10</sup>

<sup>10</sup> Liebesman and Magidor (2019) object to a certain way of using this consideration to draw metaphysical conclusions. But this does not undercut the arguments offered here, or in Chomsky (2000) and Collins (2011).



(32) There are two windows in the room.

(33) Windo(Holy) & Windo(Philly) & ~Identical(Holy, Philly)

One can insist that users of ‘Windo’ would still be able to sort/count in ways that make (32a) infelicitous, even if Philly lies uninstalled on the floor; cp. counting pairs of gloves.

(32a)  $\exists X\{\text{Two}(X) \ \& \ \forall x: Xx[\text{Windo}(x)] \ \& \ \text{InTheRoom}(X)\}$

But it’s hard to see why users of ‘Windo’ would find it difficult to count Holy and Philly as two Windos when Holy and Philly are far apart. Similarly, one would expect users of ‘Windo’ to find it easy to use ‘every Windo’ as a quantifier whose domain includes both Holy and Philly; yet I find it hard to use ‘every window’ this way, as if the requisite sortal is provided by a specific concept/sense rather than the word itself. By themselves, these points about quantifying are not decisive. But they bolster the initial arguments based on examples like (1) and (2).

(1) The window they had cut in the wall was nicer than the one they installed.

(2) The book she brought weighed two pounds, but it had been well reviewed.

Taken together, such examples tell against treating the pronouns as second occurrences of a word that has a permissive extension whose elements are as metaphysically diverse as holes and fillers, vehicles and contents, or all things that we call lines.

### 3. Recipes and Choices

Endorsing (3) need not be a matter of endorsing a proposition expressed with (3).

(3) The window that Click preferred was nicer than the one that Clack preferred.

Endorsing (3) may be more like endorsing a *recipe* that can lead to various acceptable results, depending on which choices—regarding specific ingredients—are made in executing the instructions. For example, the classic recipe for a Negroni calls for equal parts of Campari, gin, and red vermouth. But this doesn’t specify which gin, or which vermouth, is to be used. So one can recommend a Negroni without knowing which specific ingredients will be used. I return to this analogy in §4. For now, the point is just that we can describe (3) as a recipe for how to assemble a thought, as opposed to a device for expressing a particular thought.

Instead of saying that ‘window that Click preferred’ is a phrase whose head noun has an extension that the relative clause can narrow, we can describe the phrase as an instruction for how to build a complex concept. And the phrase can be a recipe for how to build a concept, even if no concept *C* is such that the phrase is a recipe for how to build *C*. If ‘window’ can be used to access a concept of certain holes *or* a concept of certain fillers, then there are at least two concepts that can be assembled by executing the phrasal recipe ‘window that Click preferred’. And likewise for ‘one that Clack preferred’, with ‘one’ understood as anaphoric on ‘window’.

Suppose that  $C_{\text{hole}}$  is the concept that applies to (and only to) things that are holes of the relevant sort and preferred by Click, and  $C_{\text{filler}}$  is the concept that applies to (and only to) things that are fillers of the relevant sort and preferred by Click. Then  $C_{\text{hole}}$  applies to Holy or nothing, depending on whether or not Click preferred Holy. Similarly,  $C_{\text{filler}}$  applies to Philly or nothing, depending on whether or not Click preferred Philly. Suppose that  $K_{\text{hole}}$  and  $K_{\text{filler}}$  are the corresponding concepts of holes/fillers preferred by Clack. Then if Click preferred Holy, ‘window that Click preferred’ can be used to build a concept that in fact applies to Holy, while ‘one that Clack preferred’ can be used to build a concept that in fact applies to Philly. Similarly, if Click preferred Philly, ‘window that Click preferred’ can be used to build a concept that in fact applies to Philly, while ‘one that Clack preferred’ can be used to build a concept that in fact applies to Holy. But it doesn’t follow that ‘window that Click preferred’ can be used to build a concept that meets *both* of the following conditions: it applies to Holy if and only if Click preferred Holy; *and* it applies to Philly if and only if Click preferred Philly.

On the contrary, if we don't have concepts that are as permissive as the invented predicate 'Windo(x)', then 'window that Click preferred' cannot be used to build a concept that is as permissive as 'Windo(x) & ClickPreferred(x)'. But we don't need such concepts to make sense of polysemous phrases like 'window that Click preferred'. We just have to drop the idea that phrases *directly express* concepts in favor of the idea that such phrases *direct assembly of* concepts, and then recognize that sometimes, a direction can be followed in two or more ways.

To repeat, we don't need to say that 'window that Click preferred' corresponds to a concept that applies to Holy or Philly depending on which of these two things Click preferred. Likewise, we don't need to say that 'one that Clack preferred' corresponds to a concept that applies to Holy or Philly depending on which of these two things Clack preferred. We can say that 'window that Click preferred' can be used to build  $C_{\text{hole}}$ , and that this same phrase can be used to build  $C_{\text{filler}}$ . Likewise, 'one that Clack preferred' can be used to build either  $K_{\text{filler}}$  or  $K_{\text{hole}}$ .

In this sense, 'window that Click preferred' and 'one that Clack preferred' do not correspond to constituents of a proposition expressed with (3). But these noun phrases can be used to build  $C_{\text{hole}}$  and  $K_{\text{filler}}$ , respectively, or  $C_{\text{filler}}$  and  $K_{\text{hole}}$ . So (3) might be used to build a mental analog of (34) and/or a mental analog of (35);

(3) The window that Click preferred was nicer than the one that Clack preferred.

(34)  $\text{Nicer}\{\text{ix}[W(x) \ \& \ \text{ClickPreferred}(x)], \text{ix}[W^*(x) \ \& \ \text{ClackPreferred}(x)]\}$

(35)  $\text{Nicer}\{\text{ix}[W^*(x) \ \& \ \text{ClickPreferred}(x)], \text{ix}[W(x) \ \& \ \text{ClackPreferred}(x)]\}$

where 'W' applies only to the relevant holes and 'W\*' applies only to the relevant fillers.

In scenarios where we knew that Click and Clack had disagreed about which of two holes was nicer, or about which of two framed windowpanes to buy from the hardware store, we might use (3) to build a mental analog of (36) and/or (37).

(36)  $\text{Nicer}\{\text{ix}[W(x) \ \& \ \text{ClickPreferred}(x)], \text{ix}[W(x) \ \& \ \text{ClackPreferred}(x)]\}$

(37)  $\text{Nicer}\{\text{ix}[W^*(x) \ \& \ \text{ClickPreferred}(x)], \text{ix}[W^*(x) \ \& \ \text{ClackPreferred}(x)]\}$

Uses of (3) need not be co-predicational. But even if uses of (3) to build thoughts like (36) or (37) are simpler, or more common, we shouldn't conclude that (3) expresses a proposition in the simpler or more common cases. Nor should we conclude that co-predicational uses are semantically complex. Such uses merely highlight the diverse ways in which natural linguistic expressions can be used. This diversity tells against the bold hypothesis that ordinary sentences express propositions. But this diversity is compatible with the modest claim that speakers often use ordinary sentences to express propositions, or least to exclude possibilities, in contexts.

Still, one might hope that the bold hypothesis can be saved given modern semantic technology. Perhaps 'window' and 'one' come with indices as suggested by (38);

(38) The *window* <sub>$\Delta$</sub>  that Click preferred was nicer than  
the *one* <sub>$\delta$</sub>  that Clack preferred.

where the italics indicate that 'window' is the antecedent for 'one', the subscript on 'window' is a device for indicating a particular sense of that noun, and the subscript ' $\delta$ ' on 'one' is a device for indicating a particular sense of the antecedent for that pronoun. The idea is to leave room for the possibility of using (38)—the alleged form of (3)—so that the sense of 'window' indicated with ' $\Delta$ ' differs from the sense of 'window' indicated with ' $\delta$ '. And if only for these purposes, let's take senses to be concepts that have extensions.

Let's say that if ' $N_{\Delta}$ ' is a noun, and ' $P_{\delta}$ ' is a pronoun whose antecedent is ' $N_{\Delta}$ ', then relative to any assignment **A** of senses to subscripts: the extension of ' $N_{\Delta}$ ' is the extension of the sense of ' $N$ ' that **A** assigns to ' $\Delta$ '; the extension of ' $P_{\delta}$ ' is the extension of the sense of ' $N$ ' that **A** assigns to ' $\delta$ ', allowing for the possibility that  $\mathbf{A}(\delta) \neq \mathbf{A}(\Delta)$ ; and **A** is *germane to a*

conversational situation  $s$  in which ‘ $N_\Delta$ ’ and ‘ $P_\delta$ ’ are used only if  $A(‘\Delta’)$  is the sense of ‘ $N$ ’ expressed with ‘ $N_\Delta$ ’ in  $s$ , and  $A(‘\delta’)$  is the sense of ‘ $N$ ’ expressed with ‘ $P_\delta$ ’ in  $s$ .<sup>11</sup>

In short, we can define an assignment-relative relation that an expression-situation pair can bear to a set. But this doesn’t show that ordinary expressions map situations to extensions. (As discussed in the next section, polysemous expressions may be more like recipes for cocktails than mappings from bottles to drinks.) Still, there can be a situation  $s1$  such that *relative to any assignment of senses to subscripts* that is germane to  $s1$ , the extension of ‘ $window_\Delta$ ’ is a set of holes, while the extension of ‘ $one_\delta$ ’ is a set of fillers. And there can be a situation  $s2$  such that relative to any assignment germane to  $s2$ , the extension of ‘ $window_\Delta$ ’ is a set of fillers, while the extension of ‘ $one_\delta$ ’ is a set of holes. Let’s grant that there can be many such situations.

If a speaker utters (1), she presumably does so in a situation that is like  $s1$  rather than  $s2$ .

(1) The window they had cut in the wall was nicer than the one they installed.

By contrast, a speaker who utters (39) presumably does so in a situation like  $s2$  rather than  $s1$ .

(39) The window they installed was nicer than the one they had cut in the wall.

So one might hope to preserve the idea that expressions have extensions relative to situations. But the situation initially envisioned—call it  $s0$ —is unlike  $s1$  and unlike  $s2$ . It’s not that *either* (i) relative to  $s0$ , ‘ $window$ ’ is used to express the hole-sense, *or* (ii) relative to  $s0$ , ‘ $window$ ’ is used to express the filler-sense. It’s rather that ‘ $window$ ’ is being used without being used to express either sense. Click’s *preference* in  $s0$ , for Holy or Philly, does not determine a *sense* for ‘ $window$ ’ in (3). No assignment of senses is germane to  $s0$ .<sup>12</sup>

(3) The window that Click preferred was nicer than the one that Clack preferred.

Correlatively, I claim, (3) can be used in  $s0$  without being used to express a proposition. Sentence (3) can be pronounced and endorsed without using it to assemble and endorse a thought that (3) can be used to assemble. To use (3) in this way, and thereby exclude the possibilities we exclude in  $s0$ , we have to use (3) *without* using ‘ $window$ ’ or ‘ $one$ ’ to express either sense. But we can indeed use (3) in this way. So it seems that in  $s0$ , (3) does not express—nor is it used to express—any proposition, even though it is used to exclude some possibilities.

There are many contexts in which (3), which has both a meaning and a pronunciation, can be used to assemble and express a proposition. But in some contexts, (3) can be used in ways that rely on not using (3) to express a thought, much less any thought that it can be used to build.<sup>13</sup> To be sure, resolving polysemy is sometimes *like* elaborating or extending an expression (e.g., replacing ‘ $window$ ’ with ‘ $windowpane$ ’, or ‘ $door$ ’ with ‘ $doorway$ ’). But (3) differs from any disjunction of elaborations or resolutions of it. Similarly, (40) differs from the disjunction of (40) and (41), even if Click can get home only by going left or going straight.

(40) Click got home.

(41) Click got home by going left.

(42) Click got home by going straight.

---

<sup>11</sup> My quote marks are polysemous. Some are like Quinean quote marks; others are like Quinean corner quotes.

<sup>12</sup> So diagnosing polysemy as homophony would not help *even if* the antecedent of ‘ $one$ ’ in (3) could be a word that does not appear in (3); cp. note 8. The second word in (3) can be used without using it to express one of its senses.

<sup>13</sup> Some context-sentence pairs make it vivid that ordinary sentences need not express propositions. Consider “contingent liar sentences;” see Parsons (1974), Kripke (1975). Let ‘ $Box$ ’ be a name for a certain box, and let ‘ $M$ ’ be a name for the sentence ‘Most of the sentences in  $Box$  are not true’. Suppose that four of the five sentences in  $Box$  are ‘ $1 = 1$ ’, ‘ $1 = 2$ ’, ‘ $2 = 2$ ’, and ‘ $2 = 3$ ’. In situations where the fifth sentence is ‘ $3 = 3$ ’,  $M$  seems false; in situations where the fifth sentence is ‘ $3 = 4$ ’,  $M$  seems true. So it can seem that  $M$  expresses a proposition that is true in some situations and false in others, until one considers the possibility that  $M$  is the fifth sentence in  $Box$ .

#### 4. Possibilities and Pragmatics

I would like to *explain* how endorsing recipes can convey information given shared assumptions. I'll settle for detaching familiar descriptions of assertion, in terms of excluding possibilities that had been left open, from the premise that declarative sentences express propositions; cp. Stalnaker (1978, 1999) and others for whom the premise is a dispensable idealization.

Imagine sitting with a friend at a bar that is tended by a robot that prepares cocktails by executing recipes. If there is a choice of ingredients to be made (e.g., selecting a brand of gin), the choice is made at random. You don't know what ingredients are available, but your friend does. And you two have some shared beliefs: if the bar has gin, it has G1 or G2 or both; if the bar has vermouth, it has V1 or V2 or both; a Martini requires gin and vermouth;<sup>14</sup> but combining G1 with V1 yields a bad Martini; combining G2 with V2 also yields a bad Martini; and your friend wants to make sure you don't get a bad drink. So if she says 'You can have a Martini', you can infer that the bar has some gin and some vermouth; and *if* it has G1, it has V2 but not V1; and *if* the bar has G2, it has V1 but not V2. Similarly, you both take it as given that a Manhattan requires bourbon (not rye) and vermouth; if the bar has bourbon, it has B1 or B2 or both; and combining B1 or B2 with V2 yields a bad Manhattan. So if your friend says, 'You can have a Manhattan', you can infer that the bar has some bourbon and V1 but not V2.

In short, if has been agreed that your friend will utter the name of a cocktail if and only if the recipe can't be executed in a way that yields a bad drink—given the ingredients at hand—then given the relevant common knowledge, your friend can use 'Martini' and/or 'Manhattan' to convey information about the available ingredients, as summarized below.

**Martini:**

- G1 v G2
- G1  $\supset$  V2 &  $\sim$ V1
- G2  $\supset$  V1 &  $\sim$ V2

**Manhattan:**

- B1 v B2
- V1 &  $\sim$ V2

It's also "common knowledge" that a Negroni requires Campari, gin, and vermouth; but Campari, G1, and V2 yields a bad Negroni, as does Campari, G2, and V1; a Boulevardier requires Campari, bourbon, and vermouth; but Campari, B2, and V1 yields a bad Boulevardier, as does Campari, B2, and V2. So your friend can use 'Negroni' and/or 'Boulevardier' to convey information about the available ingredients, as summarized below.<sup>15</sup>

**Negroni:**

- C & (G1 v G2)
- G1  $\supset$  V1 &  $\sim$ V2
- G2  $\supset$  V2 &  $\sim$ V1

**Boulevardier:**

- C & B1
- V1 v V2

With this toy example in mind, let's return to (3) and the initial situation **s0**.

- (3) The window that Click preferred was nicer than the one that Clack preferred.

---

<sup>14</sup> A vodka martini is not a martini; cp. 'fake diamond'. But neither you nor your friend insist that the vermouth be *white*. You both like an Original Martini (a.k.a. Martinez, a.k.a. Gin and It) if the right ingredients are available.

<sup>15</sup> The summaries can also be viewed as directions for how to make decent drinks. Disjoining options licenses further inferences. For example, 'You can have a Negroni or a Manhattan' would imply that the bar has C, G1, V1, and either B1 or B2. If it is safe to assume that a list of options is exhaustive, this licenses pragmatic inferences from 'not-mentioned' to 'should not have'. In which case, 'You can have a Negroni or a Manhattan' implicates that you shouldn't have a Boulevardier, and hence that the bar has B2. And if your friend says 'You can have a Manhattan or a Boulevardier'—which would imply that the bar has C, V1, B1, and either G1 or G2—taking the list of options to be exhaustive would suggest both that you shouldn't have a Martini (hence  $\sim$ G1, given V1), and that you shouldn't have a Negroni (hence  $\sim$ G2, given V1).

If ‘window’ is not used to express any sense of that word, and so (3) is not used to express any proposition, one wants to know how endorsing (3) in **s0** can be a way of making a substantive claim that excludes some possibilities while leaving open the ones left open by (28).

(28) Identical(Holy, Clickbait) & Nicer(Holy, Philly) v  
 Identical(Philly, Clickbait) & Nicer(Philly, Holy)

But once we drop the idea that (3) *represents* the possibilities we leave open by using (3), it is easier to explain how we can use (3) in **s0** to exclude the possibilities specified with (30).

(30) [Identical(Holy, Clickbait) & Nicer(Philly, Holy)] v  
 [Identical(Philly, Clickbait) & Nicer(Holy, Philly)]

The relevant possibilities are specified by the disjunction of (28) and (30). And there are two relevant ways of using (3) to build a thought: one way yields a thought that has  $C_{\text{hole}}$  and  $K_{\text{filler}}$  as constituents; the other way yields a thought with  $C_{\text{filler}}$  and  $K_{\text{hole}}$  as constituents. If the first thought is correct, then the first disjunct of (28) is correct, and the first disjunct of (30) is wrong. If the second thought is correct, then the second disjunct of (28) is correct, and the second disjunct of (30) is wrong. So (3) can be used to exclude the possibilities specified with (30), leaving the possibilities specified with (28). But there is no thought that (3) expresses; and in **s0**, there is not even a thought that (3) gets used to express. By using (3), we can indicate our view that Click was right whatever he preferred: if he preferred the hole, the hole was nicer than the filler; and if he preferred the filler, the filler was nicer than the hole. But (3) doesn’t express the thought that Click was right whatever he preferred. The meaning of (3) isn’t *that* flexible.

There are, however, no possibilities that (3) itself excludes. Expressions of a human I-language have meanings that don’t determine extensions. So the sentential expressions don’t exclude possibilities, except indirectly via thoughts expressed by using the meaningful expressions.<sup>16</sup> That said, (3) can be used to express thoughts. Or put another way, while no assignment of senses to subscripts in (38) is germane to **s0**, there are other possible situations.

(38) The  $window_{\Delta}$  that Click preferred was nicer than  
 the  $one_{\delta}$  that Clack preferred

Recall that (3) need not be used co-predicationally. It can be used to say that Click preferred one hole and Clack preferred another, or to say that Click preferred one filler and Clack preferred another. Let the **A:hh**-situations be those in which both  $window_{\Delta}$  and  $one_{\delta}$  are used to express the hole-sense, making the corresponding assignments germane. Let the **A:ff**-situations be those in which both  $window_{\Delta}$  and  $one_{\delta}$  are used to express the filler-sense, making other assignments germane. There can also be situations in which co-predication is intended in a specific way. So let the **A:hf**-situations be those in which  $window_{\Delta}$  and  $one_{\delta}$  are used to express the hole-sense and the filler-sense, respectively, making some co-predicational assignments germane. Let the **A:fh**-situations be those in which  $window_{\Delta}$  and  $one_{\delta}$  are used to express the filler-sense and hole-sense, respectively, making other co-predicational assignments germane.

For just a little while longer, let’s keep restricting attention to situations that also meet the following general conditions: Holy and Philly exhaust the candidates that Click and Clack are considering; the hole-sense and filler-sense exhaust the candidate senses for ‘window’; and any other context-sensitive aspects of (3), including any other polysemies, are resolved as in **s0**.

(3) The window that Click preferred was nicer than the one that Clack preferred.

---

<sup>16</sup> Maybe thoughts determine sets of possibilities, modulo vagueness and Liar Sentences (see note 13), and theorists can safely ignore polysemy when trying to describe thoughts that can be assembled in accord with natural sentence meanings. If so, then perhaps many versions “possible worlds semantics” should be viewed as proposals about thought contents, rather than linguistic meanings. Like most words, ‘semantics’ is polysemous; see Burgess (2008).

Then we can use the diagram below to represent the various situations in which (3) can be used.

	situations where Click preferred Holy, and Clack preferred Philly	situations where Click preferred Philly, and Clack preferred Holy
<b>s0</b>	No proposition is expressed with (3).	No proposition is expressed with (3).
<b>A:hh situations</b>	Any proposition expressed with (3) isn't true. Truth requires that Clack preferred a hole.	Any proposition expressed with (3) isn't true. Truth requires that Click preferred a hole.
<b>A:ff situations</b>	Any proposition expressed with (3) isn't true. Truth requires that Click preferred a filler.	Any proposition expressed with (3) isn't true. Truth requires that Clack preferred a filler.
<b>A:hf situations</b>	The proposition expressed with (3) is true iff the hole was nicer than the filler.	Any proposition expressed with (3) isn't true. Truth requires that Click preferred a hole.
<b>A:fh situations</b>	Any proposition expressed with (3) isn't true. Truth requires that Click preferred a filler.	The proposition expressed with (3) is true iff the filler was nicer than the hole.

Don't worry about whether any proposition gets expressed in "category mismatch" scenarios, like those where Click preferred Holy but 'window' was used to express the filler sense. Let's be concessive, focus on the co-predicational alternatives to **s0** that are not mismatch scenarios, and grant the following: relative to **A:hf**-situations in which Click preferred Holy and Clack preferred Philly, (3) is used to express a proposition that is true if and only if the window Click preferred (i.e., the hole) was nicer than the window that Clack preferred (i.e., the filler); relative to **A:fh**-situations in which Click preferred Philly and Clack preferred Holy, (3) is used to express a proposition that is true if and only if the window Click preferred (i.e., the filler) was nicer than the window that Clack preferred (i.e., the hole); and each of these propositions determines a corresponding set of possible worlds at which the proposition is true.

The first proposition, expressed in **A:hf**-situations, excludes worlds at which the filler was nicer than the hole. The second proposition, expressed in **A:fh**-situations, excludes worlds at which the hole was nicer than the filler. But we can also consider a "diagonal" set that includes the possible worlds at which the proposition expressed with (3) in **A:hf**-situations is true *and* the possible worlds at which the proposition expressed with (3) in **A:fh**-situations is true; cp. Stalnaker (1999). Modulo concerns about definite descriptions and modal profiles, the intended diagonal set might also be described as the set of worlds at which (28) is true.

(28) Identical(Holy, Clickbait) & Nicer(Holy, Philly)  $\vee$   
 Identical(Philly, Clickbait) & Nicer(Philly, Holy)

Talking about this set may be useful in describing how *our belief that Clack was right* can exclude the possibilities specified with (30), leaving the ones specified with (28).

(30) [Identical(Holy, Clickbait) & Nicer(Philly, Holy)]  $\vee$   
 [Identical(Philly, Clickbait) & Nicer(Holy, Philly)]

But we don't need to say that *sentence (3)* determines any such set in **s0**. To repeat, there is no set of worlds that (3) determines, either by itself or relative to a situation where 'window' expresses no specific sense. One can, however, use the apparatus of diagonalization to describe the possibilities left open by speakers who use (3) co-predicationally in **s0**.

More generally, the usual tools for describing *communication* can remain in place, even if we grant that the meaning of (3) does not map contexts to propositions. And if sentences like (3) do not express propositions in situations like **s0**, then tools for describing communication *should* be detached from the assumption that ordinary sentences map contexts to propositions.<sup>17</sup> But once we admit the phenomenon of “regular” polysemy—illustrated with the pane/filler equivocality of ‘window’ and vehicle/content equivocality of ‘book’—we can and should grant that polysemy is open-ended. We can use ‘window’ to talk about, via concepts that let us think about, windows that are neither panes nor fillers; we can use ‘line’ to talk about waiting lines and lines in faces; etc. So sentences like (43) are related to diagonal sets *very* indirectly.

(43) At the hardware store, the line to pick up windows was longer than the fishing line on display in the front windows.

For each speaker at each time, each lexical item will correspond to a shortish disjunction of concepts. So any one example of polysemy can be *modeled* as a case of homophony without too much distortion. But this doesn’t show that polysemy can be modeled as homophony without massive distortion, much less that sentences like (43) and (3) determine diagonal sets.

(3) The window that Click preferred was nicer than the one that Clack preferred. Examples like (3), used in contexts where the hole/filler contrast is all that matters, reflect the tip of an iceberg. Human linguistic communication almost always involves the use of expressions that are polysemous along multiple dimensions.<sup>18</sup>

Moreover, as Carston (2002, 2019, 2021) and others have stressed, it’s not just that regular polysemy shades off into extended uses that are still common enough to be noted in dictionaries. Words can be used in *ad hoc* ways that are not conventional in any interesting sense; see Wilson and Carston (2007). For example, a speaker might use ‘democracy’ to express a concept acquired only moments before, in a conversation with a radical political theorist or an annoying uncle; and the speaker might jettison the concept once the conversation ends.

Carston offers a plausible account, rooted in her cognitivist conception of pragmatics, of how lexical items come to be so flexible. The idea is that construing (utterances of) words often involves construction of new concepts, given how concepts already encoded interact with “contextual information and expectations of relevance.” My claims about what linguistic meanings are, and how they allow for conceptual equivocality, are logically independent of Carston’s specific suggestion about the etiology of polysemy; and her claims about polysemy do not presuppose a particular account of linguistic meaning. But let me end this section by noting that the proposals can certainly be combined, modulo some caveats that may be terminological.

Carston sometimes speaks of modulating lexical meanings. But while I agree that speakers often modulate the conceptual contents associated with a lexical item, I take the meaning of a human linguistic expression  $\Sigma$  to be an essential feature of  $\Sigma$ . If the meaning of ‘window in a corner’ is an instruction for how to generate a concept—to a first approximation, fetch a concept from the ‘window’-address and conjoin it with a concept built via the instruction specified with ‘in a corner’—then any expression that connects the pronunciation of ‘window in

---

<sup>17</sup> One can stipulate that meanings map contexts to sets of possible worlds, and that situations like **s0** correspond to contexts that meanings map to diagonal sets, regardless of how actual conversational situations and sentence meanings are related to contexts and meanings. But stipulations can’t show that ordinary sentence meanings determine mappings from conversational situations to sets of worlds; see Pietroski (2006) and section one above.

<sup>18</sup> If only for this reason, co-predicational construals of grammatical antecedence can vary with regard to cognitive felicity; see Murphy (2021) for illuminating experiments and discussion. For example, given (43), replacing the second occurrence of ‘windows’ with ‘ones’ would be bizarrely zeugmatic at best. But one can diagnose felicitous cases of co-predication in terms of conceptual equivocality without predicting that co-predication is unconstrained.

a corner’ with a *different semantic instruction* is a *different expression*.<sup>19</sup> And if the meaning of ‘window’ is an instruction for a how to generate a concept that can be systematically combined with others—to a first approximation, simply fetch a concept from a certain lexical address—then the modulation can be described as adding concepts to the ‘window’-family and thereby increasing the number of ways in which the lexical instruction can be executed. But the word and its meaning remain the same. Modulating the stock of available ingredients does not modulate a recipe. If you tell me to pick a bottle of gin from a shelf, your instruction doesn’t change if new options appear before I choose. Likewise, your request for a Negroni doesn’t change if it prompts a bartender to go the basement in search for a bottle that wasn’t on the shelf.

The second caveat concerns Carston’s helpful distinction between *linguistic* lexicons and *communicative* lexicons: she describes “L-lexicons” as components of I-languages; C-lexicons are said to be sets of “reasonably stable associations of particular formal elements (syntactically categorized phonetic/gestural objects) and senses/concepts, tacitly agreed on across a community of speakers (i.e. a set of conventions) and used by them as devices of communication (p.157).” Carston does not say that the “reasonably stable associations” form a theoretically interesting class, or even that a C-lexicon is a natural psychological object, as opposed to a combination of an L-lexicon with some but not all of the corresponding lexically accessible concepts; though as she suggests, speakers may assign some special status to senses/concepts that are taken to be conventional. Talk of C-lexicons can usefully highlight questions about what speakers need, in addition to L-lexicons and shared composition principles, to support shared expectations about how pronunciations will be linked to concepts in a given community. But presumably, sets of associations do not have meanings, at least not in the sense that lexical items of an I-language do.

Carston is surely right that there are norms for connecting concepts with pronunciations in ways that reduce the need for *ad hoc* modulations; see also Borg (2019). And if there are social facts in virtue of which (suitably mature) speakers ought to link their lexical items to certain concepts, we can say that speakers exhibit corresponding degrees of Lexico-Conceptual competence in using their I-languages to communicate with others; cp. Higginbotham (1989). But if having a C-lexicon amounts to having an L-lexicon and using it appropriately given local standards, we should take care to distinguish Carston’s proposal from the much less attractive idea that having meaningful lexical items is a matter of participating in community practices.<sup>20</sup>

---

<sup>19</sup> Likewise, any expression that connects the meaning of ‘window in a corner’ with a different pronunciation is a different expression. The expressions generated by human I-languages can be described as pronunciation-meaning pairs that exhibit syntactic structure. Here, we can be neutral about what pronunciations are. But one can think of them as instructions that (from a speaker’s perspective) specify how to *articulate* expressions; see Halle (1990).

<sup>20</sup> See the discussions of Lewis (1975) in Chomsky (1980) and Pietroski (2018). And even if one follows Lewis in downplaying questions about how grammars are acquired, it isn’t obvious that adherence to social norms plays an essential role in communication. Imagine two speakers of English, Dora and Dick, who are slightly deviant: they each have a noun that connects the pronunciation of ‘gasket’ with a lexical address used to access a concept of flanges, and a noun that connects the pronunciation of ‘flange’ with a lexical address used to access a concept of gaskets. I grant that if Dora utters a sentence with the pronunciation of ‘There is a gasket on the shelf’ in response to a question from a stranger, Dora *said that* there is a gasket on the shelf; and here, correction may be appropriate. But if Dora utters the same sentence in talking with Dick, she doesn’t *communicate that* there is a gasket on the shelf; cp. Davidson (1986), Pietroski (1994). More generally, I see no reason to say that Dora has attached a “wrong meaning” to her own word—whatever that would mean—as opposed to saying that she and Dick have I-languages such that (i) our word ‘gasket’ is the best translation of their nouns that have the pronunciation of ‘flange’, and (ii) our word ‘flange’ is the best translation of their nouns that have the pronunciation of ‘gasket.’ Relatedly, there may be no theoretically interesting sense in which distinct lexical items have the *same* meaning—as opposed to being decent mutual translations—even if each I-language expression has a meaning; cp. Quine (1960).



## 5. Extending the Moral

One can hypothesize that human I-languages connect pronunciations with meanings that determine extensions. But why think this hypothesis is true, given sentences like (43)?

(43) I think he used his blue pen to write about blue skies, truth, and democracy.

Prima facie, it's just not plausible that 'think', 'use', 'blue', 'write', 'about', 'sky', 'truth', and 'democracy' have extensions—much less extensions that combine in an independently plausible way to yield a truth condition for (43); see Wittgenstein (1953), Tarski (1944), Strawson (1950), Austin (1962), Chomsky (2000), Travis (2008), and many others. Similarly, it's not plausible that 'window' has an extension. So why think (3) has a truth condition?

(3) The window that Click preferred was nicer than the one that Clack preferred.

And if (3) doesn't have a truth condition, even relative to contexts, neither does (44).

(44) The window that Click liked was nice.

So even if the context-sensitivity of 'I' is not a stumbling block for the idea that words have extensions, the polysemy of 'window' is. A speaker can use and endorse (44) without knowing whether the window that Click liked was a hole, a pane of glass, a display space at the front of a store, or a window of some other sort. Similar remarks apply to endlessly many examples.

It's often said that nouns like 'water' and 'rabbit' are kind terms that have extensions. But I don't know of any good arguments for this claim. Putnam (1975) *assumed* that meanings determine extensions, while conceding that this "severe idealization" abstracts from the fact a typical word has multiple uses or senses. He used intuitions about 'water', H<sub>2</sub>O, and Twin Earth to argue for a disjunctive claim that I endorse: either meanings don't determine extensions, or meanings ain't in the head. But he offered no reason for thinking that 'water' has an extension; see Pietroski (2017a, 2018). The word 'water' may not be like 'window' ('book', 'line', etc.) in being associated with dovetailing concepts and examples of co-predication like (1).

(1) The window they had cut in the wall was nicer than the one they installed.

But prima facie, 'water' can be used to access distinct concepts on distinct occasions of use.

I grant that 'water' has a kind-sense that gets pumped to the surface by Putnam's discussion. Used this way, 'water' invites speakers of English to access a concept that applies to samples of H<sub>2</sub>O (modulo impurities) *and not* to samples of chemically distinct stuff. But why think 'water' has no other sense, perhaps corresponding to some concept that applies to samples of one or more (natural?) suitably thirst-quenching liquids, allowing for multiple realizability? Given that at least many human I-language expressions are polysemous, it would be surprising if 'water' was conceptually unequivocal, especially in a way geared for chemistry. And in any case, Diet Coke—or, as Chomsky (1995) noted, a typical cup of tea—has a higher percentage of H<sub>2</sub>O than many actual samples of water from city taps or rural wells. So even if we ignore all the ocean water and restrict attention to ordinary sources of fresh water, there is a lot of stuff that counts as water, despite having a significantly lower percentage of H<sub>2</sub>O than a lot of stuff that isn't water. This suggests that human speakers of English can use 'water' to access a concept that is not a concept of H<sub>2</sub>O (modulo impurities). But in retrospect, this conclusion should seem unsurprising unless one was trained to deny it.

Rabbits are also diverse in ways that challenge the idea that 'rabbit' has an extension. If a child acquires 'rabbit' upon encountering a few members of the seventeen species of Cottontails, but none of the European rabbits that are dominant in many places, why think the child acquired a word that has an extension? Rabbits are not a hodge-podge, but neither do the various rabbits form a kind that magnetically attracts the count noun 'rabbit'. And even if 'rabbit' has an expert sense that determines an extension, why think the noun has only this sense? Consider badgers,

who are even more genealogically diverse. As experts put it, the badgers are a polyphyletic group that is not even roughly a clade. The fifteen species of badger are spread across the family *Mustelidae*, with European badgers and ferret-badgers being more closely related to wolverines and martins than to the older line of American badgers that can be traced more directly to the first members of the family. So why think that in acquiring ‘badger’, a child acquires a word that has an extension, as opposed to an I-language expression that can be used to access a concept.

I happily grant that cognitive scientists have fruitful notions of kind-concepts (Keil 1989, Gelman 2003), and that in acquiring English, a typical child connects a pronunciation for ‘rabbit’ with some such kind-concept of rabbits. I think this is an important part of why the acquired lexical items are inter-translatable, and why they count as instances of a type that can be described as the English word ‘rabbit’; cp. Carston’s talk of C-lexical items. But in thinking about what the phrase ‘concept of rabbits’ means, we should not assume that ‘rabbit’ has an extension, or that children who form kind-concepts in response to seeing some rabbits form kind concepts that have a shared extension. If ‘concept of rabbits’ is an instruction for how to build a concept, then once a speaker knows about the diversity of rabbits, she can use ‘rabbit’ to access concepts of different species in different contexts—or use ‘rabbit’ to express a more permissive concept that is indifferent to the diversity, perhaps relying on taxonomists to have settled on a stable way of talking that can be adopted if necessary. Likewise for ‘badger’.

Moreover, even if there are some words that plausibly have extensions, this wouldn’t make it plausible that human I-language expressions have meanings that determine extensions.<sup>21</sup> There is no empirical pressure to say that ‘window’ is univocal, with a single sense that applies equally well to relevant holes and fillers. We don’t need to assume that speakers adhere to a “one word, one concept” policy, despite the diversity of word use, by continually replacing concepts with more permissive variants. We can just recognize that the “one word, one concept” policy is for ideal languages designed to avoid conceptual equivocality, and not for children who need to treat homophony as a last resort when trying to figure out how others connect pronunciations with meanings.

---

<sup>21</sup> Even if a word is introduced legalistically, it’s hard to keep the word from becoming polysemous. Consider ‘prorogue’. In the summer of 2019, the British prime minister (Boris Johnson) asked Queen Elizabeth to prorogue (i.e., discontinue the meetings of) Parliament. The Queen issued the relevant order on August 28<sup>th</sup>. This was announced in newspapers with headlines like ‘Parliament Prorogued’; and on September 9<sup>th</sup>, business in Parliament ceased. But on September 24<sup>th</sup>, the British Supreme Court found that despite the monarch’s order, Parliament had not been prorogued. As an official website explained,

...Prime Minister Boris Johnson advised the Queen to prorogue or suspend – Parliament – from the end of 9 September until 14 October. The Supreme Court has now ruled that this advice, (**and the prorogation that followed**), was unlawful because it had the “effect of frustrating or preventing, without reasonable justification, the power of Parliament to carry out its constitutional functions.” This means that **Parliament was not prorogued** on 9 September, and the **previous** parliamentary session continues....The Speaker of the House of Commons has said that parliamentary business will **resume** on 25 September at 11:30am. (<https://www.instituteforgovernment.org.uk/explainers/proroguing-parliament>, my bold font).

At least after the Court’s decision, it seems that ‘prorogue’ could be used to express (i) a concept that applies to certain orders by monarchs, including the one that Queen Elizabeth issued in 2019, and (ii) a related concept of processes in which a monarch suspends Parliament by issuing such an order—cp. Thomson (1977) on the time of a killing—but also (iii) a related normative concept of processes in which a monarch *lawfully* suspends Parliament by issuing such an order, and (iv) a related concept that applies to certain lawful orders, not including the one that Queen Elizabeth issued in 2019. Drawing on Austin (1962), one might say that the court found that a certain felicity condition on proroguing had previously gone unnoticed; though especially when used to describe court decisions regarding hard cases, the polysemous word ‘found’ differs from ‘discovered’.

It can be tempting to think that given enough ontology and appeal to contexts, one can preserve the idea that meanings determine extensions, despite examples like (1), (2), and (43).

(1) The window they had cut in the wall was nicer than the one they installed.

(2) The book she brought weighed two pounds, but it had been well reviewed.

(43) I think he used his blue pen to write about blue skies, truth, and democracy.

But the issues addressed here remain, even if the world includes holes, skies, truth, democracy, and metaphysically hybrid entities like books that are both abstract and concrete, windows that are filled holes, and so on; cp. Pustejovsky (1995), Asher and Pustejovsky (2006), Asher (2011), Ludlow (2014).<sup>22</sup> Whatever entities there are, we still face the question of whether word meanings map contexts to sets of them. And upon reflection, examples like (3) suggest that polysemous word meanings do not map contexts to senses.

(3) The window that Click preferred was nicer than the one that Clack preferred.

Indeed, we can kick away the ladder of co-predication and illustrate the point with (44).

(44) The window that Click liked was nice.

The meaning of ‘window that Clicked liked’ does not map contexts to sets of windows. The phrasal meaning, which can be used without using any mapping from contexts to sets, is more like a recipe for how to build a concept that includes an ingredient accessed via ‘window’.<sup>23</sup>

## **References**

- Asher, N., & Pustejovsky, J. (2006). A type composition logic for generative lexicon. *Journal of Cognitive Science*, 6, 1–38.
- Asher, N. (2011). *Lexical Meaning in Context*. Cambridge, U.K.: Cambridge University Press.
- Austin, J. (1962). *How to Do Things with Words*. Oxford: Oxford University Press.
- Baker, M. (2002). *The Atoms of Language*. New York: Basic Books.
- Berwick, R., Pietroski, P., Yankama, B., and Chomsky, N. (2011). Poverty of the Stimulus Revisited. *Cognitive Science* 35: 1207-42.
- Borg, E. (2019). Explanatory Roles for Minimal Content. *Noûs* 53: 513-539.
- Bromberger, S. (2011). What are words? *Journal of Philosophy*, 108:486-503.
- Burgess, J. (2008). Tarski’s Tort. In *Mathematics, Models, and Modality*. Cambridge: Cambridge University Press.
- Caplan, S., Kodner, J. and Yang, C. (2020). Miller’s monkey updated: Communicative efficiency and the statistics of words in natural language. *Cognition* 204.104466.
- Carston, R. (2002). *Thoughts and Utterances*. Oxford: Blackwell.
- (2019). Ad Hoc Concepts, Polysemy and the Lexicon. In K. Scott, B. Clark, R. Carston (eds.), *Relevance, Pragmatics and Interpretation* (pp. 150-162). Cambridge: CUP.
- (2021). Polysemy: Pragmatics and sense conventions. *Mind and Language* 36:108-133.
- Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.
- (1959). On Certain Formal Properties of Grammars. *Information and Control* 2:137-167.
- (1964). *Current Issues in Linguistic Theory*. The Hague: Mouton.
- (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- (1980). *Rules and representations*. New York: Columbia University Press.

---

<sup>22</sup> Though like Ludlow, Davidson (1986), Carston (2002, 2019), Wilson and Carston (2007), I think it is important to recognize the dynamic character of how human languages are used to express concepts.

<sup>23</sup> This paper was inspired by a series of conversations with Andy Egan and Norbert Hornstein, along with fond recollections of Car Talk, before and after giving a talk about polysemy as part of a course at the University of Arizona. Special thanks to Andy, Norbert, John Collins, and Zoltan Szabo for valuable comments on previous drafts.

- (1986). *Knowledge of Language*. New York: Praeger.
- (1995). Language and Nature. *Mind* 104:1-61.
- (2000). *New Horizons in the Study of Language and Mind*.  
Cambridge: Cambridge University Press.
- Church, A. (1941). *The Calculi of Lambda Conversion*. Princeton: Princeton University Press.
- Collins, J. (2017). The Copredication Argument. *Inquiry* 7:675–702.
- Davidson, D. (1986). A Nice Derangement of Epitaphs.  
In E. Lepore (ed.) *Truth and Interpretation*. Oxford: Blackwell.
- Devitt, M. (2020). *Overlooking conventions: The trouble with linguistic pragmatism*.  
New York, NY: Springer.
- (2021). Semantic polysemy and psycholinguistics. *Mind & Language*, 36(1), 134– 157.
- Dummett, M. (1986). A Nice Derangement of Epitaphs: Some Comments on Davidson and Hacking. In E. Lepore (ed.) *Truth and Interpretation*. Oxford: Blackwell.
- Falkum, I. & Vicente, A. (2015). Polysemy: Current perspectives and approaches.  
*Lingua* 157:1-16.
- Fodor, J. (1998). *Concepts: Where Cognitive Science Went Wrong*.  
New York: Oxford University Press.
- Gelman, S. (2003). *The Essential Child*. Oxford: Oxford University Press.
- Grice, P. (1989). *Studies in the Way of Words*. Cambridge: Harvard University Press.
- Halle, M. (1990). Phonology. In *An Invitation to Cognitive Science*, D. Osherson and H. Lasnik, eds. (Cambridge, MA: MIT Press).
- Hawthorne, J. and Lepore, E. (2011). On words. *Journal of Philosophy* 108:447-485.
- Higginbotham J. (1989). Knowledge of Reference. In A. George, ed.,  
*Reflections on Chomsky*, pp. 153-174. Oxford: Blackwell, 1989.
- Kaplan (1990). Words. *Proceedings of the Aristotelian Society, Supplementary* 64:93-119.
- Keil, F. (1989). *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press.
- King, J. (2018). W(h)ither Semantics! (?) *Nous* 52:772-795.
- Kripke, S. (1975). Outline of a theory of truth, *Journal of Philosophy*, 72: 690–716.
- (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lewis, D. (1975). Languages and Language. In K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science*, Volume 7, Minneapolis: University of Minnesota Press, pp. 3–35.
- Liebman, D. and Magidor, O. (2019). Copredication, Counting, and Criteria of Individuation: A Response to Gotham. *Journal of Semantics* 36: 549–561.
- Ludlow, P. (2014). *Living Words: Meaning Underdetermination and the Dynamic Lexicon*.  
New York: Oxford University Press.
- Murphy, E. (2021). *Linguistic Representation and Processing of Copredication*.  
Doctoral dissertation, University College London.
- Parsons, C. (1974). The Liar paradox. *Journal of Philosophical Logic* 3: 381–412.
- Pietroski, P. (1994). A Defense of Derangement. *Canadian Journal of Philosophy* 24: 95-118.
- (2005). Meaning Before Truth. In G. Preyer and G. Peters (eds.)  
*Contextualism in Philosophy*, 253-300 (Oxford: Oxford University Press).
- (2006). Character Before Content. In *Content and Modality: Themes from the Philosophy of Robert Stalnaker* (J. Thomson and A. Byrne, eds.), 34-60, Oxford University Press.
- (2010). Concepts, Meanings, and Truth: First Nature, Second Nature and Hard Work. *Mind and Language* 25: 247-78.
- (2014). Framing Event Variables. *Erkenntnis* 80:31-60.

- (2017a). Semantic Internalism. In J. McGilvray (ed.), *The Cambridge Companion to Chomsky*, 196-216, (Cambridge: Cambridge University Press).
- (2017b). Meanings via Syntactic Structures. In *Syntactic Structures after 60 Years* (N. Hornstein, H. Lasnik, C. Yang, and P. Patel, eds.) De Gruyter: Mouton (2017).
- (2018). *Conjoining Meanings: Semantics without Truth Values*. Oxford University Press.
- (2021). Fostering Liars. *Topoi* 40:5-25.
- Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Putnam, H. (1975). The Meaning of ‘Meaning’. In K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis: University of Minnesota Press).
- Quilty-Dunn, J. (2021). Polysemy and thought: Toward a generative theory of concepts. *Mind and Language* 36: 158-185.
- Quine, W. (1960). *Word and Object*. Cambridge, MA: MIT Press.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Ross, J. (1967). Constraints on variables in syntax. Doctoral dissertation, Massachusetts Institute of Technology. Published as, Ross, J. R. (1986). *Infinite syntax!* Norwood, NJ: Ablex.
- Srinivasan, M. and Berner, C. (2019). Children Use Polysemy to Structure New Word Meanings. *Journal of Experimental Psychology: General* 148: 926–942.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics* 9: 315–32.
- (1999). *Context and Content*. Oxford: Oxford University Press.
- Strawson, P. (1950). On referring. *Mind* 59:320–344.
- Tarski, A. (1944). The Semantic Conception of Truth. *Philosophy and Phenomenological Research* 4: 341-75.
- Thomson, J. (1977). *Acts and Other Events*. Ithaca: Cornell University Press.
- Travis, C. (2008). *Occasion Sensitivity*. Oxford: Oxford University Press.
- Vicente, A. & Falkum, I. (2017). Polysemy. *Oxford Research Encyclopedias*.  
<https://doi.org/10.1093/acrefore/9780199384655.013.325>.
- Wilson, D. and Carston, R. (2007). A unitary approach to lexical pragmatics: Relevance, inference and *ad hoc* concepts. In N. Burton-Roberts (ed.), *Pragmatics*, 230–259. (Basingstoke: Palgrave Macmillan).
- Wittgenstein, L. (1953). *Philosophical Investigations* (New York: Macmillan).