

## Précis of *Conjoining Meanings: Semantics Without Truth Values* (Sept. 26, 2018)

Children acquire languages that connect meanings with pronunciations in striking ways. I offer a proposal about what these meanings are, how they are related to human cognition, and how they are not related to the things we talk about by using words in contexts. In slogan form, meanings are composable recipes for how to build concepts of a special sort. The meaning of ‘green bottle’ is a tripartite instruction: access a concept via ‘bottle’, access a concept via ‘green’, and conjoin the results. Likewise, ‘my green bottle’ calls for conjoining the results of executing the meanings of ‘my’ and ‘green bottle’. Full sentences can be used to build complete thoughts, which may be true or false. But in my view, ordinary sentences don’t have truth values, not even relative to contexts. Meanings don’t determine extensions. My route to these conclusions is paved with details that surely need revision. Though a more pressing challenge is to identify the topic.

### 1. Slangs

Like most words, ‘meaning’ is polysemous. There are many kinds of significance, and many corresponding concepts. So it’s easy, even for specialists, to talk past each other when using ‘meaning’. But I think there is something like a natural kind in the vicinity.

The child-acquirable languages that connect meanings with pronunciations, spoken or signed, are distinctively human. Let’s call them Slangs. The meanings of Slang expressions are somehow compositional in a way that allows for ambiguities yet mirrors certain aspects of logical structure. For example, speakers of English can understand (1) in several ways.

(1) we watched her duck near a muddy bank

But however ‘duck’ and ‘bank’ are construed, deleting ‘muddy’ or ‘near a muddy bank’ is valid.

Such facts indicate the meanings I have in mind: Slangs connect them with pronunciations in interesting ways. One of my central claims in *Conjoining Meanings (CM)* is that ambiguity and composition, along with the compelling character of certain inferences, are natural phenomena that reflect aspects of human psychology. Correlatively, I think we should avoid stipulations about how meanings are related to truth values, possible worlds, or extensions of ideal concepts.

Chapter two offers a Fodorian account of concepts as composable mental representations with which we can think about things. This leads into a discussion of Aristotelian logic, as part of a larger argument that mental *predicates*—concepts that let us categorize—play a special role in a natural logic that vindicates an old idea: predicate reduction (e.g., replacing ‘muddy bank’ with ‘bank’) is typically *conjunct* reduction; and deleting conjuncts is valid except in specially marked environments like negation. For predicates, the default is that longer is stronger. In later chapters, I argue that phrasal meanings are instructions for how to build mental predicates, given lexical meanings that let us access a limited range of atomic concepts. But this is a hypothesis about certain natural languages, not a proposed analysis of any pretheoretic concept of meaning.

For simplicity, let’s focus on spoken Slangs. And let’s say that expressions with the same pronunciation are homophonous, suggesting that they “sound” the same, even if uttered by a soprano and a baritone. But let’s not pretend that spoken English is a single Slang.

Pronunciations vary across speakers who grew up in Brooklyn, Glasgow, Sydney, or different neighborhoods in London. In some cases, the same meaning is expressed with very different sounds, as with ‘biscuits in the lift’ vs. ‘cookies in the elevator’. Conversely, the same sound—e.g., that of ‘solicitor’ or ‘robin’—can be paired with different meanings in different places. (British solicitors refer trial work to barristers. American solicitors represent units of government and often argue cases. Similarly, British ‘robin’ and its American counterpart correspond to different species.) English Slangs also exhibit minor syntactic differences. And as games of Scrabble can reveal, many entries in the O.E.D. are not words of *my* Slang.

To be a speaker of English is to have acquired one of the many Slangs in a broad family, which can be roughly characterized in terms of paradigm cases and a vague intransitive notion of mutual intelligibility. This allows for graded notions of fluency/competence that let us distinguish young children, or an adult with a patchy vocabulary in a second language, from mature native speakers of a Slang family. But there is no communal Slang, Ideal English, that each speaker of English acquires yet never fully masters.

## 2. Kinds of Equivocality

Of course, dictionaries are useful. They help families reduce misunderstandings. They also illustrate the difference between homophony and polysemy—cases of distinct words sounding the same, as opposed to one word having “subsenses.” Polysemy may be the most interesting feature of lexical meanings, even if diagnosing examples can be hard. So let me say a little about the contrast with homophony, before turning to questions about phrasal meanings.

Words with distinct meanings can share a pronunciation. Consider ‘bank’, ‘duck’, and ‘bear malice towards a bear with bare arms’. There are several English words, spelled ‘bear’ or ‘bare’, that connect their meanings with the pronunciation /bɛr/. But we don’t expect *other* Slangs to connect these various meanings with a common pronunciation. Likewise, we don’t expect translations of French homophones to be homophonous in English. (Consider ‘seau’, ‘sceau’, and ‘saut’, whose translations are ‘bucket’, ‘stamp’, and ‘jump’.) Lexical meaning-pronunciation pairings are *arbitrary*. But ‘bear a tray of food’, ‘bear the weight of the roof’ and ‘bear the pain’ seem like examples of a single verb being used to talk about carrying or supporting or enduring. It’s not that we have three accidentally homophonous synonyms of ‘carry’, ‘support’, and ‘endure’. We gather subtly different senses of carrying, or carrying on with, under one verb.<sup>1</sup> But there is more than one way for a word to have multiple senses or uses.

The noun ‘book’ can be used to talk about spatiotemporally located things that carry inscriptions of certain contents, or abstract contents that get encoded in many ways and places. It seems that ‘book’ lets us access dovetailing concepts that can be used to think about intimately related things that differ in ontological kind; cp. ‘triangle’, which can be used to talk about perceptible inscriptions or imperceptible abstracta. We can use ‘window’ to describe an opening in a wall, a pane of glass that occupies such an opening, a display space behind such a pane at the front of a store, an opening in an envelope that makes an address visible, or a gap in a long counter at a bank; cp. ‘line’, ‘run’, and ‘set’. Even given a generous conception of homophony, a typical word exhibits a kind of equivocality, as if a word can point to a family of concepts.

If we take this idea seriously, we can view familiar examples of polysemy as special cases of a broader phenomenon: Slang lexical items are, almost always, *conceptually equivocal*. Consider the singular noun ‘fish’. Following many syntacticians, I think this word combines a count morpheme with a simpler *lexical root*—often called a mass noun—that can be used to talk about the stuff in a can of tuna. The root  $\sqrt{\text{fish}}$  is part of the singular [ $\sqrt{\text{fish}}+\text{CT}$ ], whose plural form is [[ $\sqrt{\text{fish}}+\text{CT}$ ]+PL]. I also think that *concepts* exhibit a mass/count contrast that doesn’t map cleanly onto morphological complexity. The net result, I argue, is that  $\sqrt{\text{fish}}$  is equivocal.

Suppose that FISH<sub>ONE</sub> is an atomic count-concept, WATER<sub>SM</sub> is an atomic mass-concept, [FISH<sub>ONE</sub> STUFF<sub>ONE/SM</sub>]<sub>SM</sub> is a complex mass-concept, and [WATER<sub>SM</sub> UNIT<sub>SM/ONE</sub>]<sub>ONE</sub> is a complex count-concept; where subscripts on small capitals indicate conceptual types, and a slash indicates a type-converter. A child might initially link  $\sqrt{\text{water}}$  to WATER<sub>SM</sub>, and only later introduce

---

<sup>1</sup> By contrast, no noun has subsenses corresponding to ursine animals and stock market pessimists; cp. ‘bear a weight/wait’. It can be hard to distinguish polysemy from metaphor. So in *CM*, I don’t insist on any diagnosis for any specific case. But given any plausible way of counting lexical meanings, they allow for subsenses.

[WATER<sub>SM</sub> UNIT<sub>SM/ONE</sub>]<sub>ONE</sub> to accommodate uses of [√water+CT]. The same child might link [√fish+CT] to FISH<sub>ONE</sub>, not worrying about √fish by itself until experience invites a mass concept like [FISH<sub>ONE</sub> STUFF<sub>ONE/SM</sub>]<sub>SM</sub>. But a word for fish might also be acquired in a “stuff-first” way.

A child who is often fed fish might link √fish to FISH<sub>SM</sub>, not yet realizing where fish sticks and tuna sandwiches come from; cp. √tofu and TOFU<sub>SM</sub>. But such a child can still acquire FISH<sub>ONE</sub>, perhaps upon hearing some swimming things described as fish, or after learning the truth about chicken(s). We can all entertain the thought that fish grows like wheat, while tofu comes from tofus that used to hop around. So we can acquire FISH<sub>SM</sub> and TOFU<sub>ONE</sub>, even if we already acquired [FISH<sub>ONE</sub> STUFF<sub>ONE/SM</sub>]<sub>SM</sub> and [TOFU<sub>SM</sub> UNIT<sub>SM/ONE</sub>]<sub>ONE</sub>. If we know the facts, we may limit our *use* of TOFU<sub>ONE</sub> to episodes of inventing stories, imagining nervous vegetarians, or considering logical possibility. But lexical roots are indifferent to the natures of what we talk about. We can use √fish to access FISH<sub>SM</sub> or [FISH<sub>ONE</sub> STUFF<sub>ONE/SM</sub>]<sub>SM</sub>; likewise for √rabbit, √chicken, etc.<sup>2</sup>

I think this point generalizes. However words get acquired, they tend to become conceptually equivocal. But each Slang has only finitely many atomic expressions. So whatever lexical meanings are, one wants to know how they can combine to yield boundlessly many phrasal and sentential meanings. Here too, it is useful to think about ambiguity, following Chomsky.

Words can differ inaudibly because one contains a silent plural morpheme, as with ‘fish’. Likewise, sentences composed of the same words can differ *structurally*. Many English Slangs connect a pronunciation of string (2) with two meanings, indicated below via (2a) and (2b).

- (2) the duck is ready to eat
- (2a) The duck is fit for consumption.
- (2b) The duck is prepared to dine.

But while (3) has a duck-as-eaten meaning, indicated with (3a), (3) can’t be used to express the duck-as-eater meaning indicated with (3b). By contrast, (4) only has a duck-as-eater meaning.

- (3) the duck is easy to eat
- (3a) It is easy for relevant parties to eat the duck.
- (3b) #It is easy for the duck to eat relevant stuff.
- (4) the duck is eager to eat
- (4a) #The duck is eager to be one whom relevant parties eat.
- (4b) The duck is eager to be one who eats relevant stuff.

The pattern remains the same if ‘eat’ is replaced with ‘please’ or ‘love’.

The ambiguity of (2) is not due to ‘ready’ being homophonous. On both readings, ‘ready’ has its usual meaning, akin to ‘suitably set, arranged, or equipped’. (Of course, ‘ready’ is polysemous; but so are ‘easy’ and ‘eager’.) *Given* meanings for the word-sized pronunciations in (2), an ambiguity remains: ‘the duck’ can be understood as the subject of ‘eat’ and associated with the role of eater, or as the object of ‘eat’ and associated with role of thing eaten.

Similarly, whatever ‘solicitor’ means for you, I bet you can understand (5)

- (5) a reporter phoned a solicitor from a small town
- as having the meaning indicated with (5a) or (5b) *but not* the one indicated with (5c).
- (5a) A reporter phoned a solicitor, and the solicitor was from a small town.
  - (5b) A reporter phoned a solicitor, and the phone call was from a small town.
  - # (5c) A reporter phoned a solicitor, and the reporter was from a small town.

The attested readings reflect distinct structures: [phoned [a [solicitor [from a small town]]]]; and [[phoned [a solicitor]][from a small town]]. But this doesn’t explain why in the second case,

---

<sup>2</sup> Drawing on Brendan Gillon’s work, I argue that while there is no requirement that lexical roots access mass concepts, count nouns cannot be used to access mass concepts.

‘from a small town’ is understood as restricting *events of phoning* a solicitor, as opposed to *individuals who phoned* a solicitor. One can say that [[phoned [a solicitor]][from a small town]] is relevantly like ‘ $\exists y:\text{Solicitor}(y)[\text{PastPhoningByOf}(e, x, y) \ \& \ \text{From-a-small-town}(e)]$ ’. But why *can’t* it be construed like ‘ $\exists y:\text{Solicitor}(y)[\text{PastPhoningByOf}(e, x, y) \ \& \ \text{From-a-small-town}(x)]$ ’? As the (5a)-reading illustrates, ‘from a small town’ can be heard as restricting a predicate of individuals. So what blocks a (5c)-interpretation of [[phoned [a solicitor]][from a small town]]?

In *CM*, I argue that meanings compose in ways that require phrasal meanings to be *monadic*, with the result that ‘phoned a solicitor’ has no variable for a phoner. The (5b) reading corresponds to ‘ $\exists y:\text{Solicitor}(y)[\text{PastPhoningOf}(e, y) \ \& \ \text{From-a-small-town}(e)]$ ’.<sup>3</sup> On this view, the grammatical subject of (5) is an argument of a covert verb that combines with ‘phone’.

Details aside, Slangs allow for *structural homophony* in ways that are unbounded yet tightly constrained. One string of words can correspond to two or more expression meanings, each reflecting a certain way in which the words can be arranged. To characterize these meanings, we need to distinguish the compositional character of phrasal meanings—which seems to be unlearned and common across Slangs, at least to a first approximation—from the arbitrary and often conventionalized character of lexical pronunciation-meaning ( $\pi$ - $\mu$ ) pairs.

We also need to think about what Slangs are. Following Chomsky, I argue that they are biologically implemented *procedures that generate* certain  $\pi$ - $\mu$  pairs. This involves arguing against Lewis and others who think the goal is to describe *sets* of  $\pi$ - $\mu$  pairs, each of which can be described in various ways by the members of a community who jointly “select” the set by adhering to certain conventions. I think this E-language perspective, ‘E’ connoting ‘extensional’, is deeply misguided. But at a minimum, we shouldn’t stipulate that Slangs are sets as opposed to procedures. We should instead ask what Slang meanings could be such that the lexical ones exhibit arbitrary homophony and several kinds of non-arbitrary conceptual equivocality, while the phrasal ones exhibit structural homophony in unbounded but limited ways.

### 3. What Meanings Aren’t

Even if we ignore the kinds of equivocality noted above, it seems clear that a single meaning can correspond to multiple concepts, given “Frege cases” of learning identities. Someone who learns that woodchucks are groundhogs, and that these animals are also called whistlepigs, might link at least one word to distinct concepts of the relevant rodents. But stressing Frege cases can make it seem that meanings are extensions. So I argue that a meaning can correspond to two or more concepts without being an extension that the concepts share. If meanings are “concept assembly instructions,” they are *further* removed from the environment than any assembled concepts.

One might worry that my proposal runs afoul of Putnam’s thought experiment involving Twin Earth, where in place of H<sub>2</sub>O there is a superficially similar though distinct substance XYZ. But if meanings are conceptually equivocal, the thought experiment is easily accommodated.

I grant that a speaker of English can use ‘water’ to access a kind-concept that applies—regardless of what he believes—to and only to samples of H<sub>2</sub>O (modulo slight impurities), while his Twin-Earth counterpart uses a homophonic word to access a kind-concept that applies to and

---

<sup>3</sup> However ‘duck’ is understood, ‘watched her duck’ corresponds to ‘ $\exists y:\text{Her-Duck}(y)[\text{PastWatchingOf}(e, y)]$ ’. Though there are twists. In (1), ‘near a muddy bank’ can modify ‘watched’; and typically, events of watching are co-located with the watchers. So it might *seem* that ‘watched her duck’ has a variable for watchers. Chapter six offers more evidence that phrases are used to construct monadic concepts, and that even in ‘gave a duck a dollar’, ‘gave a duck’ does not express a relation that holds between givers and things given to a duck. I also argue that the number of arguments a verb must combine with, to form an active voice declarative sentence, often differs from the adicity of the concept lexicalized. We have polyadic concepts of eating/snacking/noshing/dining. But ‘I ate’ is grammatical, and it implies more than ‘I ate something’, while ‘I snacked an apple’ is not grammatical.

only to samples of XYZ (modulo slight impurities). But ‘water’ can also be used to talk about the stuff from my well in New Mexico, even though that stuff has a lower percentage of H<sub>2</sub>O than Diet Coke, or a cup of tea on Chomsky’s desk. Water from city taps often contains fluoride or worse. Ocean water is salty. This suggests that ‘water’ can be used to access at least one concept C that applies to a lot of watery stuff that is chemically *less* like pure H<sub>2</sub>O than a lot of stuff that C doesn’t apply to. The details seem to involve notions of sources and functional role. So absent *argument* to the contrary, why deny that ‘water’ can be used to access a concept that applies to the mainly-XYZ-stuff from the Twin-Earth counterpart of my well? My own *intuitions* suggest that Twin-Earthers can water their lawns, occasionally sipping water from the hose.

I grant that ‘water’ and ‘star’ can be used, in mutually comprehensible ways, by speakers who have very different views about the nature of water and stars. We can talk about stuff/things in ways that don’t presuppose substantive conceptions of what we’re talking about, as if words let us express kind-concepts whose contents are fixed by paradigm cases and natural dimensions of similarity. But words have many uses. And we can use ‘meaning’ to express a kind-concept that applies to the interpretations, whatever they are, that Slangs connect with pronunciations.

Externalism about conceptual *contents* is compatible with Slangs being procedures that pair pronunciations with recipes for assembling *concepts*. For many purposes, my Twin and I can be described as using the same recipe for how to make an apple pie (or a Negroni), even if one of the steps calls for adding water (or ice), along with some apples (or vermouth). Recipes leave room for variation in which specific ingredients get used. Though if Earth\* has only Red Delicious apples, and Earth\*\* has only Granny Smiths, there may be contexts in which twins count as reading and following different instructions upon seeing ‘take six apples’ in a recipe book. Similarly, I think my Twin often uses ‘water from my well’ as an instruction for how to build a concept that applies to the water from his well. Though in cases where chemistry matters, ‘water’ and its Twin-Earth counterpart may not count as instances of the *same* word with the *same* meaning. This should be unsurprising given ‘solicitor’ and ‘robin’.

One can be an internalist about meanings and still say that (for many purposes) my words have the same meanings as words used by other thinkers, including my former self, who have or had different concepts. But it’s a trap to assume that *meanings* are whatever good translations have in common. It’s even dangerous to assume that meanings are what expressions with the same meaning have in common, since ‘same meaning’ can be a variant of ‘good translation’.

My word ‘water’ may count as having the same meaning as your homophonic word because each of us could add, to our own ‘water’-y address, the concepts accessed via the other ‘water’-y address. Likewise for ‘fish’, ‘tofu’, ‘rabbit’, ‘democracy’, etc. Much more needs to be said about *how* we exploit conceptual equivocality (and kind-concepts) to deal with the fact that we don’t always think about things/stuff in the same ways. But it can be tempting to think that “shared languages,” along with “speaking competently,” can ensure agreement about the truth conditions of sentences—as opposed to “mere” convergence on recipes for how to build thoughts. So in chapters 3-5 of *CM*, I argue that Slang sentences don’t have truth conditions, much less truth conditions that are determined by meanings.

Many pieces of this argument are unoriginal. Chapter three reviews Frege’s contributions to the study of logic, Tarski’s techniques for providing (consistent) truth theories for first-order fragments of Frege’s *Begriffsschrift*, and extensions of these techniques via the use of Church’s lambda calculus. I then discuss, in chapters four and five, two major difficulties for the Davidsonian Conjecture that a suitably formulated theory of truth for a Slang can serve as the core of an adequate theory of meaning for that language. First, it’s hard to see how there can be true theories of truth for Slangs given examples like (6), which is my favorite sentence.

(6) My favorite sentence is not true.

Second, nonsynonymous sentences can be truth-conditionally equivalent. So it's hard to see how any truth theory for a Slang could do double duty as a good theory of meaning.

By itself, neither difficulty is fatal for the Davidsonian Conjecture. But I argue that the best hope for replying to each is at odds with the best hope of replying to the other. In any case, we shouldn't assume that Slang sentences have truth conditions, especially not if this implies that (6) is true if and only if it isn't. And we shouldn't assume that a truth theory for a Slang can also be a plausible theory of understanding, given the objections pressed by John Foster and others. Moreover, I argue, each of these assumptions makes the other even less plausible.

Examples like (7) also tell against the idea that Slang sentences have truth conditions.

(7) Alvin chased Theodore around the tree gleefully.

Such examples illustrate the need for "event analyses," given that (7) implies both (8) and (9),

(8) Alvin chased Theodore around the tree.

(9) Alvin chased Theodore gleefully.

whose conjunction doesn't imply (7). The good idea was that (7-9) are understood as existential closures of conjunctive predicates:  $\exists e[\text{PastChaseOfTheodoreByAlvin}(e) \ \& \ \text{AroundTheTree}(e) \ \& \ \text{Gleeful}(e)]$ ;  $\exists e[\text{PastChaseOfTheodoreByAlvin}(e) \ \& \ \text{AroundTheTree}(e)]$ ; and  $\exists e[\text{PastChaseOfTheodoreByAlvin}(e) \ \& \ \text{Gleeful}(e)]$ . This was supposed to *support* the Davidsonian Conjecture. But a conjunct-reduction account of the implications doesn't require that (7-9) have *truth* conditions, much less that for each of these sentences, it is true if and only if some event satisfies the corresponding conjunction of predicates. In fact, requiring this leads to trouble, as sentences like (10) reveal.

(10) Theodore chased Alvin around the tree gleelessly.

Both (7) and (10) might be used, correctly, to describe a *single* episode of two chipmunks running around a tree. Alvin may have been happily chasing Theodore, who was unhappily chasing Alvin, while neither chipmunk realized who was chasing him.

There are many potential replies, involving tendentious claims about events and/or adverbial modification. But I argue that none of these replies is plausible given a moderately varied diet of examples. One important point is that while the grammatical Subject/Object asymmetry can be used to represent an Agent/Patient asymmetry, the major participants in an event of chasing (following, marrying, etc.) can be equally agentive. Related objections to the Davidsonian Conjecture can be illustrated with examples like (11).

(11) Today in London, the sun rose in the east, cars collided, and the sky was blue.

Even ignoring the polysemy of 'London', which can be used to talk about a movable polis or an immovable place, one wants to know what entities need to be posited in order to provide a plausible theory of truth for a Slang with words like 'rose', 'east', 'collide', 'sky', etc. If we agree to bracket these concerns, then alleged parade cases of characterizing meaning in terms of truth need to survive a little scrutiny.

One can say that any particular example introduces special complications. But in my view, the Davidsonian Conjecture has turned out to be promissory notes all the way down. Attention to details keeps revealing the need for more ancillary assumptions—in part because there is a deep tension between construing event analyses as aspects of a *truth* theory, concerning how some Slang is related to what really happened, and construing such analyses as aspects of a *psychological* theory of how speakers understand expressions. Put another way, the Conjecture makes some facts about action reports look like metaphysical puzzles, and it makes some facts about truth look like puzzles about how sentences are understood. I conclude that Slang sentences don't have truth conditions, and that the Davidsonian Conjecture was fruitful but false.

#### 4. What Meanings Are (Maybe)

Identifying meanings with “concept assembly instructions” is compatible with many proposals about the relevant combinatorial operations, which determine the possible types of inputs to those operations. Indeed, the number of meaning types can range from *one*—as in a Tarskian semantics that only assigns significance to sentential expressions, and always assigns satisfaction conditions—to *endlessly many*, as in a Frege-Church semantics that invokes two “basic” types  $\langle e \rangle$  and  $\langle t \rangle$ , along with the further types licensed by the recursive principle (R);

(R) if  $\langle \alpha \rangle$  and  $\langle \beta \rangle$  are types, so is  $\langle \alpha, \beta \rangle$

where expressions of the basic types denote entities or truth values, and an expression of the type  $\langle \alpha, \beta \rangle$  denotes (or has as its “semantic value”) a function from things denoted by expressions of type  $\langle \alpha \rangle$  to things denoted by expressions of type  $\langle \beta \rangle$ .<sup>4</sup>

In my view, Slangs are not Frege-Churchy in this respect. A few iterations of (R) generates millions of types that Slangs abhor, including some that are instantiated by concepts we can easily form (e.g., concepts of the “ancestral” relation that the predecessor-relation bears to the more inclusive relation of preceding). In chapters six and seven, I also argue that phrases and proper nouns are predicative—in Frege-Church terms, instances of type  $\langle e, t \rangle$ —and that there is little if any independent evidence for Slang expressions of type  $\langle e \rangle$  or  $\langle t \rangle$ .

I grant that Slangs can be used to build concepts that have non-predicative constituents. But these constituents may be uniformly *dyadic* and accessed by *lexical* items. As I show, this restrictive hypothesis is permissive enough to handle a wide range of constructions, including those covered by a typical first course in semantics. I posit several combinatorial operations for concepts and meanings, but only two meaning types:  $\langle M \rangle$  for monadic,  $\langle D \rangle$  for dyadic.

The simplest operation, “M-junction,” conjoins two monadic concepts to form a third. For example, M-joining BOTTLE( $\_$ ) with GREEN( $\_$ ) yields a concept of green bottles; and I claim that the meaning of ‘green bottle’ is an instruction that is executed by M-joining concepts accessed via ‘bottle’ and ‘green’. A second operation, “D-junction,” combines a dyadic concept D with a monadic concept M to form a monadic concept of things that bear the relation expressed by D to *something* that has the property expressed by M. For example, D-joining IN( $\_, \_$ ) with BOTTLE( $\_$ ) yields a concept of things in a bottle; D-joining AGENT( $\_, \_$ ) with REPORTER( $\_$ ) yields a concept of things done by a reporter. I think “complete” sentences correspond to “polarized” concepts that apply to everything or nothing. The idea is that a monadic concept, perhaps assembled by executing a phrasal meaning, can be used to form a propositional concept—much as the open sentence ‘Mx’ can be combined with a Tarskian prefix to form the closed sentence ‘ $\exists x Mx$ ’, which is satisfied by all sequences (of domain entities) or none, even if ‘Mx’ is satisfied by some sequences but not all. I also posit a limited form of abstraction on polarized concepts; cp. ‘ $\lambda y. \exists x (Dxy \ \& \ My)$ ’. But these operations are severely type-restricted. The resulting system is much less powerful—and much better suited to explaining *absences* of unattested expressions/readings—than familiar proposals that characterize meanings in terms of entities, truth values, functions, function-application, and a hierarchy of types.

On this view, meanings have *execution* conditions, and endlessly many Slang expressions have meanings that (unlike sets or truth values) have simpler meanings as *parts*. One can say that sentences have semantic values that are determined by values of the constituent words, given the relevant grammatical structure. But this determination thesis is, at best, an anemic explanandum. I think meanings are more like directions for how to build IKEA furniture: use a connector from box 1 to fasten a widget from box 4 to a gizmo from box 8; fasten the resulting unit to something

---

<sup>4</sup> Let’s not worry here about the difference, highlighted in *CM*, between denoters and “unsaturated” representations.

from box 5 by using a connector from box 2; cover the result with a cap from box 9; etc. I deny that meanings are language-independent extensions of concepts that get associated with Slang expressions. I think these expressions *are* pronounceable (grammatically structured) instructions for how to build concepts; cp. perceptible (diagrammatic) instructions for how to build desks.

The last hundred pages of *CM* addresses many details concerning variables, assignments, plurality, quantifier raising, the “conservativity” of Slang determiners, and the second-order character of the concepts I appeal to. But let me end this précis by stressing that just as a lexical meaning need not correspond to a single concept, the concept *lexicalized* may not be *accessible* for purposes of composition (with other lexically accessible concepts) via relevant operations.

On any plausible view, lexicalizing a concept **C**—linking it to a pronunciation and creating a corresponding expression that has a meaning—can involve using **C** to *introduce* a formally distinct concept **C\***. Given a concept that applies to ordered pairs  $\langle x, y \rangle$  such that  $x$  precedes  $y$ , it might be used to introduce (i) a concept that applies to ordered triples that include truth values, or (ii) a higher-order concept that can combine with concepts like EVERYTHING and SOMETHING, or (iii) a concept that applies to events of one thing preceding another. Lexicalization need not be a mere process of labeling. Frege showed us how to introduce concepts of some logically interesting types. I think children, driven by a boring natural logic that is geared to predicates and predicate reduction, use Slangs to introduce a stock of lexically accessible concepts that are systematically combinable but *much less varied* than the diverse concepts that get lexicalized.

In the end, *CM* is an argument that meanings play a large role in how humans acquire and combine the concepts we express with words, but focusing on truth/reference/communication is a distraction if we want to find out what meanings are. Public uses of Slangs are obvious, and they are often valuable. But pronunciation may be a noisy addition to procedures that generate recipes for how to build concepts.