

Semantic Internalism Paul M. Pietroski

Putnam (1975) summarized his main point memorably: cut the pie any way you like, meanings just ain't in the head! Chomsky (1964, 1965, 1977, 1995, 2000) offers a different view according to which meanings are generated by and internal to human minds.

Putnam claimed that 'water' is tied to H₂O in a way that precludes using sentence (1)

(1) That is water.

to correctly describe a sample of some other substance. Behind this claim lie familiar thoughts: the meaning of 'water' determines an extension for this word; the extension of 'water' is the set of water samples, which are samples of H₂O, modulo impurities; and an utterance of (1) is true only if the speaker indicates something in the extension of 'water'. Putnam asked us to imagine a planet like Earth, except that the H₂O is replaced with different stuff (XYZ) that is superficially similar and equally potable. He then invited us to share his judgment that (1) cannot be used to correctly describe the watery stuff on Twin Earth. Chomsky offers reasons for rejecting this judgment and denying that 'water' has an extension. Moreover, on Chomsky's view, (1) can fail to be a correct description of some stuff—e.g., weak tea—that *is* H₂O modulo a few impurities.

Chomsky usually motivates his conception of meaning by way of examples, in a style reminiscent of Austin (1961, 1962) and others who heeded Wittgenstein's (1953) advice to reflect on the many ways that expressions can be used. Chomsky's contributions to semantics, unlike his systematic work in syntax, have also tended to be critical. But he highlights a tension that others have felt: examples that are often ignored, or downplayed, motivate skepticism about the prospects for theories of meaning; yet beneath the complexities of language use, there are "meaning facts" to be described and explained, if only we had the right theoretical vocabulary. Chomsky suggests that semantics, done right, could reveal how some familiar but theoretically intractable phenomena—in particular, judgment and reference—are related to and constrained by various aspects of cognition. From this perspective, the challenge for theorists is to characterize meaning without relying on *commonsense* notions of judgment and reference, or misdescribing the phenomena by using inappropriate *technical* notions of truth and denotation.¹

1. Ambiguity as a Guide to Meanings

Despite its title, Putnam's essay isn't about the word 'meaning'. But it is hard to discuss meanings without metalinguistic comment, since the topic is not obvious. Moreover, speakers of English can use (2) to ask various questions.

(2) What are meanings?

1.1 Strings vs. Structures

The meaning of a red traffic light differs from the meaning of 'a red traffic light'. For suitably educated humans, both the light and the phrase carry significance. But the phrase is composed of meaningful parts, in ways that the light is not. We can also distinguish what a speaker meant, by using an expression, from the meaning of the expression used. The noun 'meaning' can be used to talk about sundry things. But let's focus on expressions of the "Human Languages" that children can naturally acquire given ordinary experience. I assume—along with Chomsky and Putnam, and *pace* Quine (1960)—that these expressions have meanings that we can talk about.

¹ Hornstein (1986) develops this theme. See also McGilvray (1998, this volume), Pietroski (2005), Stainton (2008) and references there.

There is another sense in which an interrogative sentence like (3) is not univocal.

(3) What are quasars?

Someone who overhears a conversation might use (3) to ask for hints about the general topic, and ideally, an identifying description of quasars. Given a suitably placed telescope, such a description might be provided via (4).

(4) They're things like the thing that looks fuzzier than the surrounding stars.

Someone else who knows what the conversation is about—and in that sense already knows what quasars are—can use (3) to ask about the nature of the quasi-stellar objects in question. Are they stars of a special kind, especially bright galaxies, or what? Similarly, we can ask which things meanings turn out to be. Are they mental objects of some kind, publicly available entities that can be signified in certain ways, or what? But initially, we don't even know if meanings are spatiotemporally located. So before asking about their nature, one might want some descriptions that help identify what we're talking about.

Chomsky (1957, 1964, 1965) noted some *data* that is directly relevant. We know a lot about meanings, despite not knowing what they are. For example, while (5) is a string of words,

(5) to persuaded woman the leave John

it does not support a sentential meaning. By contrast, (6) and (7) can be understood as sentences.

(6) the woman persuaded John to leave

(7) the woman expected John to leave

Though (6) and (7), paraphrased with (6a) and (7a), have meanings that somehow differ in kind.

(6a) The woman persuaded John that he should leave.

(7a) The woman expected that John would leave.

We also know that (8) can be understood in two ways, indicated with (8a) and (8b).

(8) the duck is ready to eat

(8a) The duck is prepared to dine.

(8b) The duck is fit for consumption.

This ambiguity is not due to any ambiguity of 'duck', 'ready', or 'eat'. Given the word meanings, there is still a *structural ambiguity*, corresponding to whether 'the duck' is understood as the subject or object of 'eat'. However, (9) has only one meaning, as does (10);

(9) Darcy is eager to please

(9a) Darcy is eager to be one who pleases relevant parties.

(9b) #Darcy is eager to be one whom relevant parties please.

(10) Darcy is easy to please

(10a) #It is easy for Darcy to please relevant parties.

(10b) It is easy for relevant parties to please Darcy.

where '#' indicates a sentential meaning that the string in question fails to have.

Examples like (8-10) illustrate what Chomsky is talking about when he talks about meanings: (8) has more than one; and the meaning of (9), unlike that of (10), associates the meaning of 'Darcy' with the subject position of 'please'. There are unboundedly many such examples of constraints on ambiguity. And we can recognize subtle distinctions. Consider the three possible interpretations of (11).

(11) the woman saw the boy walking towards the railway station

(11a) The woman saw the boy while walking towards the railway station.

(11b) The woman saw the boy who was walking towards the railway station.

(11c) The woman saw the boy walk towards the railway station.

Sentence (11a) implies that the woman walked towards the station. Both (11b) and (11c) imply

that boy walked. But (11b), unlike (11c), can be used to describe a situation in which the woman saw the boy without seeing him walk; see Chomsky (1964, p.73). So it is striking that (12)

(12) this is the railway station that the woman saw the boy walking towards
is unambiguous. It has only the meaning corresponding to (11c), with the implication that the woman saw the boy walk. In general, each string of words will have n but not $n+1$ meanings, for some number n . And whatever meanings are, (11) has three of them, while (12) has only one.

Such examples also suggest that meanings are generable mental representations of some kind. Understanding (8) in different ways seems like “seeing” a Necker cube in different ways.

(8) the duck is ready to eat

Putnam (1975) focused on word meanings rather than sentence meanings, because he felt that “our concept of word-meaning” was especially defective (p. 132). But if the goal is to figure out what meanings *are*, then constraints on structural ambiguity may be more important than facts regarding what we talk about with certain nouns. Still, focusing on *lexical* ambiguity may help.

1.2 Homophony vs. Polysemy

In one respect, (8) is like the pronunciation *ber*, which can be used to express more than one meaning. The adjective in ‘bare skin’ is a homophone that shares its pronunciation with several words, including some that get spelled differently. Whatever words of a spoken language are, they connect pronunciations with meanings of some kind. But there is no such thing as *the* meaning of ‘bear’, or *the* pronunciation of ‘sow’, which can be a rhyme for ‘hoe’ or ‘how’. So let’s say that one English word connects *ber* with the meaning **bear**¹, thereby connecting some uses of *ber* with certain ursine animals. Another word connects the same pronunciation with another meaning, **bear**², which concerns episodes or states of carrying or supporting something.

A word can also be polysemous in ways that are hard to distinguish from homophony. We speak of bearing gifts, weight, malice, pain, a likeness to relatives, children, names, false testimony; things can bear watching; citizens may bear arms. According to dictionaries, this variation reflects at least two verbs, as opposed to one verb meaning with many “subsenses” that correspond to carrying or supporting or having or yielding or tolerating or meriting something. (In saying that facts bear on a question, or that a road bears right, we use other meanings.) Yet however many words lexicographers posit, the phenomenon of polysemy remains, even if it is hard to characterize. Perhaps when we talk of someone bearing their pain, we extend a core meaning that lets us talk of posts bearing weight; and perhaps talk of bearing malice towards others, or “bearing relations” to other things, is a more extended extension. But talk of bearing pain isn’t metaphorical, at least not for living speakers of English.

We can use ‘door’ to talk about certain objects that people cannot walk through. Indeed, this feature of doors is importantly related to their usual function. But we can also use ‘door’ to talk about certain spaces, often occupied by a door, and issue executable instructions like (13).

(13) Walk through the first door, turn right, and then exit through the front door.

If it is important to disambiguate, we can call the relevant spaces doorways. Nonetheless, ‘door’ is polysemous. Words like ‘book’ and ‘country’ exhibit different kinds of concrete/abstract duality, as illustrated with (14) and (15).

(14) This book is too heavy to carry, and the other one got a good review.

(15) This country’s geography is even more spectacular than its political history.

And note the contrast with (16).

(16) This bear sold his stocks, and the other one protected her cubs.

The pronunciation *ber* can be used with a Wall Street meaning, or with the more frequently used meaning **bear**¹. But one can’t exploit this homophony—except as a joke—in sentences like (16).

The meaning of ‘book’ that we use to talk about books on a shelf, as opposed to booking a trip, is evidently the meaning we also use to talk about reviewing/writing/downloading a book. Given the many examples of this sort, we need to distinguish homophony from polysemy. This simple point becomes important when combined with three related observations.

First, lexical homophony is often conceptually *arbitrary* in ways that lexical polysemy isn’t. Correlatively, we could connect *ber* with more meanings if we want to, perhaps using it to talk about some virtue of a new technological device, as in (17).

(17) The new i-gizmos are bair.

But polysemy allows for “creative extensions” of word usage that are more constrained. The possibilities for polysemous use seem to reflect *and be limited by* how humans are apt to categorize things. Here, Wittgenstein’s (1953) talk of “family resemblance” seems apt.

Second, *structural* homophony is common but severely constrained. As (11) reminds us, (11) the woman saw the boy walking towards the railway station there are boundlessly many English homophones. Yet there are limits, as (12) shows.

(12) this is the railway station that the woman saw the boy walking towards

And while we could have used the pronunciation of ‘eager’ to express the meaning of ‘easy’, we can’t make (9) or (10) structurally homophonous in the way that (8) is.

(9) Darcy is eager to please

(10) Darcy is easy to please

(8) the duck is ready to eat

Third, polysemy is apparently confined to lexical items. Consider ‘my book’, which can be used to talk about a book I own, a book I wrote, or a book temporarily assigned to me. The word ‘my’ presumably combines ‘me’ with a polysemous possession morpheme. But combining ‘my’ with ‘book’ doesn’t seem to be a locus for further polysemy. *Prima facie*, something counts as my book (relative to a given context), if and only if it counts as both mine and a book (relative to that context). The conjunctive significance of modifying a noun with a relative clause, as in ‘book I own’ and ‘book I wrote’, also seems to be constant across cases. Far from enhancing polysemy, relative clauses often reduce the plausible options and favor a certain use of the head noun. A *book that was defaced* is spatially located, while a *book that was plagiarized* is abstract, even if the plagiarizer copied from a book on the shelf.

It is as if ‘book’ is connected with both a way of thinking about certain inscriptions of certain contents, and a way of thinking about those inscribable contents. But even if ‘book’ can be used to access either way of thinking about some things that get called books—or more briefly, either *concept*—a particular context might press in favor of accessing a particular concept. And if non-linguistic factors can create such pressure, then so can the use of certain modifiers. I’ll return to this idea. For if lexical meanings are used to access concepts, or what Chomsky calls “perspectives,” this also suggests that meanings are internal to minds.

1.3 Flexible Modifiers

Once we grant that polysemy is ubiquitous, we can avoid some implausible hypotheses regarding structural homophony. But the net result casts doubt on the idea that words have extensions.

As (18) illustrates, ‘green’ can modify a mass noun or a count noun, singular or plural.

(18) The paint is green, and so is the house, and so are the apples.

For simplicity, let’s ignore other uses of ‘green’ as in (19), and focus on the adjective.

(19) Green is his favorite color because greens look good on him.

As (20) and (21) remind us, mass/count/plural contrasts need not be phonologically marked.

(20) The raw fish is from a fine fish that lived among other fish.

(21) The green fish is from a green fish that lived among other green fish.

But ‘fish’ is not three ways homophonous in these sentences, and neither is ‘green’.

The plural noun ‘fish’ is presumably the result of combining the singular noun with an unpronounced analog of the plural morpheme in ‘fishes’; cp. ‘sheep’ and ‘deer’. If the singular noun combines a lexical root with a covert count morpheme, then the homophonous nouns in (20) can be represented as ‘fish’, ‘fish+CT’, and ‘fish+CT+PL’; see Gillon (1987), Bale and Barner (2009). It can be tempting to say that the lexical root is a “mass noun” that applies to, and only to, fish-stuff. But then it’s hard to specify a meaning for the count morpheme. I follow Gillon in thinking that the root meaning is neutral; ‘fish’ can be used to access a way of thinking about stuff as fish, *or* a way of thinking about a single thing as a fish, *or* a way of thinking about some things as fish. If the count morpheme precludes the first use, and the plural morpheme precludes the second, then absence of the count morpheme may be correlated with intentions to talk about uncountable fish-stuff. But the root can still be the least restrictive form.

This is compatible with the singular count noun being acquired first. Thinking about something as a fish may be *conceptually* basic, at least for children who don’t see fish sticks before they see any fish. A typical course of acquisition might proceed as follows: a child somehow acquires the concept FISH, with which one can think about something as a fish; connecting this concept with a pronunciation leads to acquisition of a count noun, ‘fish+CT’; the child eventually forms concepts of fishes and fish-stuff; and then the lexical root can be used to access any concept in the child’s family of ‘fish’-concepts.²

One can try to avoid psychological commitments by saying that ‘fish+CT’ is *true of* an entity *e* if and only if *e* is a fish. On this view, the count noun has an extension—{*e*: *e* is a fish}. Perhaps the meaning of the plural morpheme maps this set onto a set of “plural entities” that are composed of individual fish, so that the extension of ‘fish+CT+PL’ is {*p*: *p* is a plurality of fish}. And maybe the extension of the lexical root is either {*m*: *m* is a portion/quantity of fish} or the union of the three relevant sets; cp. Cartwright (1963), Link (1983).

As Chomsky notes, such idealizations often do no harm, given how little we know about the underlying psychology. It is simpler to say that expressions have extensions, instead of saying that expressions are used to access concepts that sometimes approximate idealized representations that have extensions. But it’s hard to see how ‘fish+CT’ could be suitably related to fish, much less the alleged *set* of fish, if not via mental representations of fish. Moreover, even if ‘fish’ has an extension that can be specified in terms of countable fish and/or the stuff that constitutes them, it’s even harder to specify the alleged extension of ‘green’. While any portion of a fish is fish, it’s not true that any portion of a green thing is green. A green house or a green apple may only be green on the outside; cp. a green cave. Green paint is green all the way through, but green pasta doesn’t have to be.

One can speculate that ‘green’ applies, primitively, to both green stuff and green things—and that the requirements on being green vary from paint to pasta, houses, apples, and caves. But why think ‘green’ *has* an extension? We don’t need to posit a set of all and only the green things, and a set of all and only the portions of green stuff, in order to say that ‘green’ can be used to access various ways of thinking about things or stuff as green.

² This may open the door to further polysemy. I assume that concepts are mental representations of a special kind; see, e.g. Fodor (197, 1986, 1998, 2003), Laurence and Margolis (1998). For ‘tofu’, typical acquisition presumably begins with acquisition of a mass concept. Though one can imagine a child who thinks that tofus are farm animals. For relevant discussions of “roots,” see Borer (2005), Halle and Marantz (1993), Harley and Noyer (2003).

2. Meanings, Concepts, and Theories

I think Putnam (1975) underestimated the importance of polysemy. But let me first note an important motivation for Putnam's view about 'water', and then say how Chomsky can accept the point.

2.1 Leaving Room for Disagreement

Whales are mammals, not fish. Stars are giant spheres of gas, not holes in the canopy. Water is a compound of hydrogen and oxygen, not a basic element.³ Such discoveries highlight Platonic questions that Putnam (1962, 1975) addressed: how can we talk about the same things/stuff, across generations, in ways that allow for deep disagreements about the nature of the things/stuff in question? If scientists know enough to know what they're talking about, how can they be so spectacularly wrong about their subject matter? How can we think and talk about a common subject matter that each of us has misdescribed in our own way?

Kripke (1980) and others noted that part of the answer lies with causal connections. Someone who perceives a few whales (stars, etc.) is positioned to make claims about them, and then spread the word, even if the claims are badly mistaken; cp. Kuhn (1962). This point is often expressed in terms of alleged extensions/denotations of *nouns* like 'star' and 'Aristotle'. But this is inessential simplification. A gold star on a homework assignment is not a counterexample to modern astronomy, and there has been more than one Aristotle. One can reply by positing more homophony. But Kripke's points may reflect the character certain *concepts* that we can access and express with certain uses of nouns; see, e.g., Evans (1982).

Put another way, one can deny that words have extensions, but grant that words are sometimes used to articulate scientific ways of thinking about things/stuff. Indeed, when Chomsky talks about science, he often stresses that humans can introduce *technical* notions that let different people think about—and disagree about—the nature of the things/stuff in question; see chapter one of McGilvray (2014). Children may well acquire some “natural kind concepts” that approximate the idealized terminology that Putnam had in mind; see, e.g., Keil (1992).

In some contexts, an utterance of (1) is relevantly like an utterance of (22).

(1) That is water.

(22) That is a sample of H₂O, perhaps modulo some impurities.

For example, one might tell a story in which atomic structure is important, thereby creating a context in which 'water' is used to access a natural kind concept. Relative to such a context, (1) might be as wrong as (22) if the indicated stuff is not H₂O. But it doesn't follow that (1) is false relative to *every* context in which the indicated stuff differs atomically from H₂O. Putnam offered a particular story, in the context of his essay, and invited us to conclude that the watery stuff on Twin Earth *can't* be correctly described (by speakers of Earth-English) as water. My own judgment is that this modal conclusion is implausible, and that far from being “intuitive,” it reflects the assumption/idealization that 'water' is not polysemously connected with a family of concepts. But in any case, Putnam's claim about 'water' is a hypothesis, not a datum.

By contrast, it is a datum that (9) is unambiguous as indicated below.

(9) Darcy is eager to please

(9a) Darcy is eager to be one who pleases relevant parties.

(9b) #Darcy is eager to be one whom relevant parties please.

Introspection may reveal that 'water' *can* be used to access a concept that would not apply to

³ Fire, air, and earth are more even complicated. Of course, 'H₂O' is not a word. But the important point is that generic claims like 'Salt is sodium chloride' and 'Water is (di-)hydrogen oxide' seem *right*, and not just in the intuitive but statistically incorrect sense that 'Mosquitoes carry the West Nile virus' seems right; see Leslie (2007).

samples of Putnam's imagined substance XYZ. But why think that 'water' can *only* be used to access a concept that applies *only* to samples of H₂O, modulo impurities? As Putnam recognizes, and Chomsky often stresses (citing Aristotle), many nouns can be used to talk about things/stuff in terms of their functional roles and perhaps a paradigmatic appearance. Consider 'house', 'glue', 'table', etc. So why not think that 'water' is like these other nouns, in that it *can* be used to talk about any stuff that humans can use/perceive in certain ways, and polysemous in that 'water' *can also* be used to talk about stuff that has a certain essence?

Suppose that on another planet, Fraternal-Earth, doppelgangers of our scientists discover that all the stuff they call 'mud' has a common molecular structure. There is, in effect, a Platonic form of Frat-mud. Those scientists could use 'mud' to express a natural kind concept (NKC) and become expert at applying this concept. Other speakers might defer to these experts if mud becomes valuable (and counterfeited). But why think the Frat-Earth scientists, who grew up as our children do, *couldn't* use their word 'mud' to talk about our diverse samples of mud? Why think their word has a restrictive *meaning* just because their mud is especially uniform? Prima facie, the Frat-Earth scientists supplement their ordinary concepts with a new NKC. Similarly, we can have a NKC of water that is not the only concept we can access with our word 'water'.

2.2 Modulo Impurities

Chomsky highlights the functional dimensions of 'water', as opposed to 'H₂O', in another way that is closer to home. Even in the contexts Putnam had in mind, (1) and (23) are not equivalent.

(1) That is water.

(23) That is a sample of pure H₂O.

The stuff that comes from our taps is water, despite unintended impurities and intended fluoride. Rivers and lakes can be less pure.

This might seem like a small and uninteresting complication. But there are important differences between (1) and (22).

(22) That is a sample of H₂O, perhaps modulo some impurities.

Chomsky (1995) offers an illuminating example.

Suppose cup-1 is filled from the tap. It is a cup of water, but if a tea bag is dipped into it, that is no longer the case. It is now a cup of tea, something different.

Suppose cup-2 is filled from a tap connected to a reservoir in which tea has been dumped (say, as a new kind of purifier). What is in cup-2 is water, not tea, even if a chemist could not distinguish it from the present contents of cup-1....

In cup-2, the tea is an "impurity" in Putnam's sense, in cup-1, it is not, and we do not have water at all (except in the sense that milk is mostly water, or a person for that matter). If cup-3 contains pure H₂O into which a tea bag has been dipped, it is tea, not water, though it could have a higher concentration of H₂O molecules than what comes from the tap or is drawn from a river.

These observations can be bolstered with actual chemical analyses. I happen to have a well, whose output was analyzed by experts at "National Testing Laboratories, Ltd." According to their "Quality Water Analysis," the stuff from my well is about 99.4 percent H₂O. Ocean water of average salinity is about 96.5 percent H₂O. To put these numbers in context, the water contents of some familiar substances are listed below (data from the Department of Agriculture); where for these purposes, water contains only trace impurities.⁴

⁴ The US government does not often confirm Chomsky's claims. But see: ndb.nal.usda.gov/ndb/foods/show/4240 for club soda; ...4253 for non-cola diet soda; ...4337 for tea; ...4361 for diet cola; ...4287 and ...4288 for coffee and espresso; ...4159 and ...4156 for Michelob Ultra and Bud Light; ...283 for vinegar.

Club Soda	99.9
Diet soda, not cola	99.8
Tea	99.7
Diet Cola	99.54
<i>stuff from my well</i>	99.4
Coffee	99.39
Espresso	97.8
<i>ocean water, avg. salinity</i>	96.5
Michelob Ultra	95.4
Bud Light	95.0
Distilled vinegar	94.78

Diet Coke has a higher percentage of H₂O than the stuff from my well, which delivers stuff with the granularity of coffee. Tea is practically H₂O, as Chomsky suggested. Diet Sprite and Club Soda are even more like H₂O. An espresso has a higher percentage of H₂O than a typical sample of ocean water. Indeed, distilled vinegar and Bud Light are much closer to ocean water—by an H₂O modulo impurities test—than ocean water is to H₂O. One can say that some impurities are especially polluting. But why ignore fluoride, yet stress the neither-hydrogen-nor-oxygen components of Diet Sprite? As Chomsky observed, there are ordinary uses of ‘water’ such that what counts as water *in this human sense* is a complicated matter having to do with sources and intended purposes. There are also scientific uses of ‘water’ that abstract from these complexities, with the result that what counts as water *in this specialized sense* is clear and uncomplicated: samples of H₂O, allowing for some stipulated range of impurities that are acceptable for the purposes at hand. But the meaning of ‘water’ makes this *range of uses* possible. One can hypothesize that scientific uses determine the “real extensions” of ordinary words. But this hypothesis has consequences that seem absurd.

Again, Putnam did highlight an important point. We can use ‘water’ to access and express a scientific notion. But this does not define ‘water’. Competent speakers of English know that the following argument is not valid: water is H₂O; the water from my well has a high mineral content; so the H₂O from my well has a high mineral content. Likewise, the following argument is invalid: water is H₂O; so ‘water’ is true of an entity *e* if and only if *e* is a sample of H₂O. The premise is about water, while the conclusion is about a word. If the word is polysemous, then the premise is true only if the polysemy is resolved in favor of scientific usage. So even if the conclusion is true when restricted to scientific usage, it is a fallacy to drop the restriction and conclude that however the polysemy is resolved, ‘water’ is true of all and only samples of H₂O.

Chomsky can thus accommodate Putnam’s point about natural kind *uses* of ‘water’. Indeed, Putnam initially formulates his main claim in way that Chomsky could endorse: whatever meanings are, they are not “in the head” *and also* things that determine extensions. But Putnam then slides into assuming that meanings determine extensions.⁵ This reflects one use of

⁵ Putnam concedes that talk of words having extensions is a “very severe” idealization (p. 133). But he then argues 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. Putnam says that (i) is to be understood in terms of methodological solipsism: no psychological state “presupposes the existence any individual other than the subject to whom that state is described.” He then takes (ii) as a premise in arguing against (i). Yet he concludes (p. 165), “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 why think ‘water’ has an extension, much less one that is determined socially?

‘meaning’ by many philosophers and some linguists. But if ‘meaning’ is polysemous, then in the spirit of Putnam’s observation that NKC’s permit theoretical disagreements that can run deep, we should ask if there is a NKC of meaning—a concept that lets us think about some things that we call meanings, and think about them in a way that lets us offer theories of what these meanings *are*, even if it turns out that such theories abandon the idea that meanings determine extensions.

Prima facie, this is what Chomsky was doing by offering examples like (11),

(11) the woman saw the boy walking towards the railway station

and noting that structural homophony is a special kind of ambiguity. The idea was that a *scientific* concept of meaning would be a concept of certain mental representations that Human Languages connect with pronunciations, in ways that yield (a) interestingly constrained structural homophony, (b) relatively unconstrained lexical homophony, and (c) the vaguer though still interesting phenomena of lexical polysemy. I haven’t yet said what Human Languages *are*. I simply characterized them as the languages that children can acquire given ordinary experience. But Chomsky’s views about meanings are connected with an independently plausible conception of Human Languages as biologically implemented generative procedures.

2.3 ‘I’ before ‘E’

Like ‘meaning’, ‘language’ is polysemous. We speak of mathematical languages, bee languages, languages of thought, etc. So let’s adopt a generous conception of languages that covers anything that somehow connects interpretations of some kind with signals of some kind. Human Languages can be described as special cases that connect interpretations of a particular sort (meanings) with signals of a particular sort (pronunciations). This leaves room for many kinds of languages, and many proposals about the respects in which Human Languages are distinctive. But as stressed in section one, the languages that children naturally acquire connect meanings with pronunciations in ways that are unbounded and yet constrained.

This suggests that in acquiring a Human Language, a child acquires—i.e., comes to implement—a *procedure* that can *generate* expressions. As a simple example of a generative procedure, consider the operation of “adding one.” Given the number 1, this operation delivers the next number 1’ (a.k.a. 2); given 1’, applying the same operation delivers 1’’; and so on. There is a corresponding set of boundlessly many *pairs* of numbers: (1, 1’); (1’, 1’’); etc. We can’t list all these pairs. But we can think of them as the pairs determined by the operation “ $x + 1$,” thereby encapsulating an infinite list of pairs in a compact way. Of course, the pairs don’t determine a unique procedure. Consider the following complex process: subtract 1; then double the result; then add 4; then divide by 2. This yields the same results as adding one. There are many ways of pairing each of boundlessly many numbers with its successor. But if a finite mind performs this trick, it does so via some finitely specifiable procedure. This suggests a model for how a mind might connect meanings with pronunciations in an open-ended way.

Chomsky (1986) introduces the technical notion of I-language to talk about generative procedures, with ‘I-’ connoting ‘intensional’ (and ‘internal’). Given an I-language that generates certain meaning-pronunciation pairs, one can talk about those pairs; and one might call the set of those pairs an E-language, with ‘E-’ connoting ‘extensional’.⁶ But when a child acquires a Human Language, she acquires an I-language that connects unboundedly many meanings with pronunciations in accord with substantive constraints on homophony. We can introduce a corresponding notion of I-meanings to talk about the interpretations, whatever they are, that

⁶ Church (1941) distinguished functions-in-intension (procedures) from functions-in-extension (sets), stressing the primacy of the former. Lewis (1972) focused on sets of sentences.

Human I-languages connect with pronunciations. Using this technical notion, we can say that (11) has three I-meanings while (12) has only one.

(11) the woman saw the boy walking towards the railway station

(12) this is the railway station that the woman saw the boy walking towards

We can and should ask what I-meanings and their constituents are, just as chemists asked what hydrogen and oxygen are, and what it is for elements to combine in certain ratios. But these are scientific questions to be posed in the vocabulary of our best theories, not questions to be settled by “intuition” or hunches about how ideal expressions are related to extensions. One can *hypothesize* that I-meanings determine extensions. But if ‘water’ is one of the best cases for an affirmative answer, then there is little empirical support for this claim about I-meanings. To repeat, Chomsky thinks there are scientifically interesting notions in the vicinity of the ordinary words ‘language’ and ‘meaning’. The ordinary words are polysemous and otherwise unsuitable for scientific discourse. But we can develop NKC’s of I-languages and I-meanings. And on Putnam’s own view, such concepts may have extensions that are discovered by developing theories that abandon the idea that meanings determine extensions.

3. I-Meanings and Human Concepts

If the relevant targets of scientific inquiry are I-meanings, then instead of assuming that they determine extensions—and asking how this alleged explanandum could be true—we might more profitably ask how I-meanings are related to human concepts. The facts turn out to be interesting.

3.1 Polysemy Redux

In an essay for the New York Times (Nov 25, 2012), James Atlas wrote that there is “a good chance that New York City will sink beneath the sea.” But

...the city could move to another island, the way Torcello was moved to Venice, stone by stone, after the lagoon turned into a swamp and its citizens succumbed to a plague of malaria. The city managed to survive, if not where it had begun.

Chomsky (1995) offered similar examples in criticizing the idea that names have denotations. Quite typically, words offer conflicting perspectives.... We have no problem understanding a report in the daily press about the unfortunate town of Chelsea, which is “preparing to move”...with some residents opposed because “by moving the town, it will take the spirit out of it”, while others counter that “unless Chelsea moves, floods will eventually kill it”. There is a city called both “Jerusalem” and “al-Quds”, much as London is called “London” and “Londres”.... The government that claims it as its capital city has been considering plans to move al-Quds, while leaving Jerusalem in place.

He concludes this passage by drawing a philosophical moral.

The discussion would pose puzzles...if, failing to observe some of Wittgenstein’s good advice, we were to suppose that words like “London” or “Jerusalem” refer to things in the world in some public language, and were to try to sharpen meanings and ideas for conditions under which the presuppositions of normal use do not hold. (Chomsky 1995, p.??)

As an example of the puzzles that one *would* face if one held that city-names were logical constants that denote cities, consider the argument displayed below as (24).

(24) Torcello was moved to Venice.

Venice is a nice place.

Venice will be moved.

Torcello was moved to a nice place that will be moved.

Since this argument is not valid, it does not have a valid logical form like (25).

(25) MovedTo(t, v)
NicePlace(v)
WillBeMoved(v)
 $\exists x[\text{MovedTo}(t, x) \ \& \ \text{NicePlace}(x) \ \& \ \text{WillBeMoved}(x)]$

Likewise, since (26) is not valid, it does not have a valid logical form like (27).

(26) France is hexagonal. (27) H(f)
France is a republic. R(f)
There is a hexagonal republic. $\exists x[\text{H}(x) \ \& \ \text{R}(x)]$

But there are no *paradoxes* here. The point is that proper nouns like ‘Venice’ and ‘France’ are polysemous, and in this respect like ‘book’, which can be used to talk about spatially located things or more abstract intentionally characterized things. Trouble attends the hypothesis that words have extensions/denotations. And we can decline invitations to posit the ever more subtle logical forms required to keep the trouble at bay; see Pietroski (forthcoming) for related discussion of (28), with ‘Linus’ used as a name for (28).

(28) Linus is not true.

Chomsky does not, however, conclude that there are no systematically related meanings to study. He denies that Human Languages generate sentences that have compositionally determined truth conditions; cp. Davidson (1984), Lewis (1972). But he embraces the idea that expressions of a Human Language have (I)-meanings that are composable in interestingly constrained ways.

[W]e cannot assume that statements (let alone sentences) have truth conditions. At most, they have something more complex: ‘truth indications’, in some sense. The issue is not ‘open texture’ or ‘family resemblance’ in the Wittgensteinian sense. Nor does the conclusion lend any weight to the belief that semantics is ‘holistic’ in the Quinean sense that semantic properties are assigned to the whole array of words, not to each individually. Each of these familiar pictures of the nature of meaning seems partially correct, but only partially. There is good evidence that words have intrinsic properties of sound, form, and meaning; but also open texture, which allows their meanings to be extended and sharpened in certain ways; and also holistic properties that allow some mutual adjustment. The intrinsic properties suffice to establish certain formal relations among expressions, interpreted as rhyme, entailment, and in other ways by the performance systems (Chomsky 1996, p. 52)

This combination of skepticism about extant proposals, combined with a guarded optimism about the prospects for doing better—if we reject troublemaking assumptions that are empirically unmotivated—permeates the essays in Chomsky (1977).

We can make true claims with sentences like (29) and (30).

(29) Beavers are mammals.

(30) Beavers build dams.

Prima facie, asserting (29) implies that *all* beavers are mammals, perhaps apart from a few *recherché* cases. But asserting (30) does not carry a correspondingly strong commitment. We can use (30) to say, roughly, that dam-building is characteristic of healthy beavers that have access to the relevant materials. Though as Chomsky notes, (31) is importantly different.

(31) Dams are built by beavers.

A typical claim made with this passive sentence would imply that all, or nearly all, dams are

built by beavers. In this respect, (31) is akin to (29). Yet (32) and (33) are equivalent.

(32) Beavers built this dam.

(33) This dam was built by beavers.

So the I-meanings of (29-33) are interestingly related, in ways that theorists can try to describe and explain. But these sentences may not have truth conditions that await specification.

Contextualized uses of sentences, in acts of assertion or judgment, can be true or false; cp. Strawson (1950). But an action can be right or wrong without having a rightness condition; And even if particular acts of using (29-33) have truth conditions, these truth conditions may not be systematically related in the ways that the sentential I-meanings are.

Similarly, claims made with (34) and (35) differ in kind.

(34) Poems are written by fools like me.

(35) Mountains are climbed by fools like me.

Asserting (34) implies that all poems are written. But asserting (35) does not imply that all mountains are climbed. In each case, the truth-condition can be indicated roughly as follows: all the NOUNs that are VERBED are VERBED by fools like the speaker. Since all poems are written (or at least created intentionally), but not all mountains are climbed, the difference in the *claims* can be coherently described. But for those who think that ‘write’ and ‘climb’ have extensions, the challenge is to show how any alleged truth-conditions for *sentences* (34) and (35) are compositionally determined via some algorithm that doesn’t mischaracterize (29-33).

One can always speculate that particular examples are complex in ways that have not yet been understood. But Chomsky’s point is not merely that some relatively short sentences present difficulties for the hypothesis that the I-meanings of sentences determine truth conditions. Rather, it seems that few if any expressions of a Human Language fit the model of words having extensions that determine the truth conditions of sentences.

3.2 Assembly Instructions

That said, we do sometimes use some words to express concepts that approximate ideal concepts that do have extensions, at least relative to contexts. For some purposes, it can be useful to imagine languages whose expressions connect such concepts with pronunciations in one-to-one fashion. These expressions would not offer conflicting perspectives. And for just this reason, they may be poor models of human linguistic expressions. But we want some account of how words that allow for conflicting perspectives can also allow for “natural kind uses.” Moreover, absent a detailed alternative to the accounts of linguistic meaning that Chomsky criticizes, one might worry that alternatives will be worse. So let end by sketching a conception of I-meanings in the spirit of Chomsky’s remarks, his examples, and the general goal of describing meanings without (i) relying on *commonsense notions* of judgment and reference, or (ii) mischaracterizing the phenomena by using inappropriate *technical notions* of truth and denotation.

Suppose that in the course of acquiring the noun ‘book’, a child connects this noun with a pair of concepts: one that can be used to talk about certain inscribable *contents*, and one that can be used to talk about *inscriptions* of certain contents. If (copies of) both concepts are stored at the same “lexical address,” then instead of identifying the meaning of ‘book’ with any particular concept, we can say that this lexical meaning—call it μ (‘book’)—is an instruction for how to access a concept from the relevant address, which is also connected to a certain pronunciation, π (‘book’). Indeed, one might think of μ (‘book’) as an instruction of the following form: `fetch@ADDRESS`. If there are two or more concepts at the address in question, then the instruction can be executed in more than one way.

If ‘green’ is also a pronounceable instruction of the same form, then combining ‘green’ with ‘book’ to form phrase could be an instruction for how to form a complex concept by joining—perhaps by simply conjoining—two concepts fetched from the two lexical addresses. On this view, the meaning of ‘green book’ is a recipe of the following sort: fetch a concept from the ‘green’-address; fetch a concept from the ‘book’-address; and conjoin the fetched concepts. There might be only one way to execute this instruction, since the concept of book-contents may not be conjoinable with any concept at the ‘green’-address. That is, modifying a noun with an adjective may call for a combinatorial operation that applies to boundlessly many pairs of concepts that can be “fit together” in a certain conjunctive way, but only within certain constraints. In this respect, ‘green book’ may be like ‘book which he defaced’; cp. ‘book which he plagiarized’. But ‘green book that he plagiarized’ is not meaningless. We know what kind of concept is being called for; and we can use ‘which’—the head of the relative clause—to the exploit the dual perspective that ‘book’ provides.

Similarly, while (36) and (37) both seem fine, (38) is anomalous in a way that (39) is not.

(36) France is hexagonal.

(37) France is a republic.

(38) France is a hexagonal republic.

(39) France is hexagonal, and it is republic.

Yet (38) is still comprehensible, as is (40); see Chomsky (1965).

(40) Colorless green ideas sleep furiously.

It is as if ‘hexagonal republic’ is a recipe, but not one that we can use to make a good concept. By contrast, each sentential constituent of (39) is fine, as if we can use ‘it’ to access a concept of France that is *not* accessed via the occurrence of ‘France’ in (39).

The details are as complicated as they are. Combining ‘heard’ with ‘Fido’ to form a phrase calls for some operation(s) other than conjunction. Likewise for combining ‘Fido’ with ‘bark’, ‘heard’ with ‘Fido bark’, combining the lexical root ‘fish’ with a count morpheme, etc. But the hard questions, posed regularly by semanticists, are not made easier by supposing that words have extensions. To repeat an earlier point, one can talk about operations that map sets onto sets. But it’s hard to see how grammatical modes of combination could be suitably related to such operations if not via psychological correlates that can be invoked by theories that describe meanings as instructions for how to assemble concepts; see Pietroski (2010, 2011).

We can imagine a mind that connects each of its lexical items with exactly one concept, and each mode of grammatical combination with exactly one combinatorial operation, chosen from a stock of ideal concepts and operations. In this very special case, it might seem that the corresponding expression meanings are either ideal concepts or their extensions. Such a mind might connect ‘water’ with a concept of H₂O. But even if we regard such thinkers as idealized versions of ourselves, crucial empirical questions remain. Could the idealized minds just as well—and perhaps more naturally—connect ‘water’ with several concepts, including one that applies to the stuff from my well, but not Diet Coke, and perhaps the watery stuff on Twin Earth? Or are we cognitively noisy simulacra of ideal beings who reside only in Plato’s heaven, yet somehow determine that our words are like theirs?

For purposes of scientific inquiry, Chomsky urges us to characterize Human Languages as biologically implemented I-languages (generative procedures) that children can acquire given ordinary experience. Perhaps ideal thinkers could use these I-languages in ways that approximate the conjecture that I-meanings determine extensions—*ceteris paribus*, allowing for grains of salt while disallowing other impurities. But we actually use our I-languages in more interesting ways.

References

- Austin, J. (1961). *Philosophical Papers*. Oxford: Oxford University Press.
- (1962). *How to Do Things with Words*. Oxford: Oxford University Press.
- Bale, A. & Barner, D. (2009). The interpretation of functional heads: Exploring the mass/count distinction. *Journal of Semantics* 26(3): 217-252.
- Borer, H. (2005). *Structuring Sense* (volumes I and II) Oxford: Oxford University Press.
- Chomsky, N. (1957). *Syntactic Structures*. The Hague: Mouton.
- (1964). *Current Issues in Linguistic Theory*. The Hague: Mouton.
- (1965). *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
- (1977). *Essays on Form and Interpretation*. New York: North Holland.
- (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.
- Cartwright, H. (1963). Classes, Quantities, and Non-singular Reference.
Dissertation, Univ. of Michigan.
- Church, A. (1941). *The Calculi of Lambda Conversion*. Princeton: Princeton University Press.
- Davidson, D. (1984). *Essays on Truth and Interpretation*. Oxford: Oxford University Press.
- Evans, G. (1982). *Varieties of Reference*. Oxford: Oxford University Press.
- Fodor, J. (1975). *The Language of Thought*. New York: Crowell.
- (1986). *Psychosemantics*. Cambridge, MA: MIT.
- (1998). *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.
- (2003). *Hume Variations*. Oxford: OUP.
- Gillon, B. (1987). Readings of Plural Noun Phrases in English. *Linguistics and Philosophy* 102:199-219.
- Halle, M. & Marantz, A. (1993). Distributed Morphology and the Pieces of Inflection.
In *The View from Building 20*, ed. K. Hale and S. J. Keyser. Cambridge, MA: MIT Press.
- Harley, H. & Noyer, R. (2003). Distributed Morphology. In L. Cheng and R. Sybesma, eds.,
The Second Glot International State-of-the-Article Book, Berlin: Mouton de Gruyter.
- Keil, F. (1992) *Concepts, Kinds, and Cognitive Development*. Cambridge, MA: MIT Press.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Kuhn, T. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Hornstein, N. (1986). *Logic as Grammar*. Cambridge, MA: MIT Press.
- Laurence, S. & Margolis, E. (1999). *Concepts: Core Readings*.
- Leslie, S. (2007). Generics and the Structure of the Mind. *Philosophical Perspectives* 21:375-403.
- Lewis, D. (1972). General Semantics. In D. Davidson and G. Harman (eds.),
Semantics of Natural Language (Dordrecht: Reidel).
- Link, G. (1983). The Logical Analysis of Plurals and Mass Terms: A Lattice-Theoretic Approach. In R. Bäuerle, et.al., eds., *Meaning, Use, and Interpretation of Language* (Berlin: deGruyter).
- McGilvray, J. (1998). Meanings Are Syntactically Individuated and Found in the Head.
Mind and Language 13: 225-80.
- (2014). *Chomsky*, 2nd edition, Polity.
- (this volume). Meaning and Creativity.

- Pietroski, P. (2005). Meaning Before Truth. In *Contextualism in Philosophy* (G. Preyer and G. Peters, eds.). Oxford: Oxford University Press.
- (2010). Concepts, Meanings, and Truth: First Nature, Second Nature and Hard Work. *Mind and Language* 25: 247-78.
- (2011). Minimal Semantic Instructions. In C. Boeckx (ed.), *The Oxford Handbook of Linguistic Minimalism* (Oxford: Oxford Univ. Press).
- (forthcoming). I-Languages and T-sentences. In B. Armour-Garb (ed.), *The Relevance of the Liar* (Oxford: Oxford University Press).
- Putnam, H. (1962). It Ain't Necessarily So. *Journal of Philosophy* 59:658-671.
- (1975). The Meaning of 'Meaning'. In K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis: University of Minnesota Press).
- Quine, W. (1960). *Word and Object* (Cambridge, Mass.: MIT Press).
- Stainton, R. (2008). Meaning and Reference: Some Chomskyan Themes. In E. Lepore and B. Smith (eds.), *Handbook of Philosophy of Language* (Oxford: Oxford University Press).
- Strawson, P. (1950). On referring. *Mind*, 59(235): 320–344.
- Wittgenstein, G. (1953). *Philosophical Investigations*. New York: Macmillan.