Describing I-junction

Paul M. Pietroski, Univ. of Maryland

The meaning of a noun phrase like ‘brown cow’, or ‘cow that ate grass’, is somehow conjunctive. But conjunctive in what sense? Are the meanings of other phrases—e.g, ‘ate quickly’, ‘ate grass’, and ‘at noon’—similarly conjunctive? I suggest a possible answer, in the context of a broader conception of natural language semantics. But my main aim is to highlight some underdiscussed questions and some implications of our ignorance.

1. Varieties of Conjunction

The ampersand of a propositional calculus, as in (1), can be characterized via *truth* tables.

(1) P & Q

In a first-order predicate calculus that includes open sentences like (2-5),

(2) Bx **&** Cx (3) Bx **&** Cy

(4) Axy **&** Gx (5) Fxy **&** Uzw

the ampersand can be characterized in terms of Tarski’s notion of *satisfaction* by sequences (assignments of values to variables), with sentences like (1) as special cases. One can also imagine a more restricted language that does not generate open sentences like (3-5), yet does generate predicates like (6), with each open “slot” linked to the other.

(6) B(_) $\underline{\&}$ C(_)

The ampersand of such a language might be characterized in terms of intersection. Another familiar option involves appeal to functions, types, and truth *values* as in (7);

(7) $^{\langle eT, \langle eT, eT\rangle\rangle}\&(\lambda x.\text{T iff Bx}, \lambda x.\text{T iff Cx})$

where (7) describes the same function in extension as (8), which has (2) as a constituent.

(8) λx.T iff Bx **&** Cx

There are many ampersands. Of course, there are also many conceptions of “logical form;” and some notational variation is just that. For some purposes, (2), (6), (7), and (8) provide equivalent idealized ways of representing the meaning of ‘brown cow’. But the bold ampersand has a striking property illustrated with (3) and (5): it can be used to form an open sentence that has *more* free variables than any conjunct. And prima facie, combining expressions in a human language—a language that human children can naturally acquire in conditions typical for our species—cannot *increase* semantic adicity in this way. Note that ‘brown cow’ is not ambiguous, with (3) as a potential reading. Likewise, ‘from under’ cannot be understood as an open sentence like (5). So if one describes the meaning of ‘brown cow’ with the bold ampersand of (2), one needs to say more about why conjunctive meanings are so limited in human languages.

Many things might be said, in this regard, and I won’t try to rebut them here. My initial point is simply that the bold ampersand—unlike the conjoiners of (1), (6), and (7)—allows for considerable freedom: ‘**&**’ can link sentences that have arbitrarily many free variables, and any number of shared variables. Human linguistic meanings may not be well characterized with a conjunction operator that affords such freedom. The bold ampersand may instead let theorists describe a space of logically possible construals, for spoken expressions, and then discover that human linguistic meanings constitute a proper subset of that larger space. Indeed, given Chomsky-style “poverty of stimulus” considerations—for review, see Berwick et. al. (2011) and references there—it should be unsurprising if human languages employ a formally constrained operation of conjunction.

From a certain perspective, this talk of human languages "employing" operations might seem misplaced. If one takes languages to be sets of formula-interpretation pairs, then one might not care about how the relevant interpretations are *specified*. And if these interpretations are individuated extensionally, then the distinction between (7) and (8)

(7) $^{<eT,\, <eT,\, eT>>}\&(\lambda x.\mathrm{T\ iff\ } Bx,\ \lambda x.\mathrm{T\ iff\ } Cx)$

(8) $\lambda x.\mathrm{T\ iff\ } Bx\ \mathbf{\&}\ Cx$

might be merely notational. Likewise, given a simple object language that has monadic but no polyadic predicates, using a *meta*language that generates (2-5) might be harmless,

(2) Bx **&** Cx (3) Bx **&** Cy

(4) Axy **&** Gy (5) Fxy **&** Uzw

so long as one never uses the "excess" expressive power of the metalanguage to ascribe meanings to expressions of the object language. In that case, using (2) instead of (6)

(6) $B(\_)\ \underline{\&}\ C(\_)$

might be justified on grounds of notational convenience, or as a way of describing the object language as a set of <formula, interpretation> pairs that respect certain constraints that can be formulated in the metalanguage. But from another perspective, any such description of a human language is fundamentally misguided.

When a child acquires a human language, she acquires a biologically implemented *procedure* that generates unboundedly many expressions. For some purposes, one can adopt the idealization that each such procedure determines an infinite set of expressions. And for some purposes, one can describe each such procedure in terms of a function that maps phonological inputs (PHONs) to semantic outputs (SEMs)—or semantic inputs to phonological outputs. So following Chomsky (1986, 1995), one can apply Church's (1941) distinction between functions-in-intension and functions-in-extension: human I-languages are natural acquirable *procedures* that map PHONs to SEMs; human E-languages, if such there be, are the corresponding *sets* of PHON-SEM pairs.

Frege (1892) contrasted Functions, as procedures, with Courses-of-Values. Church located this intuitive distinction in the context of a post-Turing conception of computation. We have since grown used to thinking of functions primarily as sets of ordered pairs, and using lambdas so that identity claims like (9) are true.

(9) $\lambda x.|x - 1| = \lambda x.^{+}\sqrt{(x^2 - 2x + 1)}$

But when Church invented the lambda calculus, he wanted to describe the space of *computable* functions. So he wanted a notation that could be used to describe procedures in a way that would make (9) false but (10) true.

(10) $\mathrm{Extension}[\lambda x.|x - 1|] = \mathrm{Extension}[\lambda x.^{+}\sqrt{(x^2 - 2x + 1)}]$

Given arithmetic procedures, one can talk about corresponding sets of ordered pairs. But if one is concerned with computation, one need not talk about "computable sets."

Likewise, if the goal is to describe the space of naturally acquirable generative grammars (as opposed to the corresponding space of generated sets), one needs something like Chomsky's procedural notion of an I-language. Reference to E-languages *may* have utility for certain limited purposes. But if children acquire E-languages, they presumably do so *by* acquiring I-languages. Simplifying greatly, one can also think of an I-language as itself a possible output of a biologically implemented procedure ("Universal Grammar") for acquiring procedures (for paring PHONs with SEMs), given a normal course of experience.

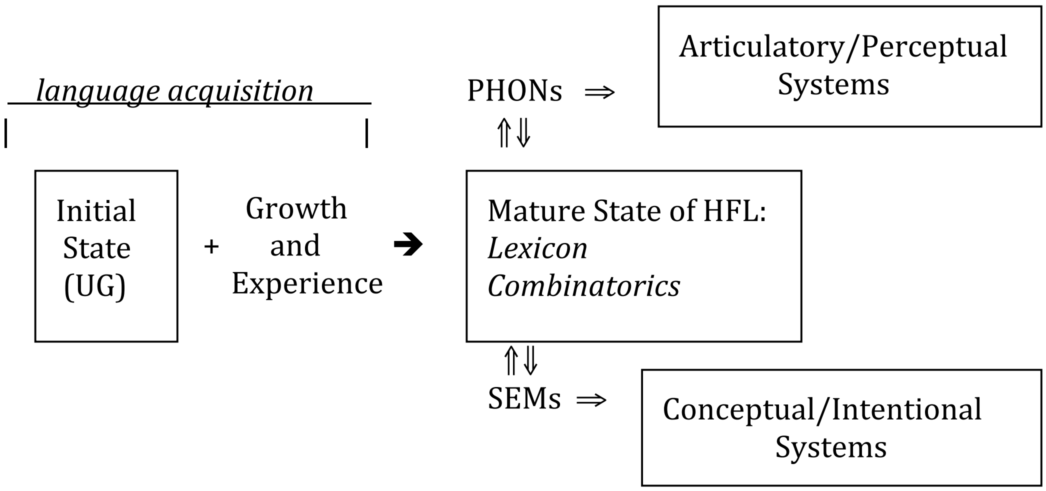

On any such view, PHONs and SEMs are posited as generable by naturally acquirable, biologically implemented procedures. This does not preclude externalist (or extensional) conceptions of interpretations for human linguistic expressions. One can, for example, hypothesize that SEMs *are* representations of Tarski-style satisfaction conditions; see Higginbotham (1985), Larson and Segal (1995), Ludlow (2011). In my view, this hypothesis is false; see Pietroski (2003, 2005, 2010). I suspect that SEMs are internalistically individuated "instructions" for how to assemble concepts by (i) fetching atomic concepts from certain lexical addresses, and (ii) employing certain combinatorial operations. But for now, let's remain agnostic about how SEMs are individuated. Let's even allow for the possibility of molecular duplicates generating different SEMs by virtue of inhabiting different environments. Perhaps in a world with cow look-alikes that are not cows, my I-language would be different, at least with regard to the SEM of 'cow'; and maybe there are worlds in which a molecular duplicate of me pairs the PHON of 'cow' with a different SEM. My point here is orthogonal to these issues (or perhaps non-issues).

For whatever one says about the meanings of 'brown' and 'cow', and whether the significance of combining these expressions is externalistic, one faces questions about the *form* of conjunction exhibited by 'brown cow'—at least given an I-language perspective. If the SEM of 'brown cow' is an instruction whose execution leads to the assembly of a conjunctive concept, whose conjuncts were fetched via 'brown' and 'cow', then one central question is easily posed: what form does the concept of conjunction exhibit? Is it like the bold ampersand of (2), in being able to freely link polyadic predicates?

(2) Bx **&** Cx

Or does executing the SEM of 'brown cow' result in a conjunctive concept like (6),

(6) B(_) <u>&</u> C(_)

whose conjoiner is limited to monadic conjuncts? Or is the relevant conjunctive concept more formally permissive than '<u>&</u>', yet more restrictive than '**&**'?

Put another way, when 'brown' is syntactically joined with—i.e. adjoined to—'cow', how is the conjunctive significance of the (ad)junction represented? Perhaps speakers vary in this I-language respect. But each speaker presumably represents the conjunction somehow: flexibly as in (2), inflexibly as in (6), or in some other way.

For some purposes, these details may not matter. One might pursue the project of saying what knowledge would *suffice* for representing the truth conditions—i.e., conditional assignments of truth/falsity—of the boundlessly many declarative sentences of English, without commitment to how (or if) actual speakers represent these truth conditions; cp Davidson (1967a). In terms of Marr's (1982) distinction between levels of explanation, one might try to describe an abstract function-in-extension that maps syntactic structures to compositionally representable Tarski-style satisfaction conditions, without yet offering any hypothesis about how this mapping is computed. But failure to advance a hypothesis about how the alleged truth conditions are represented, or how the alleged mapping is computed, is not an argument that meanings are representation-neutral; see, e.g., Evans (1981), Peacocke (1985), Davies (1987).

The point is obvious with regard to overt lexical items. If one refuses to countenance mental representations, one might insist that 'rabbit' and 'undetached-rabbit-part' are synonymous if extensionally equivalent across all possible situations; cp. Quine (1960). But given mental representations, one can distinguish extensionally equivalent semantic hypotheses. And recent psycholinguistic studies suggest that at least for lexical items like 'most', which arguably belongs to the logical vocabulary, arithmetically equivalent specifications of satisfaction conditions differ with regard to their plausibility as descriptions of linguistic meanings; see Hackl (2009), Pietroski *et.al*. (2010), Lidz *et.al*. (2011). Consider, for illustration, three candidate analyses of (11).

(11) Most of the dots red.
(11a) $\#[\text{DOT}(\text{X}) \,\&\, \text{RED}(\text{X})] > \#[\text{DOT}(\text{X}) \,\&\, {\sim}\text{RED}(\text{X})]$
(11b) $\#[\text{DOT}(\text{X}) \,\&\, \text{RED}(\text{X})] > \#[\text{DOT}(\text{X})]/2$
(11c) $\#[\text{DOT}(\text{X}) \,\&\, \text{RED}(\text{X})] > \#[\text{DOT}(\text{X})] - \#[\text{DOT}(\text{X}) \,\&\, \text{RED}(\text{X})]$

Provably, each candidate will agree with the others with regard to truth/falsity. But read as procedures (or algorithms) for determining truth/falsity, they differ: each appeals to at least one operation not appealed to by the others; cp. Gallistel & King (2009) on fundamental cognitive operations. Lidz *et.al*. (2011) argue that (11c) best reflects human semantic competence. And while they may be wrong about the details (cp. Hackl 2009), the proposal is perfectly intelligible: 'most' accesses a concept that includes a representation of cardinality-subtraction, as opposed to predicate-negation, or division by 2. But just as one can ask how the nonconjunctive aspects of (11) are represented, so one can ask how the conjunctive aspects are represented. Again, there are many ampersands.

To be sure, one can stipulate of a notion of meaning* such that the meaning* of (11) is captured equally well—and so equally badly—with each of (11a-c). But that is irrelevant if the question concerns the mental representations accessed by speakers who understand (11) in the way(s) that speakers of English do. Likewise, one can stipulate that (2) and (6)

(2) Bx **&** Cx (6) B(_) $\underline{\&}$ C(_)

capture the meaning* of 'brown cow' equally well. But that is irrelevant if the question concerns the conjunctive representation accessed by speakers of a human language.

2. Minimal Dyadicity

The underlined ampersand of (6) only permits monadic conjuncts and conjunctions. If this were the only conjunctive operation that human I-languages can invoke, that would explain why 'brown cow' and 'from under' cannot have polyadic interpretations like (3) and (5).

(3) Bx **&** Cy (5) Fxy **&** Uzw

But verb phrases like ‘ate grass’ and ‘ate at noon’ also seem to be conjunctive, at least in the sense that each is more restrictive than ‘ate’, much as ‘brown cow’ is more restrictive than ‘cow’. And for many reasons, including the availability of passive constructions like (12),

(12) Grass was eaten at noon.

it has become common to represent the meanings of phrases like ‘ate grass at noon’ as in (13) or (13’);

(13) PAST(E) & ∃P[EAT(E, P) & GRASS(P)] & ∃P[AT(E, P) & NOON(P)]

(13’) PAST(E) & EAT(E) & ∃P[THEME(E, P) & GRASS(P)] & ∃P[AT(E, P) & NOON(P)]

where PAST(E) and EAT(E) are concepts of events, and the dyadic concepts relate events to participants.

The difference between (13) and (13’) concerns whether the variable corresponding to ‘grass’ is *separated* from the concept corresponding to the verb; cp. Schein (1993, 2002), Kratzer (1998). But for present purposes, this is a small difference concerning *which* dyadic concept is conjoined with the monadic GRASS(P). The crucial contrast is with EAT(E, X, Y), which applies to triples consisting of an event of eating, an eater, and a thing eaten; cp Davidson (1967b). The assumption here is that at least the variable corresponding to the *subject* in (14) is separated from the concept corresponding to the verb, as in (15).

(14) cows eat grass

(15) ∃P[EAT(E, P) & GRASS(P)] & ∃P[AGENT(E, P) & COWS(P)]

And note that (14) can be embedded as in (16), which can be analyzed as in (17),

(16) hear cows eat grass

(17) ∃P{HEAR(E, P) & ∃P’[EAT(P, P’) & GRASS(P’)] & ∃P’[AGENT(P, P’) & COWS(P’)]}

with an event as the “thing” heard; cp. Higginbotham (1983). This captures the ambiguity of (18), since either the hearing or the eating could be in the barn.

(18) I could hear cows eating grass in the barn.

Again, we need to ask what kind of conjunction is really being posited in (13, 15, 17). But if any such “event analyses” capture how speakers understand verb phrases, then we need to posit a concept of conjunction that is at least minimally more permissive than the underlined ampersand of (6).

(6) B(_) <u>&</u> C(_)

So let ‘^’ indicate a conjoiner that can also link a single dyadic concept to a single monadic concept, subject to the following constraint: the second variable of the dyadic concept must be linked to (the variable of) the monadic concept, and *immediately* existentially closed, yielding a complex monadic concept. This permits (2a) and (4a),

(2a) BROWN(X)^COW(X) (4a) ∃Y[ABOVE(X, Y)^COW(Y)]

but not (4b*), (4c*), (4d*), or (5*).

(4b*) ∃X[ABOVE(X, Y)^COW(X)] (4c*) ∃Y[ABOVE(X, Y)^COW(X)]

(4d*) ∃X[ABOVE(X, Y)^COW(Y)] (5*) FROM(X, Y)^UNDER(X, Y)

Given the posited constraint on the new conjoiner, (13) can be rewritten as (19),

(13) PAST(E) & ∃P[EAT(E, P) & GRASS(P)] & ∃X[AT(E, P) & NOON(P)]

(19) PAST(_)^∃[EAT(_, _)^GRASS(_)]^∃[AT(_, _)^NOON(_)]

with the existential quantifier always binding the *monadic* concept in its scope, and the second variable of any *dyadic* concept always linked to the *adjacent* monadic concept. In this sense, there is no “choice” with regard to which variable ‘∃’ binds. Despite the absence of variable letters, (19) is unambiguously a concept of past events of eating grass at noon.

Verbs and prepositions clearly introduce a kind of (at least thematic) dyadicity. But even with regard to 'brown cow', one might think that the relevant conjunction involves a little more than monadicity. For (6) fails to capture the idea that 'brown cow' applies to things that are cows and *brown-for-cows*. The point is clearer with examples like 'big ant', 'talented musician', or 'ate quickly'. But as many authors have noted, it is far from obvious that 'brown' itself applies equally to brown cows, brown houses, fruit that went brown, and so on. So perhaps human I-language conjunction regularly allows for one dyadic concept, as in (20), which could be rewritten as (21);

(20) ∃X[BROWN-ONE(x, X) & COW(X)]

(21) ∃[BROWN-ONE(_, _)^COW(_)]

where BROWN-ONE(x , X) is a formally dyadic concept that applies to a thing, x, and one or more things, the Xs, iff x is both one of the Xs and a brown one.

Developing this idea in detail would require a digression into Boolos' (1998) interpretation of the *second-order* monadic predicate calculus, in which a (capitalized) variable can have *many* values relative to one assignment of values to variables; see Pietroski (2005, 2006, 2011) for discussion. Such digression would have independent interest, not least because it would raise related questions about where human I-languages lie with respect to *two* dimensions of variation exhibited by familiar invented languages: those with only singular/first-order variables vs. those with plural/second-order variables; and those that are fundamentally monadic vs. those exhibit unrestricted relationality.

For present purposes, though, it is enough to note that human I-languages may invoke a concept of conjunction that permits only a minimal and "short lived" kind of dyadicity. This raises the empirical question of whether human I-languages invoke any concept of conjunction that is *more* permissive, like the bold ampersand of (2) and (3).

(2) Bx **&** Cx (3) Bx **&** Cy

Indeed, given the minimally nonmonadic conjoiner of (19),

(19) PAST(_)^∃[EAT(_, _)^GRASS(_)]^∃[AT(_, _)^NOON(_)]

one might wonder if human I-languages invoke *any other* form of composition.

3. Composition Constrains

As mentioned above, it is at least generally true that combining expressions of a human I-language does not *increase* the number of free variables, as in (3). It is often assumed that human I-languages employ a recursive "saturation" operation that *reduces* the number of free variables, on the model of (8) to (8a), (22) to (22a), and so on.

(8) λx.T iff Bx **&** Cx (8a) T iff Ba **&** Ca

(22) λy.λx.T iff Rxy (22a) λx.T iff Rxa

But this is not obvious, given the possibility of "thematic separation" as in (15),

(15) ∃P[EAT(E, P) & GRASS(P)] & ∃P[AGENT(E, P) & COWS(P)]

which can be rewritten as (23).

(23) ∃[EAT(_, _)^GRASS(_)]^∃[AGENT(_, _)^COWS(_)]

And while one can think of '∃' as reducing the adicity of AGENT(_, _)^COWS(_), this formally dyadic predicate is not saturatable; its second variable must be *closed* prior to any other combination operation. So one can think of ∃[AGENT(_, _)^COWS(_)] as a monadic concept that is a "thematic conversion" of COWS(_); while the latter concept applies to cows, the former applies to events done by cows. In this sense, monadicity is preserved.

Elsewhere, I have shown how a wide range of examples—including names, plural nouns, ditransitive verbs, quantificational determiners, as in (24)—

(24) Tyler gave every philosopher three dollars

can be accommodated in these terms; see Pietroski (2005, 2010, 2011). And there are at least three kinds of motivations for so accommodating, if possible.

First, various empirical considerations suggest that human I-languages abhor polyadicity. For example, the apparently triadic concept BETWEEN is expressed with constructions like (25), and not with a predicate that takes three arguments as in (25a).

(25) Bob is between Al and Carl.

(25a) *Bob [[betweens Al] Carl]

But why should the circumlocution of (25) be required if (25b) can be a lexical meaning?

(25b) λz.λy.λx.T iff BETWEEN(X,Y,Z)

Such examples are, I suspect, ubiquitous.

(26) Al is taller than Bob. (26a) *Al talls Bob.

(27) Al is from Chicago. (27a) *Al froms Chicago.

(28) Al sold a car to Bob for a dollar.

(28a) *Al [[[sold a car] Bob] a dollar]

The need for such circumlocution—along with a general preference for subject-predicate structures (with subjects analyzed, in Aristotelian fashion, as containing predicates)—is unsurprising if relational *concepts* must be *expressed* via constituents that can be *conjoined* by a constrained operator like that of (23). Polyadicity would then have to be expressed via separated thematic conjuncts, with an existential closure for each variable beyond the first.

Second, it is independently plausible that human I-languages treat names like 'Al'—and deictic expressions like 'she'—as complex monadic predicates, perhaps with indices as constituents, *and not* as primitive expressions of type <e>. Examples like (29)

(29) Every Al at the party met the Bob that arrived with his wife.

are familiar, at least since Burge (1973), and commonplace in many languages.

Third, if only because human I-languages are *biologically implemented* procedures, one wants to describe human semantic competence in terms of the simplest combinatorial operations that accommodate the facts. It will, I suspect, be hard enough to figure out how biology can arrange for representations like (19).

(19) PAST(_)^∃[EAT(_, _)^GRASS(_)]^∃[AT(_, _)^NOON(_)]

By contrast, while (2) *looks* simpler,

(2) Bx **&** Cx

the bold ampersand might be too unconstrained to be implementable with the biology available to children (i.e., human language acquisition devices).

On any view, it must be possible to use a lexical item like 'cow' in combination with different variables, and likewise for 'brown', given expressions like (29).

(29) …brown cow that saw a brown dog that chased a yellow cow that …

Indeed, drawing on both Frege and Tarski, one might think of COW(_) as an "unsaturated" concept that can be used to form boundlessly many open sentences: COW(X); COW(X'); etc. But given polyadic concepts like BETWEEN(_, _, _) and TALLER(_, _), allowing replacement of each "slot" with any variable would quickly lead to an explosion of potential adicities and variable-binding for complex concepts: BETWEEN(X, Y, Z) **&** TALLER(X, Y); BETWEEN(X, Y, Z) **&** TALLER(Y, Z); BETWEEN(X, Y, Z) **&** TALLER(Z, W); BETWEEN(X, Y, Z) **&** TALLER(W, V); BETWEEN(X, Y, Z) **&**

TALLER(W, X); BETWEEN(X, Y, Z) **&** TALLER(Y, X); etc. And while this explosion may help logicians explore the space of coherent possibilities, it may also reflect a kind of representation system not found in nature, at least not prior to its explicit invention by theorists.

4. Some Implications and Speculations

It is well known that *describing* the meaning of a word (e.g., 'brown') is harder than *assigning* an interpretation to a symbol (e.g., 'B') as part of a model. The same point applies to the semantic contribution of *combining* 'brown' with 'cow'. It will take a lot of work to determine the meaning of this mode of composition. By contrast, it is easy to stipulate an interpretation for 'BC' in terms of interpretations assigned to 'B' and 'C'. In one sense, this is obvious, since one can invent a language in which concatenation signifies disjunction. But even *given* that combining 'brown' with 'cow' signifies a kind of conjunction, there is an important sense in which we have not yet begun to describe the meaning of 'brown cow'.

Until we can seriously defend a specific hypothesis about what the relevant operation of conjunction can conjoin—and I don't pretend to have done this here—how can we claim to have addressed the real questions concerning how the meaning of a complex expression is compositionally determined? And if we don't even know which kind of conjunction we understand 'brown cow' to be, then in my view, we do not know the first thing about how human linguistic meaning is related to truth and logic. (Imagine your response to a student who didn't how the ampersand worked, in some interpreted formal language, yet still claimed to know how truth was related to the interpretation function.)

Put another way, we don't know what kind of concept we express with 'brown cow'. I sketched a "minimalist" proposal, and I don't know of a better one. But that's a low standard, given how little attention has been devoted to the question. Philosophers and linguists have devoted a lot of energy to saying how an ideal logician might *recursively specify* plausible interpretations for human linguistic expressions, in ways that would allow specifications of many other interpretations that no human child could come up with. This has often led to insights with regard to characterizing an abstract function-in-extension that maps syntactic structures to compositionally representable Tarski-style satisfaction conditions, without yet offering any hypothesis about how this mapping is computed. In my view, we still know next to nothing about semantic *composition*. But at least the questions we face in this regard are tolerably clear, and there is an obvious research strategy: start with the dumbest, most restrictive conception of semantic composition that has a prayer of descriptive adequacy; find out where it really is inadequate; and adjust accordingly.

By contrast, various philosophical projects—including Davidson's (1967a) program for semantics—involve "regimenting" the concepts that we naturally express with (meaningful) expressions of a human I-language; where the preferred regimentation is often motivated by certain normative aims concerning truth and/or uses of language in a scientific description of the language-independent world. But if we still don't have a clear view about even the *form* of the concept we express with 'brown cow', the idea of regimenting such concepts strikes me as premature at best. And the idea of starting with a conception of how human linguistic meaning is related to truth strikes me as misguided.

Until we can defend a specific hypothesis about the sense(s) in which 'brown cow' and 'ate quickly' are conjunctive, and what the relevant operation(s) of conjunction can conjoin, I don't see how we can have a defensible view about the lexical meanings of 'brown' or 'cow' or 'ate'. I would have thought that views about how meaning is related to truth should be informed by discoveries concerning lexical meanings and how they compose, not the other way around. If that isn't provocation enough, let me end with a little more.

Dummett (1994) suggested that the distinctive contribution of 20th century analytic philosophy was to challenge the classical idea, indicated on the left,

World ⇑ Thought ⇐ Language	World ⇑ Thought ⇒ Language

that overt language is related to the world via inner thought. According to Dummett, analytic philosophy suggests a different picture and invites a distinctive methodology: thought is related to the world via overt language; so the primary task in semantics is to figure out how public language is related to the world, without mental intermediaries, and then say how inner thought is related to public language. Whatever one thinks of this idea, as history or policy, my sense is that many philosophers who would count themselves part of the Frege-Russell-Wittgenstein-Carnap tradition would reject Dummett's centerpiece of analytic philosophy. I would. But a desire for something like Dummett's centerpiece may still motivate the strategy of starting with assumptions about how human linguistic meaning is related truth, adopting a Tarskian conception of semantic composition (i.e., recursive specification of satisfaction conditions), and letting the chips fall where they may regarding the semantic contributions of words. But this methodological perspective, indicated on the right, seems bizarre if we are concerned with human I-languages.

Meaning/Truth ⇑ Words ⇐ Combinatorics	Meaning/Truth ⇓ Words ⇐ Combinatorics

If word meanings are composable in ways that human biology can implement, then perhaps a better way forward is to start with questions like the following: in what sense is 'brown cow' conjunctive, and is 'ate quickly' relevantly similar? Plausible answers will, I suspect, suggest severe constraints on the space of possible word meanings for humans. But if human lexical items are biologically constrained, in ways that dovetail with our limited modes of semantic composition, what are the odds that meaning is related to truth in any simple way? Perhaps the moral is that describing meanings is hard, and saying how meaning is related to normative notions is even harder.

References

Berwick, R., Pietroski, P., Yankama, B. and Chomsky, N. 2011: Poverty of Stimulus Revisited. *Cognitive Science* 35: 1207-42.

Boolos, G. 1998: *Logic, Logic, and Logic*. Cambridge, MA: Harvard University Press.

Burge, T. 1973: Reference and Proper Names. *Journal of Philosophy* 70:425-39.

Carlson, G. 1984: Thematic Roles and their Role in Semantic Interpretation. *Linguistics* 22:259-79.

Chomsky, N. 1986: *Knowledge of Language* (New York: Praeger).

—1995: *The Minimalist Program*. Cambridge, MA: MIT Press.

—2000b: *New Horizons in the Study of Language and Mind*. Cambridge: CUP.

Church, A. 1941. *The Calculi of Lambda Conversion*. Princeton: Princeton University Press.

Crain, S. and Pietroski, P. 2001: Nature, Nurture, and Universal Grammar. *Linguistics and Philosophy* 24: 139-86.

Davidson, D. 1967a: Truth and Meaning. *Synthese* 17:304-23.

—1967b: The Logical Form of Action Sentences. In The Logic of Decision and Action, ed. N. Rescher, Pittsburgh: University of Pittsburgh Press.

Davies, M. 1987: Tacit knowledge and semantic theory: Can a five per cent difference matter? *Mind* 96: 441–62.

Dummett, M. 1994: *The Origins of Analytic Philosophy*. Cambridge, MA: Harvard.

Evans, G. 1981: Semantic theory and tacit knowledge. In S. Holtzman and C. Leich (eds), *Wittgenstein: To Follow a Rule* (London: Routledge and Kegan Paul).

Frege, G. 1892: Function and Concept. In Geach and M. Black (trans.), *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell (1980).

Gallistel, C., & King, A. 2009: Memory and the computational brain. New York: Blackwell.

Hackl, M. 2009: On the Grammar and Processing of Proportional Quantifiers: Most versus More Than Half. *Natural Language Semantics.*

Heim, I. and Kratzer, A. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Higginbotham, J. 1983: The Logical Form of Perceptual Reports. *J. of Philosophy* 80:100-27.

— 1985: On Semantics. *Linguistic Inquiry* 16: 547-93.

Hornstein, N. and Pietroski, P. (2009): Basic Operations. *Catalan Journal of Linguistics* 8:113-39.

Horty, J. 2007: *Frege on Definitions: A Case Study of Semantic Content*. Oxford: OUP.

Larson, R. & Segal, G. 1995: *Knowledge of Meaning*. Cambridge, MA: MIT Press.

Lidz, J., Hunter, T., Pietroski, P., and Halberda J. (2011, in press). Interface Transparency and the Psychosemantics of Most. Natural Language Semantics.

Ludlow (2011). The Philosophy of Generative Grammar. New York: Oxford University Press.

Kratzer, A. 1996: Severing the External Argument from its Verb. In J. Rooryck and L. Zaring (eds.), Phrase Structure and the Lexicon. Dordrecht (Kluwer Academic Publishers).

Larson, R. 1988: forthcoming: *On Shell Structure*. Routledge.

Marr, D. 1982: *Vision.* WH Freeman and Company, New York NY.

Parsons, T.1970: Some problems concerning the logic of grammatical modifiers. *Synthese* 21:320-34.

—1990: *Events in the Semantics of English*. Cambridge, MA: MIT Press.

Partee, B. 2006: Do we need two basic types? In Hans-Martin Gaertner, et.al. (eds.) *40-60 puzzles for Manfred Krifka*, http://www.zas.gwz-berlin.de/40-60-puzzles-for-krifka.

Peacocke, C. 1986: Explanation in computational psychology: Language, perception and level 1.5. *Mind and Language* 1:101-123.

Pietroski, P. 2005a: *Events and Semantic Architecture*. Oxford: Oxford University Press.

—2005b: Meaning Before Truth. In G. Preyer and G. Peter (eds.), *Contextualism in Philosophy*

(Oxford: Oxford University Press).
—2006: Induction and Comparison. *Univ. of Maryland Working Papers in Linguistics* 15:157-190.
—2008: Minimalist Meaning, Internalist Interpretation. *Biolinguistics* 4: 317-41.
—2010: Concepts, Meanings, and Truth. *Mind and Language*
—2011: Minimal Semantic Instructions. *Oxford Handbook of Minimalism*.
Pietroski, P., Lidz, J., Hunter, T., and Halberda, J. 2009: The Meaning of 'Most': Semantics, Numerosity, and Psychology. *Mind and Language* 24:554-85.
Poeppel, D. and Embick, D. 2005: The relation between linguistics and neuroscience. In A. Cutler (ed.), *Twenty-First Century Psycholinguistics: Four Cornerstones*. Lawrence Erlbaum.
Preyer, G. and Peter, G., eds. 2002: *Logical Form and Language*. Oxford: OUP.
Putnam. H. 1975: The Meaning of 'Meaning'. In K. Gunderson (ed.), *Language, Mind and Knowledge* (University of Minnesota Press).
Quine, W. 1960: *Word and Object*. Cambridge, MA: MIT.
Schein, B. 1993: *Plurals and Events*. Cambridge, MA: MIT Press.
—2002: Events and the Semantic Content of Thematic Relations. In Preyer and Peter (2002).
—forthcoming. *Conjunction Reduction Redux*. Cambridge, MA: MIT Press.
Tarski, A. 1933: The Concept of Truth in Formalized Languages. Reprinted in *Logic, Semantics, Metamathematics* 2nd edition), tr. J.H. Woodger, ed. J. Corcoran (Indianapolis: Hackett).