

Chapter One: Introduction (5 September 2020)

Humans regularly acquire and use languages of a special sort. These languages—let’s call them Slangs—have expressions that are meaningful, pronounceable, and combinable in distinctive ways. Slangs can be described as generative procedures, internal grammars, or “I-languages” in Chomsky’s (1986) sense. These psychologically encoded procedures generate expressions that connect meanings of some kind with pronunciations that may be spoken or signed. Indeed, a typical Slang generates endlessly many such expressions from finitely many lexical items that have been memorized. This invites questions about how complex expressions and their meanings are related to lexical items, human cognition, and the use of Slangs in communication.

1. Agenda

Pietroski (2018) argues that lexical meanings are biologically realizable instructions for how to access concepts of a special kind, and that phrasal meanings are composite instructions for how to assemble concepts from lexically accessible components.¹ According to this “internalist” proposal, phrases are pronounceable recipes for how to make complex concepts, and a single lexical item can be associated with a diverse family of concepts. As we’ll see, lexical items are polysemous in ways that tell against two “externalist” claims, either of which can be viewed as a corollary of the other: Slang meanings determine extensions; and declarative sentences have compositionally determined truth conditions. These bold claims became standard in part because they were adopted by Davidson (1967a, 1973), Lewis (1970, 1975), Montague (1974), and other influential authors. But many facts suggest that meanings don’t determine extensions, and that *pace* Putnam (1975), the meanings of Slang expressions just ain’t outside the head. After a dark age of behaviorism, Chomsky (1957, 1959a, 1959b, 1964, 1965, 1966) had revived the idea that meanings are mental representations. And this idea can be developed with a twist invited by studies of grammar: meanings are instructions, generated by Slangs, for how to build concepts.

This proposal may now sound radical. But I think it is attractively retro. So in the pages that follow, there will be two interwoven themes. One concerns the nature of human linguistic expressions and their meanings. Another concerns intellectual history and the implausibility of some currently common assumptions about how meaning is related to truth.

1.1 A Call for Restoration

In chapter two, I summarize some points from Chomsky (1957, 1959a, 1964) in a way that includes the largely ignored conception of meaning that Chomsky combined with his celebrated account of the syntactic structures exhibited by ordinary sentences. His suggestions about how to characterize meanings—partly in terms of how complex expressions are generated, and partly in terms of constraints on how lexical items can be used in acts of reference—were programmatic. But in my view, they were superior to the subsequent truth-theoretic alternatives. Chomsky embedded his claims about meaning in an independently attractive conception of the grammars that children naturally acquire and the adequacy conditions on proposed descriptions of these grammars. By contrast, I think Davidson (1967a) initiated an unfortunate tradition of assuming that Slang expressions have truth-theoretic properties, and that characterizing these properties is a free-standing enterprise which can be conducted in abstraction from empirical studies of human psychology.

¹ At least for these purposes, sentences are special cases of phrases. I take concepts to be composable mental symbols with which thinkers can think about things; see, e.g., Fodor (19xx, 20xx), Laurence and Margolis (20xx), Rey (201x), and chapter two of Pietroski (2018). This leaves room for many proposals about the types of concepts that correspond to Slang expressions. Though in my view, many considerations favor a rather spare typology.

Things could have gone differently. I'll argue that Davidson's (1967b) discussion of action reports like (1-4) did not support externalism about Slang meanings.

- (1) a man buttered some toast at midnight with a knife
- (2) a man buttered some toast at midnight
- (3) a man buttered some toast with a knife
- (4) a man buttered some toast

His insights regarding such sentences cohered better with the internalist perspective outlined by Chomsky (1964), who had offered relevant observations about examples like (5) and (6).

- (5) a woman saw a boy walking towards the train station
- (6) a woman saw a boy walk towards the train station

Davidson didn't address these observations. But his proposed "event analyses" could have been profitably developed within Chomsky's framework.

Davidson highlighted implications. Deleting 'with a knife' from (1) seems valid, as if (1) and (2) are strings of words that somehow get understood as abbreviations of (1a) and (2a).

- (1a) Some event was such that it was a buttering of some toast by a man,
and it occurred at midnight, and it was done with a knife.
- (2a) Some event was such that it was a buttering of some toast by a man,
and it occurred at midnight.

As discussed in chapter three, the plausible idea that many grammatical modifiers correspond to conjuncts fits with an old "natural logic" tradition that stressed the role of conjunct reduction in valid inferences like the following: a woman from Boston won a big prize; hence, a woman won a prize. This tradition can be fruitfully extended, in ways that apply to (1-4), by analyzing sentences in terms of quantification over events. But this suggests that the relevant intuitions regarding compelling inference reflect derivational properties of sentences, as opposed to alleged truth conditions, especially given Chomsky's examples of how modification can be construed.

Note that (5) can be understood in either of the ways indicated with (5a) or (5b).

- (5a) A woman saw a boy who was walking towards the train station.
- (5b) A woman saw an event of a boy walking towards the train station.

But this kind of ambiguity is limited in Slangs. String (6) can only be understood in the second way, with the implication that an event of walking was seen; and (7) is likewise unambiguous.²

- (7) this the station that a woman saw a boy walking towards

Chomsky emphasized such constraints in his derivational conceptions of syntactic structure and phrasal meaning. He didn't offer specific accounts of the attested implications. But he did suggest an analogy to certain relations that are exhibited by pronunciations, as in (8) and (9).

- (8) a baker buttered a bun in the middle of the night
- (9) but he liked the bun so much that he never took a bite

Slang expressions can share parts that interface with human perceptual and articulatory systems in ways that are experienced as alliteration or rhyme. Corresponding intuitions—e.g., that 'bite' rhymes with 'night'—can be explained internalistically. Competent speakers of English represent 'bite' and 'night' as sharing a final phonological part, despite variation in the acoustic waves produced when these words are uttered. Similarly, (1) and (2) may share parts that interface with human concepts of events in ways that speakers of English experience as felt implications, however the expressions are related to potential contexts of use.

² Unlike (4) or (5), (3) also has a reading on which the woman saw the boy while she was walking. Davidson (1967a, n.5 and n.14) mentions some of Chomsky's work approvingly, though he doesn't cite Chomsky (1964); cp. Higginbotham (1983), who offered event analyses of perceptual reports like (4).

I think this psychologistic conception of the linguistic phenomena was well motivated given work by Chomsky (1957, 1959b, 1964, 1965, 1966) and others, while Davidson and many others remained unduly influenced by Quine's (1960) unjustified behaviorism. But in any case, we shouldn't assume that intuitions of implication provide evidence for describing meanings in terms of truth. Speakers often make truth-evaluable claims by pronouncing meaningful sentences in suitable contexts. Correlatively, many sentential utterances are true or false. But it doesn't follow that the corresponding sentential meanings map contexts of utterance onto truth or falsity.

I'll argue that there is a sense in which (1) implies (2), and (10) implies (11).

(1) a man buttered some toast at midnight with a knife

(2) a man buttered some toast at midnight

(10) a girl who sang won a big prize

(11) a girl won a prize

But like most words, 'implies' is polysemous. For purposes of theorizing, and clarifying the vague idea that one Slang sentence can imply another, we can and should distinguish a technical internalist notion of "intailment"—exhibited by certain instructions for how to build human concepts—from various externalist notions of entailment that are exhibited by truth-evaluable assertions, propositional contents, invented formulae, or mental representations that can be constructed by executing linguistically encoded instructions in careful ways.

The idea will be that typically, deleting grammatical modifiers yields examples of intailment, some of which correspond to examples of entailment; though interesting notions of entailment are not limited in the ways that Slangs are. I think the phenomenon of intailment reflects a natural logic. But valid inference is a broader notion, as Frege (1879, 188x) stressed. Many notions of entailment also cover cases in which premises somehow "necessitate" conclusions that don't follow as a matter of logic. (Consider: there is some water on Hesperus; hence, there is some H₂O on Phosphorus.) In my view, the more inclusive notions are neither needed nor wanted in theories of meaning for Slangs. We need a notion of intailment according to which (12) intails (13), but not (14) or (15), whatever we say about truth and necessity.

(12) a unicorn that flew to Hesperus saw a big centaur

(13) a unicorn saw a centaur

(14) a unicorn saw a centaur and there are many prime numbers

(15) a unicorn that flew to Phosphorus saw a big centaur

I'll return to the different idea of defining "models for Slangs" so that if a model makes (1) true, it makes (2) true. But it's hard to see how any such stipulations can explain anything. One can also define models so that some of them make (1) and (10) true while making (2) and (11) false. As we'll see, this is part of what led Lewis (1970) and Montague (1974) to describe languages like English and Japanese in terms of a "general" framework that also covers formal languages in which the analogs of (10) and (16) are instances of the same combinatorial form.

(16) an alleged criminal bought a fake diamond

In such languages, the analogs of (10) fail to imply analogs of (11), much as the Slang sentence (16) fails to imply the Slang sentence (17).

(17) a criminal bought a diamond

Or as I'll put it, we can invent languages in which the best translation of (10) doesn't intail the best translation of (11). But (10) still intails (11), and (1) intails (2). I'll argue that *pace* Lewis and Montague, theories of meaning should help explain why endlessly many phrases are actually understood as conjunctive predicates, and theorists shouldn't assume that Slang expressions are importantly like the expressions that logicians have invented to study truth.

This will connect back to the main point of chapter two. Slangs are languages of a special kind. They generate meaningful expressions in distinctively human ways, yet they are importantly unlike the notational systems we invent for scientific purposes. Slang expressions are not mere formulae that can be interpreted and recursively combined. The lexical items that children acquire have meanings of their own; and while it's hard to describe the relevant modes of composition correctly, it's easy to invent languages that are recursive in other ways. This invites Chomsky's (1957, 1959a) project of trying to characterize the expression-generating procedures that we naturally acquire, without insisting that these procedures be described in ways that highlight analogies to artificial languages that were designed for various purposes.

Davidson, Lewis, and Montague suggested a very different project: assume that languages like spoken English are sets of sentences, and that these sentences are strings of sounds that have truth conditions; then try to characterize the truth conditions, using whatever vocabulary seems best, without worrying about any procedures internalized by actual speakers. Many philosophers, influenced by Quine (1960), were skeptical about the prospects of finding a psychological basis for linguistic meaning. This fed a suspicion that any real semantic properties of Slang expressions are semantic in Tarski's truth-theoretic sense, and that theories of meaning should not invoke Fregean senses or mentalistic meanings that speakers associate with certain sounds; cp. Frege (1892a), Locke (17xx), Katz and Fodor (1963), Katz and Postal (1964).

Given this starting point, every Slang sentence provides an opportunity for productive work. I think the insights obtained can be preserved, and described more accurately, without externalist rhetoric. But my aim isn't to criticize particular proposals that happen to be formulated in truth-theoretic terms. My objections will be targeted at the less anthropocentric conception of meaning that accompanied the proposed replacement of Chomsky's project with less mentalistic approaches to describing the meanings that Slangs connect with pronunciations.

If we describe Slangs as generative procedures that have semantic components whose character has to be discovered—and we view theories of meaning as proposals about aspects of the grammars that speakers use in producing and understanding speech—then as Chomsky stressed, the adequacy conditions on theories of meaning (syntax, phonology, acquisition, etc.) are demanding. Claims about meaning should cohere with discoveries regarding the range of grammars that children can naturally acquire. Hypotheses in this domain are exposed to a host of empirical risks, including many that theorists may not yet envision, since the enterprise is a branch of inquiry regarding certain aspects of nature. Explananda are not given in advance, and this makes the job hard, in ways familiar from the history of science. But theorists are not hostage to “philosophical” stipulations about what the explananda are.³

Nonetheless, early advocates of truth-theoretic semantics insisted that theories of meaning specify truth conditions for Slang sentences, and that such theories need not reflect the ways that speakers generate or understand these sentences. Chomsky's claims about syntactic structure were declared to be, at best, partial descriptions of how humans happen to connect sentential sounds with truth conditions that could be specified equally well in other ways by aliens who would count as equally competent speakers of the languages (i.e., sets) in question. Given a language, *L*, the goal was sometimes described as formulating a theory, *Θ*, such that knowledge of *Θ* would *suffice* for understanding the sentences of *L*; where it was assumed that the actual speakers of *L* need not know *L*, or even represent its sentences in the way that *Θ* does. The primary targets were the alleged truth conditions, not the minds of actual speakers.

³ Chomsky (1957, 1959a, 1959b, 1964) provides a much better example of Quinean “naturalized epistemology,” applied to the study of human languages, than Quine (1960); see Hornstein (20xx) for enlightening discussion.

I think this anti-psychologistic project was, from its inception, a bad idea. The operative notions of language and understanding were too far removed from the natural phenomena that Chomsky and others had already identified. Insisting on a framework that is geared to invented languages, and trying to make Slangs fit the mold, seems like insisting that biologists make room for unicorns and centaurs at the risk of misdescribing horses and humans. Instead of starting with ground rules about how expression meanings are related to truth, we should return to the idea that hypotheses about Slangs require defense, and then *ask* whether or not Slang sentences have truth conditions that theories of meaning should specify.

1.2 Resisting Externalism about Slang Meanings

Focusing on truth rather than psychology did invite further questions that Davidson (1973, 1976) and Lewis (1970, 1975) addressed: what does a theory of meaning for a language need to provide in addition to mere specifications of the relevant sentences and truth conditions; and in virtue of what do sentences have their truth conditions? In my view, these were artificial questions that led to implausible speculations about how meaning is related to logic, necessity, imagined practices of “radical interpretation,” and conventional uses of Slangs. But in debates about these topics, it was often conceded—at least for purposes of argument—that declarative sentences have truth conditions; and repeated concession can foment orthodoxy.

Other factors also made it *seem* less implausible that each declarative sentence has a semantic character that determines (or perhaps just is) a mapping from contexts to truth values. Partee (1975) and other linguists showed that by judiciously employing the formal apparatus from Lewis and Montague (1974), one can formulate and sometimes resolve puzzles concerning many specific constructions.⁴ Moreover, the project of truth-theoretic semantics seemed to cohere with Kripke’s (1970) discussion of reference and Putnam’s (1975) slogan that “any way you cut the pie, meanings just ain’t in the head.” But the project also faced trenchant objections; see, e.g., Harman (1974), Foster (1976), Dummett (1976), Chomsky (1975, 1977, 1980). And in retrospect, I think it’s clear that externalism about linguistic meaning was far less motivated than externalism about the contents of thoughts often expressed by using meaningful expressions.

In chapter four, I’ll argue that Kripke’s (1973) discussion of nouns like ‘unicorn’ and ‘Hamlet’ could have and should have been read as a reminder that (i) lexical items are polysemous, and so (ii) we shouldn’t assume that Slang expressions have extensions. We can use words like ‘horse’ and ‘Secretariat’ to express concepts, some of which may have extensions or denotations. But this doesn’t show that Slang expressions, or even nouns, have meanings that

⁴ Davidson’s (1967b) insights about implication were often ignored, in part because it was assumed that Slang sentences have representation-neutral truth conditions that *can* be specified without quantifying over events, regardless of how humans happen to understand these sentences. To some extent, Davidson encouraged this reaction by (i) focusing on the alleged truth conditions, (ii) trying to be agnostic about how speakers associate these truth conditions with pronunciations, and (iii) not explicitly embracing the hypothesis that speakers recognize the relevant implications as instances of conjunct reduction in the scope of suitable quantifier. But it still seems odd that while Lewis (1970) and many others assumed that a verb phrase like ‘poked Porky with a pencil’ indicates some function *F* that maps entities to truth values, and that the prepositional phrase indicates a higher-order function that maps functions like *F* to functions of the same type, the alleged higher-order functions went unspecified. It was thought sufficient to say that ‘with a pencil’ indicates a function that maps the function indicated with ‘poked Porky’ onto a function that maps each entity, *e*, to truth or falsity depending on whether or not *e* poked Porky with a pencil’. Though even if this function exists, any such description of it is woefully indefinite. In retrospect, one wonders why it wouldn’t have been equally (un)satisfactory to remain silent about the meaning of ‘everyone’ but insist that the direct object in ‘poked everyone’ indicates a function that maps the function indicated with ‘poked’ onto a function that maps each entity, *e*, to truth or falsity depending on whether or not *e* poked everyone. I’ll return to the contrast between recursive specifications of extensions and the mere existence of abstract mappings from sets to sets.

determine extensions. I grant that words like ‘horse’ and ‘unicorn’ can be used to express atomic concepts whose contents are rigid (in Kripke’s technical sense) and determined externalistically. Perhaps these contents can be described in terms of functions, each of which maps every possible world to a set of possible things that “exist at” that world. But even if we often use words to express concepts whose contents are plausibly identified with such functions, it doesn’t follow that the words have meanings that determine these contents/functions relative to contexts of use.

In my view, Putnam’s (1975) influential Twin-Earth thought experiment never supported the assumption that meanings determine extensions. On the contrary, Putnam simply adopted this assumption in the course of making a point that is better described as an insight about certain conceptual contents.⁵ We can use ‘water’ to express a concept whose content is fixed, externalistically, in a way that has a striking consequence: this concept applies only to samples of H₂O, modulo trace impurities, even for thinkers who have no concepts of Hydrogen or Oxygen. But this doesn’t show—or even suggest—that we can’t use ‘water’ to express concepts that apply to other stuff; see Pietroski (201x) and the discussion of polysemy in chapter four.

It’s important to distinguish causal conceptions of reference, which were often motivated by criticisms of earlier “descriptivist” conceptions, from a much less plausible idea: the words used in acts of reference have extensions that are somehow determined by causal connections between the acts and corresponding referents.

⁵ While it can be hard to follow the dialectic in Putnam’s paper, he concedes that talk of words having extensions is a “very severe” idealization (p. 133), and he seems to grant that he is abstracting away from polysemy. He then says that two familiar “assumptions” are “not jointly satisfied by *any* notion, let alone any notion of meaning” (135-36): (i) “knowing the meaning of a term is just a matter of being in a psychological state,” and (ii) the meaning of a term determines its extension; where (i) is to be understood in terms of methodological solipsism, so that no psychological state “presupposes the existence any individual other than the subject to whom that state is described.” Putnam then takes (ii) as a premise in arguing against (i). Thirty pages later (165), his conclusion is unconditional: “The traditional problem of meaning splits into two problems. The first is to account for the *determination of extension*. Since, in many cases, extension is determined socially and not individually, owing to the division of linguistic labor, I believe this problem is properly a problem for socio-linguistics.” But at no point does Putnam argue that ‘water’ has an extension, much less one that is determined socially.