Minimal Semantic Instructions
Paul M. Pietroski
University of Maryland

Chomsky's (1995, 2000a) Minimalist Program (MP) invites a perspective on semantics that is distinctive and attractive. In section one, I discuss a general idea that many theorists should find congenial: the spoken or signed languages that human children naturally acquire and use—henceforth, human languages—are biologically implemented procedures that generate expressions, whose meanings are recursively combinable *instructions to build concepts* that reflect a minimal interface between the Human Faculty of Language (HFL) and other cognitive systems. In sections two and three, I develop this picture in the spirit of MP, in part by asking how much of the standard Frege-Tarski apparatus is *needed* in order to provide adequate and illuminating descriptions of the "concept assembly instructions" that human languages can generate. I'll suggest that we can make do with relatively little, by treating all phrasal meanings as instructions to assemble number-neutral concepts that are monadic and conjunctive. But the goal is not to legislate what counts as minimal in semantics. Rather, by pursuing one line of Minimalist thought, I hope to show how such thinking can be fruitful.[1]

1. Procedural Considerations

For better and worse, we can use 'language' and 'meaning' to talk about many things. As an initial guide to the topic here, let's tentatively adopt two traditional ideas: languages, via their expressions, connect signals of some kind with interpretations of some kind; and expressions of a human language have meanings—semantic properties that are recognized when the expressions are understood. Following Chomsky, I take each human language to be a state of HFL that generates expressions that pair phonological structures (PHONs) with semantic structures (SEMs), via which HFL interfaces with other cognitive systems that let humans perceive/articulate linguistic signals and assemble/express corresponding interpretations.[2] While the signals are plausibly gestures or sounds, in some suitably abstract sense, I assume that the interpretations are composable mental representations that may be individuated externalistically. On this view, SEMs can be characterized as instructions to assemble concepts, and meanings can be identified with such instructions in the following sense: to have a meaning is to be a certain kind of instruction, and thus to have a certain "fulfillment" condition; and semantic theories for human languages are theories of the concept assembly instructions that HFL can generate. (Readers who find this banal may wish to skim ahead to section two.) This mentalistic perspective permits versions of Truth Conditional Semantics. But the idea is that central questions about meaning concern the concepts and composition operations invoked via SEMs.

1.1 I-languages and Interpretations

We need to distinguish generative procedures from generated products. So following Chomsky (1986), let's say that I-languages are *procedures that generate* expressions, while E-languages are *sets of* expressions; cp. Frege (1892) and Church (1941) on functions as *intensions* vs. *extensions*. As an analogy, note that '$|x - 1|$' and '$^+\sqrt{(x^2 - 2x + 1)}$' suggest different algorithms for determining a value given an argument, with 'x' ranging over whole numbers; yet each procedure determines the same set of argument-value pairs. We can use lambda-expressions to denote sets, and say that $\lambda x.|x - 1| = \lambda x.^+\sqrt{(x^2 - 2x + 1)}$. Or we can use such expressions to denote procedures, and say that $\lambda x.|x - 1| \neq \lambda x.^+\sqrt{(x^2 - 2x + 1)}$, adding that Extension$[\lambda x.|x - 1|] =$ Extension$[\lambda x.^+\sqrt{(x^2 - 2x + 1)}]$. But whatever our conventions, different algorithms can have the same input-output profile. Likewise, distinct I-languages can in principle generate the same

expressions. And in practice, speakers may implement distinct I-languages whose expressions associate signals with interpretations in ways that support workable communication.[3]

At least for purposes of studying the natural phenomena of human linguistic competence, including its acquisition and use, I-languages are importantly prior to E-languages. Each normal child acquires a language with unboundedly many expressions. So to even say *which* E-language a child allegedly acquires, one needs a generative procedure that specifies that set. And if a child acquires a set with unboundedly many elements, she presumably does so by acquiring (an implementation of) a procedure. Moreover, a biologically implemented procedure may not determine a set of expressions; but even if it does, there is no reason for taking this set to be an interesting object of study. Indeed, the acquired procedures may already lie at some remove from any stable target of scientific inquiry: the real generalizations may govern HFL, the faculty that lets humans acquire and use certain I-languages. But in any case, the theoretical task is not merely to specify the generable expressions that speakers can use. The task is to specify the expression-generating procedures that speakers implement.[4]

We begin, however, in ignorance. With regard to expressions of a human I-language (henceforth, "I-expressions"), we don't know what the relevant interpretations are, or how they relate to reference and communication. But if spoken I-expressions connect sounds with mind-independent things, they presumably do so via mental representations. And for present purposes, I take it as given that human infants and many other animals have *concepts* in a classical sense: mental representations that can be *combined* in ways that can be described in terms of conceptual *adicities*; see, e.g., Frege (1884, 1892) and Fodor (1975, 1986, 2003). So if only for simplicity, let's suppose that spoken I-expressions connect (representations of) sounds with composable concepts, allowing for concepts that are distinctively human.

Matters are hard enough, even with this assumption, in part because a single I-expression may be linked to more than concept, as suggested by the phenomenon of polysemy. But even setting aside examples like 'book'—which illustrates an abstract/concrete contrast that may distinguish kinds of concepts *and* kinds of things we can think about—it seems that a single lexical meaning can correspond to more than one concept. A speaker who knows that Venus is both the morning star and the evening star may have more than one concept of Venus, no one of which is linguistically privileged. Likewise, a speaker may have many ways of thinking about water. And as Chomsky (2000b) stresses, it is hardly obvious that some set is the extension of each 'water'-concept, given what competent speakers call 'water' when they are not doing science; cp. Putnam (1977). At a minimum, it would be rash to insist that each meaning privileges a single concept, or that concepts linked to a single meaning must share an extension. So let's say, tentatively, that each expression of a human I-language links a single phonological structure (PHON) to a single semantic structure (SEM); where each SEM determines (and perhaps just is) a meaning that need not determine a single concept.[5]

1.2 Meanings as Instructions

Chomsky describes PHONs and SEMs as instructions via which HFL interfaces with human articulatory/perceptual systems and conceptual/intentional systems. If we focus on comprehension, as opposed to production, words seem to invoke concepts that can be combined via operations invoked by phrasal syntax. So especially if a word can invoke different concepts on different occasions, one might describe each lexical SEM as an *instruction to fetch* a concept that meets a certain condition. Then a phrasal SEM can be characterized as an *instruction to*

combine, in a certain way, concepts fetched or assembled by executing the constituent SEMs. The interest of this claim lies with the details: which concepts and combination operations are invoked by I-expressions? And eventually, the instruction metaphor must be replaced with something better, perhaps via analogies to programming languages and compilers. But the idea is that SEMs are Janus-faced: they are grammatical objects, whose composition (from a lexicon of atomic expressions) can be described in terms of formal operations like concatenation and labeling; yet they can direct construction of concepts, whose composition can be described in terms of semantic operations like saturation or conjunction. Or put another way: SEMs are generated, hence they exhibit a syntax; but these expressions are also apt for use in concept construction, allowing for an overtly mentalistic/computational version of the idea that meanings are "directions for the use of expressions;" cp. Strawson (1950).

This leaves room for various conceptions of what these directions require. For example, one can hypothesize that 'brown cow' is (an I-expression whose SEM is) the following tripartite instruction: fetch a concept that applies to x iff x is a cow; fetch a concept that applies to x iff x is brown; and conjoin these concepts. This says nothing about where the concepts must come from. A speaker who links the words to suitable concepts, COW(X) and BROWN(X), might well conjoin those very concepts; but the instruction could be fulfilled by fetching any extensionally equivalent concepts. Or perhaps the instruction is to fetch a concept that applies to brown things, form a corresponding higher-order concept like &[BROWN(X), X(X)] and *saturate* it with a concept that applies to cows to obtain a concept like &[BROWN(X), COW(X)]. Fulfilling this instruction requires a certain process, culminating in the construction of a concept with a certain form. Or perhaps 'brown cow' calls for concepts from specific lexical addresses, but without imposing conditions on what the concepts apply to. Then twins might use the same I-expression to construct concepts that differ extensionally; although theorists can add that an I-language is an idiolect of English only it meets certain externalistic conditions.[6]

Thus, many theorists should be able to adopt the idea that HFL generates concept assembly instructions, and that part of the task in semantics is to describe the "I-concepts" that can be constructed by executing these instructions. Put another way, at least part of our job is to say which "I-operations" are invoked by phrasal syntax and what *kinds* of concepts can be combined via these operations. We should not assume, a priori, that all human concepts are combinable via I-operations. The best overall theory may be one according to which few if any of our "pre-lexical" concepts are combinable via the operations that I-expressions can invoke; see Pietroski (2010). But in any case, semanticists face a task that invites a Minimalist question: what is the sparest inventory of operations and conceptual types that allows for rough descriptive adequacy with regard to characterizing the concept assembly instructions that HFL can generate? And here, we need to consider not just the syntactic operations employed in generating SEMs, but also the conceptual operations employed in executing SEMs.

My specific suggestion, developed and contrasted with others below, has two main aspects. First, an open-class lexical SEM is an instruction to fetch a *monadic* concept that need not be the concept lexicalized. Second, a phrasal SEM is an instruction to build a *conjunctive* monadic concept via I-operations that are limited to (i) highly restricted forms of conjunction and existential closure, and (ii) a few ways of converting one monadic concept into another. These conversion operations presuppose (a) some thematic concepts, associated with prepositions or certain grammatical relations, and (b) an analog of Tarski's (1933) treatment of "closed"

3

sentences as satisfied by everything or nothing, along with a number-neutral version of his appeal to sequences and variants. This is still a lot to posit, since concept-construction has to be implemented biologically. Moreover, to even pose tractable implementation questions, we need theoretical notions of appropriate "granularity" (Poeppel and Embick [2005]); and while (i) *may* be in the right ballpark, at least some of (ii) seems worryingly grand. But I don't know how to make do with less—even ignoring lots of interesting details, in order to focus on highly idealized elementary constructions. And it is all too easy to posit far more: a richer typology of I-concepts, corresponding to abstracta like truth values and high-order functions; additional composition operations; type shifting, etc. But in the spirit of MP, we can try to formulate the sparest proposals that have a prayer of descriptive adequacy, highlighting further assumptions that may be notationally convenient but replaceable with more economical alternatives.

2. Monadic Mentalese

In this section, I describe a *possible* mind with an I-language whose expressions can only be used to construct monadic concepts. Such a mind exhibits no semantic typology of the usual sort, though it deploys concepts of various types. Applying this model to human minds, given known facts, requires appeal to some additional operations for converting one monadic concept into another. But such appeal may be unavoidable and independently plausible. And in any case, it can be instructive to see which facts can be accommodated *without* assuming that human I-expressions/I-concepts exhibit a wide range of Fregean types.

2.1 Possible Psychologies

For initial illustration, imagine a language whose syntax is exhausted by a unit-forming operation, UNIFY, corresponding to a *single* operation of concept composition. In such a language, every complex expression is of the form $[\alpha\ \beta]$, and the meaning of every expression can specified as follows: SEM($[\alpha\ \beta]$) = **O**[SEM(α), SEM(β)]; where '**O**' stands for a "macro" instruction to *execute* the two enclosed subinstructions, thereby obtaining two concepts, and then *compose* these concepts via the one invokable operation.

      For example, suppose that brown and cow are atomic combinables whose meanings are instructions to fetch concepts from certain lexical addresses. Then SEM([brown cow]) = **O**[SEM(brown), SEM(cow)] = **O**[fetch@brown, fetch@cow]. Likewise, if Bessie and cow are atomic combinables, SEM([Bessie cow]) = **O**[fetch@Bessie, fetch@cow]. And if Bessie can be combined with [brown cow], SEM([Bessie [brown cow]]) = **O**[SEM(Bessie), SEM([brown cow])] = **O**[fetch@Bessie, **O**[fetch@brown, fetch@cow]]. If the operation invoked is *monadic concept-conjunction*, then buildable concepts will all be of the following form: •[Φ(X), Ψ(X)]; where '•' stands for a dyadic concept that can (only) connect two monadic concepts to yield a third, which applies to whatever both constituent concepts apply to. We can represent the meaning of [brown cow], for such a mind, as shown below.

        SEM([brown cow]) = **CONJOIN**[fetch@brown, fetch@cow]

        SEM([brown cow]) ➜ •[BROWN(X), COW(X)]

Here, '**CONJOIN**' indicates a kind of instruction, and '➜' indicates the sort of concept that results from executing a given instruction. Ignoring polysemy for now, suppose that each lexical item is linked to exactly one fetchable concept, and hence that each lexical SEM is executable in exactly one way—viz., by fetching the corresponding concept.

      Given this simple combinatorics, atomic SEMs have to be instructions to fetch monadic concepts. In this language, Bessie must be linked to a monadic concept, BESSIE(X). But this

concept might apply to x iff x is a certain cow, who we call 'Bessie'; cp. Quine (1963), Burge (1973). And given a monadic concept for <u>Bessie</u> to fetch, executing SEM([<u>Bessie</u> [<u>brown</u> <u>cow</u>]]) is a way of constructing the concept •[BESSIE(X), •[BROWN(X), COW(X)]].

By contrast, suppose the sole operation of concept composition is *saturation*, which can combine a monadic concept like COW(X) with a singular concept like BESSIE to form COW(BESSIE). Given this language, <u>Bessie</u> can fetch BESSIE, while <u>cow</u> can fetch COW(X). Executing SEM([<u>Bessie</u> <u>cow</u>]) is then a way of constructing COW(BESSIE), which is a "propositional" concept. The singular constituent, BESSIE, is of a different type. Given these two types, $\langle t \rangle$ and $\langle e \rangle$, one can say that COW(X) is of type $\langle e, t \rangle$ and the propositional concept is of the following form: $\langle e, t \rangle (\langle e \rangle)$. More generally, the buildable concepts will exhibit the following abstract type: $\langle \alpha, \beta \rangle (\langle \alpha \rangle)$, indicating that a concept of adicity *n* is formed by saturating a concept of adicity *n+1* with an argument of an appropriate type. (Concepts of type $\langle t \rangle$ and $\langle e \rangle$ have adicity zero.) If [<u>brown</u> <u>cow</u>] is also an expression of this language, this instruction cannot be executed by using <u>brown</u> to fetch BROWN(X), a concept of type $\langle e, t \rangle$. But <u>brown</u> might fetch •[BROWN(X), X(X)]; where this higher-order concept of type $\langle\langle e, t \rangle, \langle e, t \rangle\rangle$ was previously defined in terms of BROWN(X) and linked to <u>brown</u> as a second fetchable concept. Then executing SEM([<u>brown</u> <u>cow</u>])—i.e., **SATURATE**[fetch@<u>brown</u>, fetch@<u>cow</u>]—could be a way of constructing the concept •[BROWN(X), COW(X)]; cp. Parsons (1970), Kamp (1975).

This second language, familiar in kind, permits lexical expressions that fetch dyadic concepts like CHASE(X, Y), which can be saturated by a singular concept to form a complex monadic concept like CHASE(X, BESSIE). Indeed, the operation of saturation itself imposes *no* constraints on which concepts can be fetched and combined with others: a concept of type $\langle \alpha, \beta \rangle$ can be combined with *either* a concept of the "lower" type $\langle \alpha \rangle$, thereby forming a concept of type $\langle \beta \rangle$, *or* any "higher" type $\langle\langle \alpha, \beta \rangle, \gamma \rangle$ such that $\langle \gamma \rangle$ is also a possible concept type.

If only for this reason, we should ask if we need to posit saturation as a composition operation in theories of I-languages. Even setting aside empirical arguments against such appeal (see Pietroski [2005a, 2010]), one might prefer to explore hypotheses according to which there are severe restrictions on the concepts that can be fetched by atomic I-expressions. For even if the specific proposals explored are wrong, seeing why can provide insights about the actual typology. A theory that imposes few constraints on the fetchable concepts may be harder to disconfirm. But "negative" facts, concerning nonexistent types and nonexistent meanings within a type, are relevant. And in any case, compatibility with facts is not the only theoretical virtue.

That said, incompatibility with facts is a vice. And monadic concept-conjunction cannot be the *only* operation invoked by I-expressions for purposes of combining fetchable concepts. Expressions like 'chase Bessie'—'chase every cow', 'saw Aggie chase Bessie', 'did not chase Bessie' etc.—are not simply instructions to conjoin monadic concepts fetched with the lexical items. But given neo-Davidsonian proposals, one can plausibly say that 'Aggie chase Bessie' is used to build a multi-conjunct concept: a concept that applies to things that have Aggie as their Agent, are chases, and have Bessie as their Patient; see Parsons (1990), Schein (1993, 2002). In my view, this model of composition is basically correct and extendable to other cases. By way of exploring this idea, according to which I-languages differ in just a few ways from the first "Conjunctivist" language imagined above, let me describe a possible range of atomic I-concepts and I-operations that permit construction of complex monadic concepts. In section three, I'll

offer a proposal about how such concepts could be fetched and combined as suggested, given a syntax that adds a *labeling* operation to UNIFY; cp. Hornstein (2009). The resulting account may be compatible with the facts.

2.2 Lexicalization

Imagine an initial stage of lexical acquisition in which many concepts are paired with phonological forms, so that certain perceptions (of sounds/gestures) reliably invoke certain lexicalizable concepts. During a second stage, each <PHON, concept> pair is assigned a lexical address that is linked to a bin, which may eventually contain one or more concepts that can be fetched via that address. But if B is the bin that is linked to address A, then a concept C can be added to B only if C is a *monadic* concept that is the result of applying an available "reformatting operation" to some concept *already* linked to A. Only a few reformatting operations are available. So there are constraints on which concepts can be fetched via any one lexical address.

Binned concepts must be monadic, because the computational system we are considering can only operate on concepts of this kind. The imagined mind has a language faculty that generates instructions to create complex concepts from simpler ones. But this modest faculty can only generate instructions of two kinds: those that call for *conjunction* of two monadic concepts, and those that call for *conversion* of one monadic concept into another. And while this limits the faculty's utility, the surrounding mind may be able to invent monadic analogs of nonmonadic concepts, thereby making the faculty more useful than it would otherwise be; cp. Horty's (2007) discussion of Frege on definition. For example, a dyadic concept like KICK(X, Y) might be used to introduce a monadic concept KICK(E), perhaps by introducing a triadic concept KICK(E, X, Y) such that KICK(X, Y) ≡ ∃E[KICK(E, X, Y)] and KICK(E, X, Y) ≡ AGENT(E, X) & KICK(E) & PATIENT(E, Y). Then given a proto-word of the form <PHON, bin-address, KICK(X,Y)>, the analytically related concept KICK(E) can be added to the bin, which will not contain the lexicalized dyadic concept.

More generally, this mind might create formally new monadic analogs of lexicalizable concepts as follows: use a concept C^n of adicity n to introduce a concept C^{n+1} of adicity $n+1$; and use C^{n+1}, along with n "thematic" concepts that are independently available, to introduce a monadic concept C^1. Suppose that given a singular concept like BESSIE, this mind can also create an analog monadic concept. For illustration, IDENTITY(X, BESSIE) will do. But given a proto-word of the form <PHON, bin-address, BESSIE>, one can imagine forming the corresponding monadic concept CALLED(PHON, X), which applies to anything called with the PHON in question. And if CALLED(PHON, X) is added to the bin, it might later be fetched and conjoined with another concept—perhaps demonstrative—so that at least in the context of use, the resulting concept of the form •[CALLED(PHON, X), Φ(X)] applies to exactly one individual, like the one mentally denoted with BESSIE; see, e.g., Burge (1973), Katz (1994), Longobardi (1994), Elbourne (2005).

When a monadic concept is lexicalized, it may be added to its own bin. But this does not guarantee conjoinability with other concepts. Suppose the concept lexicalized with (the PHON of) 'brown' is a concept of surfaces, while the concept lexicalized with 'house' is not. Then the proto-word <PHON, bin-address, BROWN(S)> may, given suitable prompts, lead to introduction of a concept that applies to the *brown-surfaced*: BROUNS(X) ≡ ∃S[SURFACE(S, X) & BROWN(S)]. In which case, BROUNS(X) could be added to the bin, making it possible to coherently conjoin HOUSE(X) with a concept fetched via the address initially linked to BROWN(S); cp. Chomsky (2000b). As this example suggests, lexicalization might lead to a polysemous word that bins several monadic concepts. Whatever is initially lexicalized with 'book', the end result may be a

6

lexical item that can be used to fetch any of several concepts—including at least one that applies only to certain spatiotemporal particulars created by publishers, and another that applies only to certain abstracta created by authors. Likewise, for mature speakers, the bin for 'country' may include at least two concepts: one that applies to the French *polis*, but not to the *terrain* inhabited by the citizens of France; and a distinct concept that applies to this terrain, but not the occupying polis. And perhaps only the polis-concept can be coherently conjoined with at least one concept fetched via the word 'republic', while only the terrain-concept can be coherently conjoined with at least one concept fetched via the word 'hexagonal'.[7]

2.3 Number Neutrality

Especially given the possibility of reformatting, we need to be clear about kinds of *variables* that can appear in the concepts fetched/assembled via I-expressions. Other things equal, one wants to posit as few kinds as possible. I see no way of avoiding appeal to various sortals, including sortals for "eventish" things that can have other things as participants. Whatever one says about I-operations and adicities, we need distinctions among predicates; see, e.g., Vendler (1959), Dowty (1979, 1991), Baker (1997), Svenonius (forthcoming). But one can try to minimize the number of logical types posited. And this quickly leads to questions about whether to accommodate plurality with one kind of variable or two.

One traditional approach treats all conceptual variables as singular, but sorted in a way that is usually interpreted in terms of a split-level domain: first-order variables range over whatever they do; and second-order variables range over "plural entities"—sets, collections, or mereological sums—whose elements are things over which the first order variables range.[8] And we can certainly imagine a mind with I-concepts like $COW(X_{-PL})$ and $COW(X_{+PL})$; where the fomer applies to cows and the latter applies to sets of cows. Or in more explicitly Tarskian terms: $COW(X_{-PL})$ is satisfied by a sequence σ of domain entities iff the entity that σ assigns to the unplural variable is a (basic entity that is a) cow; and $COW(X_{+PL})$ is satisfied by σ iff the entity that σ assigns to the plural variable is a plural entity whose every element is a cow. From this perspective, a word like 'three' is used to fetch a nondistributive concept like $THREE(X_{+PL})$, which is satisfied by σ iff the entity that σ assigns to the plural variable is a plural entity with three elements. But at least if the focus is on I-concepts, which have whatever character they do, one need not think of each assignment of values to variables as a Tarskian *sequence* that assigns exactly *one* value to each variable.

Following Boolos (1998), theorists can allow for assignments that assign many values to a variable. Correlatively, we can imagine a mind with number-neutral concepts like $COW(\mathbb{N})$, which applies to one or more things iff each of them is a cow; where 'thing*s*' exhibits *grammatical* agreement with 'one or more', with no suggestion of more than one. That is, an assignment **A** satisfies $COW(\mathbb{N})$ iff the *one or more things* that **A** assigns to the number-neutral variable are such that each of them is a cow. From this perspective, 'three' is used to fetch a concept that applies to one or more things iff they are three (and hence more than one). Some things are three iff they correspond one-to-one with the points of a triangle, the words in the series 'one, two, three', etc. So given three cows, no one or two of them are three; but each is a cow, and any two of them are cows. So the concept •[$THREE(\mathbb{N})$, $COW(\mathbb{N})$] is as well-formed as •[$BROWN(\mathbb{N})$, $COW(\mathbb{N})$]. The former concept does not apply to any one or two cows; though likewise, the latter concept does not apply to any red or green cows. The concepts $ONE(\mathbb{N})$ and

~ONE(ℕ)—a.ka. ~PLURAL(ℕ) and PLURAL(ℕ)—can combine with COW(ℕ) to form the "singular" concept •[ONE(ℕ), COW(ℕ)] and the "plural" concept •[~ONE(ℕ), COW(ℕ)]; where the former applies to one or more things iff each is a cow and there is only one of them, while the latter applies to one or more things iff each is a cow and there are more than one of them.

For many purposes, we can adopt either view of I-concept variables: as essentially singular, always taking a single value relative to any assignment, but with variables of one sort ranging over sets of entities over which variables of the other sort range; or as number-neutral, ranging over whatever variables range over, but allowing that a concept can apply to some things without applying to any one of them. Following Schein (1993, 2002, forthcoming), I think the number-neutral approach is empirically superior; see Pietroski (2005a, 2006, 2010). But here, the more important point is that the sorted approach is not mandatory. And if we want to locate the sparest plausible assumptions about I-concepts, one might well start with the hypothesis that all I-concepts are number-neutral, allowing for specific concepts like ONE(X)/~ONE(X) that are not neutral. For if distinctively plural variables are required, there should be evidence of this.[9]

This point is especially important in the current setting, because Boolos ingeniously explored the resources available within *monadic* second-order logic. If these resources suffice for human I-language semantics, that is worth knowing. Assuming that the resources of first-order logic are inadequate, adopting a "singularist" perspective on I-concept variables certainly invites—and it may require—the standard Fregean typology, given the limitations imposed by Tarskian sequences. Especially when asking *if* the typology is required, as opposed to convenient, we must not assume models of plural locutions that were designed to fit the typology. So in what follows, I will assume that appeal to number-neutral variables is legitimate, especially in a theory that already posits monadic reformatting as part of lexicalization.

2.4 Extending Monadicity

It cannot be that *all* I-concepts are monadic. We can express relational thoughts. But this does not require a recursive combination operation, like saturation, that can take polyadic concepts as inputs. Conjunction can yield a simulacrum of polyadic thought given repeated—though not necessarily recursive—appeal to a severely restricted kind of dyadicity.

Earlier, I restricted '•' to combination of monadic concepts. But imagine a mind that allows for one exception: instances of ∃•[θ(E, X), Φ(X)] are well-formed; where θ(E, X) is a dyadic concept whose second variable is *the* variable of the monadic concept Φ(X), and this variable is immediately *closed* to further conjunction, leaving a concept that applies to one or more things iff they bear the relation in question to one or more things to which Φ(X) applies. This permits concepts like ∃•[AGENT(E, X), COW(X)], which applies to one or more events iff some cows were the agents of those events; cp. Carlson (1984), Schein (2002). Simplifying a little, we can say that some cows are the agents of some events iff each cow is the agent an event, and each event has a cow as its agent; where each event has at most one agent. The more complex concept •[KICK(E), ∃•[PATIENT(E, X), •[COW(X), PLURAL(X)]]] applies to one or more events iff: each is a kick; and one or more things are such that they are the patients of those events, and they are cows (i.e., each of them is a cow, and they are not one). Likewise, •[∃•[AGENT(E, X), •[RANCHER(X), FIVE(X)]], •[BRAND(E), ∃•[PATIENT(E, X), •[COW(X), FIFTY(X)]]]] applies to one or more events iff in those events, five ranchers branded fifty cows.

Concepts of events are in no sense true or false. And *perhaps* concepts of type <e> and

<t> will have to be introduced eventually, along with concepts of higher types. I discuss quantification in section three. But Tarski (1933) provided a semantics for the first-order predicate calculus without appeal to truth values, and without treating closed sentences as instances of special type <t>, by effectively treating sentences as predicates: expressions satisfied by sequences of entities. So let's be clear that "propositional concepts," which can be negated and conjoined, need not be concepts of truth/falsity.

Consider a pair of operators, ↑ and ↓, that create monadic concepts from monadic concepts; where for any one or more things, ↑Φ(X) applies to them iff Φ(X) applies to one or more things, and ↓Φ(X) applies to them iff Φ(X) applies to nothing.[10] More briefly, without fussing about number neutrality: for each entity, ↑Φ(X) applies to it iff Φ(X) applies to something; and ↓Φ(X) applies to it iff Φ(X) applies to nothing. One can think of ↑ and ↓ as polarizing operators that convert any monadic concept into a concept of everything or nothing, perhaps akin to EXIST(X) and ~EXIST(X). For example, given any entity, ↑COW(X) applies to it iff COW(X) applies to something; so ↑COW(X) applies to you, and likewise to me, iff there is a cow. By contrast, ↓COW(X) applies to you (and me) iff nothing is a cow. And for each thing, either ↑COW(X) or ↓COW(X) applies to it—since it is either such that there is a cow, or such that there is no cow. This mode of composition clearly differs from the always *restricting* operation of conjunction. But correlatively, nothing is such that both ↑COW(X) and ↓COW(X) apply to it. Hence, nothing is such that •[↑COW(X), ↓COW(X)] applies to it.

Given a suitable metalanguage, we can say: $\uparrow\Phi(X) \equiv \exists Y[\Phi(Y)]$, and $\downarrow\Phi(X) \equiv \sim\exists Y[\Phi(Y)]$. But the idea is not that '↑' and '↓' are abbreviations for existential closure and its negation. For example, '↑BETWEEN(X, Y, Z)' is gibberish, as is '↓AGENT(E, X)'. The idea is rather that certain I-expressions, perhaps associated with tense and/or negation, invoke "closure" operations that convert a monadic concept (say, of events) into a concept of all or none. So let's say that any concept of the form ↑Φ(X) or ↓Φ(X) is a *T-concept*, with 'T' connoting Tarski, totality, and truthy. Note that for any concept Φ(X) and any entity e, ↑↑Φ(X) applies to e iff ↓↓Φ(X) does, since each of these concepts applies to e iff ↑Φ(X) does—i.e., iff Φ(X) applies to something. Likewise, ↑↓Φ(X) applies to e iff ↓↑Φ(X) does, since each of these concepts applies to e iff ↓Φ(X) does—i.e., iff Φ(X) applies to nothing. And while ↑•[Φ(X), Ψ(X)] applies to e iff something falls under the conjunctive concept •[Φ(X), Ψ(X)], which applies to e iff e falls under both conjuncts, [↑Φ(X), ↑Ψ(X)] applies to e iff (e is such that) something falls Φ(X) *and* something falls under Ψ(X). Thus, ↑•[BROWN(X), COW(X)] is a more restrictive concept than •[↑BROWN(X), ↑COW(X)], much as the more familiar ∃Y[BROWN(Y) & COW(Y)] implies ∃Y[BROWN(Y)] & ∃Y[COW(Y)] but not conversely. Correlatively, ↓•[BROWN(X), COW(X)] applies to e iff nothing is both brown and a cow, while •[↓BROWN(X), ↓COW(X)] applies to e iff (e is such that) nothing is brown *and* nothing is a cow. So •[↓BROWN(X), ↓COW(X)] is a more restrictive than ↓•[BROWN(X), COW(X)], and ↓COW(X) is more restrictive than ↓•[BROWN(X), COW(X)].

The basic idea is medieval: the default direction of inference is conjunction *reduction*—e.g., from •[BROWN(X), COW(X)] to COW(X); but in the presence of a negation-like operator, this default is reversed.[11] And note that when the concepts conjoined are both T-concepts, which apply to all or none, "closing up" has no effect. If **P** and **Q** are T-concepts, and so each is of the form ↑Φ(X) or ↓Φ(X), then ↑•[**P**, **Q**] is logically equivalent to •[**P**, **Q**]: ↑•[**P**, **Q**] applies to e iff something/everything falls under both **P** and **Q**; •[**P**, **Q**] applies to e iff e/everything falls under both **P** and **Q**. By contrast, ↓•[↓**P**, ↓**Q**] applies to e iff: nothing falls under both ↓**P** and ↓**Q**; i.e.,

9

nothing is such that both **P** and **Q** are empty; i.e., something falls under **P** *or* something falls under **Q**. So propositional disjunction can be characterized, *a la* de Morgan, given T-concepts.

More generally, T-concepts provide resources for accommodating the meanings of sentential I-expressions without supposing that they exhibit a special semantic type <t>. So we should pause before assuming that HFL generates expressions of this type, as opposed to expressions that can be used to construct T-concepts, which can bear an intimate relation to existential *thoughts* of type <t>. While "post-linguistic" cognition may traffic in complete thoughts, in which each monadic concept is saturated or quantificationally bound, HFL may interface with such cognition via formally monadic T-concepts. The notion of 'sentence' has always had an unstable place in grammatical theory. And especially within MP, one might want to preserve the old idea that each I-expression is (labeled as) an instance of some grammatical type exhibited by some *atomic* expression. One can stipulate that sentences are projections of some functional category. But no such stipulation seems especially good. So perhaps we should drop the idea that HFL generates expressions of type <t>, and adopt a more Tarskian type-free approach to human I-language semantics.[12]

2.5 Abstracting

At this point, our imagined mind can form many systematically related concepts. It can also convert an ordinary monadic concept—one that can apply to some but not all things—into a T-concept that must apply to all or none. But it cannot yet do the converse. And this is arguably the most interesting respect in which human thought is recursive. Given •[**P**, **Q**], the capacity to form •[**P**, •[**P**, **Q**]] is not that impressive. And while I cannot discuss the semantics of complementizer phrases (see Pietroski [2005a] for a Conjunctivist analysis), the kind of recursion exhibited 'Aggie thinks that Bessie said that Aggie saw Bessie' is many ways less interesting than the kind exhibited by 'who saw the cow that Aggie saw'. Embedding one sentence in another is a good trick. Using a sentence to create a concept that can apply to some but not all things is a great trick. Clearly, this requires more than mere conjunction. But as Tarski (1933) showed us, the requisite machinery is relatively simple, even if it initially seems complex. It involves a distinctive kind of composition, though not one that depends on an operation of saturation. So if appeal to this kind of composition is unavoidable, as I suspect, our question is whether we should appeal to it and saturation and conjunction.

Let's assume that our imagined mind can deploy *indices*, like '1' and '2', that can be used in two ways: deictically, as devices for temporarily tracking salient things perceived (cp. Pylylshyn [2007]); or anaphorically, as devices for temporarily tracking things independently described. Some of these indices may be singular, but suppose that some are number-neutral. Let's also suppose that this mind also has some concepts like FIRST(X) and SECOND(X), which apply to whatever the corresponding indices are tracking. Such concepts are, in an obvious sense, context sensitive in a way that concepts like COW(X) are not. In an equally obvious sense, COW(X) applies to different things at different times, as cows come and go. But indices are, so to speak, designed as temporary tracking devices with no independent content of their own. So as an idealization, we can say that COW(X) simply applies to cows, without relativization to anything else; although CALF(X) is already more complicated. By contrast, FIRST(X) doesn't apply to anything *tout court*: FIRST(X) is satisfied by an assignment **A** iff the one or more things that **A** assigns to the conceptual variable are whatever **A** assigns to the first index; COW(X) is satisfied by **A** iff the one or more things that **A** assigns to the conceptual variable are cows.[13]

This allows for concepts like ∃•[INTERNAL(E, X), FIRST(X)], which is satisfied by **A** iff: whatever things **A** assigns to the first index, those one or more things are the internal participants of whatever **A** assigns to the free conceptual variable. Likewise, the T-concept ↑•[∃•[EXTERNAL(E, X), SECOND(X)], •[SAW(E), ∃•[INTERNAL(E, X), FIRST(X)]]] is satisfied by **A** iff: whatever **A** assigns to the second index saw whatever **A** assigns to the first index; or more longwindedly, (all things are such that) there were one or more events of seeing whose external participants are whatever **A** assigns to the second index and whose internal participants are whatever **A** assigns to the first index. But let's suppress the eventish and conjunctive substructure, abbreviating this T-concept as follows: **2SAW1**. For regardless of whether T-concepts are formed by conjunction or saturation, T-concepts with constituents like FIRST(X) are concepts ripe for abstraction.

Given any index *i* and T-concept **P**, which can evaluated relative to any assignment **A**, let **TARSKI**{*i*, **P**} be the semantic concept indicated below;

$$∃\mathbf{A}^*{:}\mathbf{A}^*≈_1\mathbf{A}\{\text{ASSIGNS}(\mathbf{A}^*, \text{X}, 1) \ \& \ \text{SATISFIES}(\mathbf{A}^*, \mathbf{P})\}$$

where ASSIGNS(**A***, X, 1) applies to one more things iff they are the things that **A*** assigns to the first index, and '**A***≈₁**A**' means that **A*** differs from **A** at most with regard to what it assigns to the first index. To be sure, any natural concept of satisfaction is likely to differ from Tarski's. But the idealization is that a suitably equipped mind can use a T-concept with FIRST(X) as a constituent to form a concept that applies to one or more things (relative to a certain assignment of values to indices) iff making them the values of the first index (and holding everything else constant) satisfies the T-concept. Likewise, given a T-concept with SECOND(X) as a constituent, one can form a concept that applies to one or more things (relative to a certain assignment of values to indices) iff making them the values of the second index (and holding everything else constant) satisfies the T-concept. One can think of this as number-neutral lambda abstraction. But a Church-style construal of 'λX.Φ(X)' presupposes sequence variants and a Tarski-style construal of 'Φ(X)'. And the goal here is to be explicit about theoretical commitments.

In the context of our example, relative to any assignment **A**:

$$\mathbf{TARSKI}\{1, \mathbf{2SAW1}\} = ∃\mathbf{A}^*{:}\mathbf{A}^*≈_1\mathbf{A}\{\text{ASSIGNS}(\mathbf{A}^*, \text{X}, 1) \ \& \ \text{SATISFIES}(\mathbf{A}^*, \mathbf{2SAW1})\}$$

and this concept applies to one or more things iff they were seen by whatever **A** assigns to '2', since if **A***≈₁**A**, then both assignments assign the same one or more things to '2'; similarly,

$$\mathbf{TARSKI}\{2, \mathbf{2SAW1}\} = ∃\mathbf{A}^*{:}\mathbf{A}^*≈_2\mathbf{A}\{\text{ASSIGNS}(\mathbf{A}^*, \text{X}, 2) \ \& \ \text{SATISFIES}(\mathbf{A}^*, \mathbf{2SAW1})\}$$

and this concept applies to one or more things iff they saw whatever **A** assigns to '1', since if **A***≈₂**A**, then both assignments assign the same one or more things to '1'. I readily grant that this kind of concept construction—from **2SAW1** to **TARSKI**{1, **2SAW1**} or **TARSKI**{2, **2SAW1**}—is more sophisticated than conjunction. Indeed, Tarskian composition violates some conceivable compositionality constraints respected by conjunction. Relative to **A**, the T-concept **2SAW1** may apply to nothing, while **TARSKI**{1, **2SAW1**} applies to many things; see Salmon (2006). Suppose that whatever **A** assigns to '2', it/they saw many things, but *not* whatever **A** assigns to '1'. Then relative to a single assignment: **2SAW1** is false of each thing, and in that sense false, yet **TARSKI**{1, **2SAW1**} is true of many things; hence, ↑**TARSKI**{1, **2SAW1**} is true of each thing.

In this sense, **2SAW1** can be false while ↑**TARSKI**{1, **2SAW1**} is true. And since whatever **A** assigns to '2' might have seen nothing, **2SAW1** and ↑**TARSKI**{1, **2SAW1**} can both be false. Like it or not, this kind of "non-truth-functional" composition is available to a mind that is

equipped to perform the Tarski trick. And my claim is *not* that this kind of abstraction can be reduced to anything else. On the contrary, I see no way to avoid positing a capacity for such composition in the construction of I-concepts. But I also see no way to avoid appeals to a more mundane operation of conjunction, and at least a few thematic concepts; see Baker (1997). In my view, the question is whether we *also* need Fregean typology and an operation of saturation.

3. Back to SEMs

Having described a possible mind with the capacities needed to construct a wide range of potential I-concepts, let me turn to the task of showing how I-expressions might be systematically described as instructions to build such concepts. In section 2.1, I imagined a language whose syntax is exhausted by a unit-forming operation, UNIFY. Let's now suppose that human I-languages make it possible to unify/concatenate expressions and *label* them so that a complex operation MERGE can be defined as follows (see Hornstein [2009]): MERGE(α, β) = LABEL{UNIFY(α, β)} = LABEL{[α, β]}; where the new operation (deterministically) selects one of the two expressions just unified and appends a copy to the unified expression. The idea is that if α has the right properties to be the "head" of [α β], then LABEL{[α, β]} = [α, β]$_\alpha$. In which case, MERGE(α, β) = [α, β]$_\alpha$, as desired. But what kind of instruction is [α, β]$_\alpha$?

3.1 Conjunction and Conversion

For any instructions **I** and **I***, let +[**I**, **I***] be a "macro" instruction to execute the two subinstructions and conjoin the results, thereby creating a concept of the form •[Φ(x), Ψ(x)]. Then examples like [brown$_A$ cow$_N$]$_N$ and [cow$_N$ [that Aggie saw]$_C$]$_N$, ignoring structure within the relative clause, conform to a very simple view: for any expressions α and β, SEM([α, β]$_\alpha$) = +[SEM(α), SEM(β)]. One might well endorse the medieval suspicion that modulo special expressions like negation, the general trend is for [α, β]$_\alpha$ to be more restrictive than its constituents. This trend would be surprising if concatenation signifies an operation (like saturation) that is indifferent to whether or not a complex expression carries more information than its parts. And even many apparent counterexamples suggest complications of the trend, as opposed to wholesale departures. A big ant is an ant that meets a further condition; and even a fake diamond is a fake of a certain sort. Likewise, a chase of Bessie is a chase. But phrases like [chase$_V$ Bessie$_N$]$_V$ suggest that the *phrasal label*—or more precisely, a *mismatch* between the phrasal label and the *other* constituent label—can play a significant role.[14]

    At least for cases of combining constituents that correspond to concepts of different sorts, like a concept of events and a concept of an animal, a natural thought is that the phrasal label invokes an *adapter* that combines with one concept to form a concept of the *same sort* as the other. In terms of [chase$_V$ Bessie$_N$]$_V$, perhaps the phrasal label V is an instruction to use the result of executing SEM(Bessie$_N$) in creating a concept that is sure to be conjoinable with the concept obtained by executing SEM(Bessie$_N$). There are various ways of encoding this idea. But consider the following principle of composition: SEM([α, β]$_\alpha$) = +[SEM(α), **ADAPT**[SEM(β), α]]; where '**ADAPT**' stands for a macro instruction to execute the subinstruction and use the resulting monadic concept to form another, via some operation determined by the label α and the available conversion operations. Obviously, the work lies with specifying the specific instances of '**ADAPT**' in an empirically adequate and motivated way. But for [chase$_V$ Bessie$_N$]$_V$, in which Bessie$_N$ is effectively classified as the internal argument of chase$_V$, we already have what is needed. Suppose that classifiying an argument as internal is an instruction to use the argument to construct a concept Φ(x), and then a concept of things whose "internal participants" fall under

12

the concept. More explicitly, one can adopt the hypothesis below.

$$\textbf{ADAPT}[\text{SEM}(\beta), V] = \textbf{INTERNALIZE}{:}\text{SEM}(\beta)$$

For any expression β, let **INTERNALIZE**:SEM(β) be the macro instruction to execute SEM(β) and use the resulting concept $\Phi(X)$ to create a concept of the following form: $\exists\bullet[\text{INTERNAL}(E, X), \Phi(X)]$. The idea is that INTERNAL(E, X) is a "thin" but formally thematic concept that groups together PATIENT(E, X), perhaps THEME(E, X), and any other "thick" thematic concept—with independent conceptual content—that can be introduced by classifying an expression as the internal argument of a predicate. One way or another, the lexical item 'chase' can indicate that any internal participants of chases are *patients*, making it possible to replace $\bullet[\text{CHASE}(E), \exists\bullet[\text{INTERNAL}(E, X), \Phi(X)]]$ with $\bullet[\text{CHASE}(E), \exists\bullet[\text{PATIENT}(E, X), \Phi(X)]]$.[15] In any case, [chase$_V$ Bessie$_N$]$_V$ can direct construction of $\bullet[\text{CHASE}(E), \exists\bullet[\text{INTERNAL}(E, X), \text{BESSIE}(X)]]$.

This kind of "thematic conversion," invoked to preserve a fundamentally Conjuncivist conception of semantic composition, is formally similar to a more familiar kind. If the sole combination operation is saturation, then faced with examples like [brown$_A$ cow$_N$]$_N$, one might adopt some version of the following view.

$$\text{SEM}([\alpha, \beta]_\alpha) = \textbf{SATURATE}[\text{SEM}(\alpha), \textbf{ADAPT}[\text{SEM}(\beta), \alpha]]$$
$$\textbf{ADAPT}[\text{SEM}([\dots]_A), N] = \textbf{LIFT}{:}\text{SEM}([\dots]_A)$$
$$\text{SEM}([\text{brown}_A \text{ cow}_N]_N) = \textbf{SATURATE}[\text{fetch@cow}_N, \textbf{LIFT}{:}\text{fetch@brown}_A]$$
$$\rightarrow \bullet[\text{BROWN}(X), X(X)]\{\text{COW}(X)\} = \bullet[\text{BROWN}(X), \text{COW}(X)]$$

The idea here is that classifying brown$_A$ as the "inferior" constituent of [brown$_A$ cow$_N$]$_N$ is an instruction to use the concept fetched (or constructed) via this consitutent into an analytically related concept of the higher type $\langle\langle e, t\rangle, \langle e, t\rangle\rangle$, which can be saturated by the "head" concept fetched (or constructed) via the noun. But at least as a theory of I-languages, this assumes the availability of higher types, as well as an operation of conjunction.

From an E-language perspective, one can be less committal and say merely that words indicating two *functions* of type $\langle e, t\rangle$—$\lambda x.T$ *iff x is a cow* and $\lambda x.T$ *iff x is brown*, where '*T*' stands for a certain truth value, and '*x*' is a singular variable—are combined to form a phrase that indicates a third function of the same type, $\lambda x.T$ *iff x is both a cow and brown*. From this extensional perspective, corresponding roughly to Marr's (1982) computational Level One, saturating-and-lifting is equivalent to conjoining. But from an I-language/procedural perspective, corresponding more closely to Marr's algorithmic Level Two, these are distinct operations: the former presupposes the latter as a subpart; and while conjunction might be described as a very special case of saturation, restricted to concepts of one type, it might also be described as a basic operation. Moreover, since nouns can be modified with relative clauses, the requisite lifting operation would have to be available as a *recursive* operation. So especially if many adverbial modifiers have to be diagnosed in terms of monadic concepts of events, it seems that the requisite lifting operations will encode a Conjunctivist principle of semantic composition *and more*. In which case, theorists may as well posit more than one basic composition operation; cp. Higginbotham (1985), Larson and Segal (1995), Heim and Kratzer (1998). By contrast, appeals to thematic relations seem unavoidable, if only to formulate the *constraints* on how they can project to grammatical relations in human I-languages. This invites the Minimalist project of making do with conjunction by assigning thematic significance to certain cases of labeling.

Assigning such significance leaves room for the possibility that some I-expressions are

unlabeled instances of the form [α, β]; cp. Chametsky (1996) on adjunction. If [brown cow$_N$] or [cow$_N$ [that I saw]] is an example, with the unlabeled constituent not being a candidate for the phrasal head, perhaps SEM([α, β]) = +[SEM(α), SEM(β)]; see Hornstein and Pietroski (forthcoming) for discussion. One can add that SEM([α, β]$_α$) = +[SEM(α), **ADAPT**[SEM(β), α]]. Alternatively, one can say that all phrases are labeled, and say that some phrasal labels call for a "null" adapter. Perhaps **ADAPT**[SEM([…]$_C$), $_N$] = **RETURN**:SEM([…]$_C$). In which case, SEM([cow$_N$ [that I saw]$_C$]$_N$) = +[SEM(cow$_N$), SEM([that I saw]$_C$)]. If [brown$_A$ cow$_N$]$_N$ really means something like *is a cow that is brown for a cow*, there may be few if any cases of pure adjunction apart from relative clauses. But that would still leave endlessly many cases.

A related point is that words like cow$_N$ and chase$_V$ may already be combinations of lexical roots with functional items that serve as labels. If cow$_N$ = [√cow $_N$], then perhaps SEM(cow$_N$) = +[SEM(√cow), SEM($_N$)] = +[fetch@√cow, fetch@$_N$]; where $_N$ is a device for fetching a functional monadic concept like INDEXABLE, while $_V$ is a device for fetching a concept like TENSABLE, thus allowing for a distinction between SEM(chase$_V$) and SEM(chase$_N$). I cannot pursue these issues here, but simply raise them to note the kinds of resources still available *without* appeal to Fregean typology; see Hornstein and Pietroski (2010), drawing on Marantz (1984), Halle and Marantz (1993), Baker (2003), and Borer (2005).

Let's return to the idea that [chase$_V$ Bessie$_N$]$_V$ is an instruction to build a concept like •[CHASE(E), ∃•[INTERNAL(E, X), BESSIE(X)]], with the grammatical object of the verb used to fetch or construct a concept that can restrict the participant variable of a "thin" thematic concept. There is an obvious analog for subjects, as in [Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_V$. Suppose we have a formally thematic concept, EXTERNAL(E, X), that groups together AGENT(E, X), EXPERIENCER(E, X), and any other "thick" thematic concepts—with independent conceptual content—that can be introduced by classifying an expression as the external argument of a predicate. One way or another, 'chase' can indicate that any external participants of chases are agents, making it possible to replace •[∃•[EXTERNAL(E, X), Φ(X)], •[CHASE(E, X), …] with •[∃•[AGENT(E, X), Φ(X)], •[CHASE(E, X), …]].

The requisite conversion operation is easily defined. For any expression β, let **EXTERNALIZE**:SEM(β) be the macro instruction to execute SEM(β) and use the resulting concept Φ(X) to create a concept of the following form: ∃•[EXTERNAL(E, X), Φ(X)]. But we can't say both of the following, at least not without qualification.

> **ADAPT**[SEM(β), V] = INTERNALIZE:SEM(β)
> **ADAPT**[SEM(β), V] = EXTERNALIZE:SEM(β)

Correspondingly, we can't say that SEM([Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_V$) = +[**ADAPT**[SEM(Aggie$_N$), V], +[SEM(chase$_V$), **ADAPT**[SEM(Bessie$_N$), V]. This doesn't make it clear *which* conversion operation, **EXTERNALIZE** or **INTERNALIZE**, goes with *which* grammatical argument. But there are three obvious possibilities to consider.

Perhaps labels should be viewed, not as atomic elements, but as stand-ins for the entire "head expression;" cp. Chomsky (1995). If [Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_V$ = [Aggie$_N$ [chase$_V$ Bessie$_N$]$_{CHASE}$]$_{CHASE\_BESSIE}$, which can be abbreviated as [Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_{V(N)}$, then SEM([Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_V$) = +[**ADAPT**[SEM(Aggie$_N$), V(N)], +[SEM(chase$_V$), **ADAPT**[SEM(Bessie$_N$), V]. This effectively classifies external arguments as such, allowing for the obvious rules.

> **ADAPT**[SEM(β), V] = **INTERNALIZE**:SEM(β)

14

ADAPT[SEM(β), V(N)] = **EXTERNALIZE**:SEM(β)

Or perhaps the syntax is nuanced in a different way—independently suggested by many authors, including Chomsky (1995) and Kratzer (1996)—with external arguments as arguments of an independent verbal element, as in [Aggie$_N$ [$\underline{v}$ [chase$_V$ Bessie$_N$]$_V$]$_{\underline{V}}$]$_V$. For these purposes, 'V(N)' can be replaced with '$\underline{V}$'.

ADAPT[SEM(β), $\underline{V}$] = **EXTERNALIZE**:SEM(β)

There are real empirical issues here. For example, is there a covert internal argument in 'Aggie counted'? But for better or worse, the Conjunctivist framework does not force a particular stand on these issues.

  A third option is that the conversion operation is not determined by head label *alone*. Perhaps cyclicity plays a role here: within a given "phase" of instruction *execution*, the first occurrence of label 'V' triggers the first operation (**INTERNALIZE**), and the second occurrence of 'V' triggers the second operation (**EXTERNALIZE**). If there are at most two grammatical arguments per phase/cycle/whatever, one might imagine a binary "switch" that gets "reset" to its initial state at the start of each cycle; cp. Boeckx (2008). If some such thought is correct, perhaps we can make do with formally thematic concepts that are super-thin: ON(E, X) and ~ON(E, X), instead of EXTERNAL(E, X) and INTERNAL(E, X).

  In any case, an expression like [Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_V$ can be an instruction to build a concept like •[∃•[EXTERNAL(E, X), AGGIE(X)], •[CHASE(E), ∃•[INTERNAL(E, X), BESSIE(X)]]]. Adding adverbs and prepositional phrases is not without difficulties. But the leading idea, unsurprisingly, is that I-expressions like 'yesterday'—'on Tuesday', 'with a stick', etc.—are instructions to fetch/construct additional conjuncts. Prepositions, as functional elements, can be viewed as instructions to fetch adapters and convert concepts like STICK(X) into concepts like ∃•[INSTRUMENT(E, X), STICK(X)]. This provides a way of describing the massive polysemy of prepositions: there need not be a single "thematizing" operation that 'with' invokes. And prepositional phrases may well have internal Conjunctivist structure. I cannot pursue this rich topic here; but see Svenonius (forthcoming).

  At this point, let me offer an explicit treatment of sentential expressions and relative clauses, before turning to quantificational constructions, which pose the most obvious challenge for a Conjunctivist semantics. For simplicity, let's ignore tense. Eventish treatments are familiar; see Higginbotham (1985), Parsons (1990). It is also worth remembering that 'Aggie chase Bessie' can appear as an internal argument of 'see'. And the current proposal lets us treat both 'see trees' and 'see Aggie chase Bessie' as instructions to build concepts of seeings whose internal participants are one or more things that meet a certain condition: being trees, or being chases of Bessie by Aggie; cp. Higginbotham (1983). But at some point, a clause is treated as sentential. And if the concept built via [Aggie$_N$ [chase$_V$ Bessie$_N$]$_V$]$_V$ is prefixed with ↑, the result

    ↑•[∃•[EXTERNAL(E, X), AGGIE(X)], [CHASE(E), ∃•[INTERNAL(E, X), BESSIE(X)]]]

is a T-concept that applies to all or none, depending on whether or not Aggie chased Bessie. Let's abbreviate this concept as **ACHASEB**, recalling the discussion at the end of section two.

  There are many ways of encoding the idea that a tensed version of 'Aggie chase Bessie' can be an instruction to create a T-concept, depending on what one thinks about sentences and sentential negation. For instead of thinking about sentences as a special kind of grammatical category, headed by a special functional item, one might think of sentences as results *in thought* of "spelling out" tensed instructions; cp Uriagereka (1999). If a sentence corresponds to a cycle

(or phase) of interpretation, the relevant I-expression may direct construction of a monadic concept that can be true of some things but not others. But this concept may be converted to a T-concept at the Conceptual-Intentional interface, given the demands of the judgment systems of external to HFL. And instead of thinking about overt negation in human I-languages as a modifier of sentences, as in familiar formal languages, one can hypothesize two "modes of closure" at the interface: in the absence of an instruction to the contrary, use ↑; but given overt negation, use ↓. That said, one can also think of sentential classification as an instruction to adapt a subsentential instruction by invoking the positive T-operator.

$$\text{SEM}([\ldots[\text{Bessie}_N \, [\text{chase}_V \, \text{Aggie}_N]_V]_V\ldots]_S) =$$
$$\textbf{UP}:\text{SEM}(\ldots[\text{Aggie}_N \, [\text{chase}_V \, \text{Bessie}_N]_V]_V\ldots)$$

For present purposes, let's remain neutral about the details, and just say that executing the sentential instruction 'Bessie did chase Aggie' results in construction of the concept **ACHASEB**. And one can still treat 'Bessie did not chase Aggie' as an instruction to construct ↓**ACHASEB**, which is a T-concept of the form ↓↑•[...]—or as an instruction to construct a logically equivalent concept of the form ↓•[...], which applies to none or all, depending on whether or not •[...] applies to one more things.

The more important point here is that if pronouns and traces of movement are instructions to fetch concepts like $\text{FIRST}(X)$ and $\text{SECOND}(X)$, as suggested at the end of section two, then relative clauses are easily accommodated. Recall that relative to any assignment **A**, $\text{FIRST}(X)$ applies to whatever **A** assigns to the first index; likewise for $\text{SECOND}(X)$. And suppose that 'which she chased' is classified as the result of combining a displaced index-bearing expression with the very sentential expression from which it was displaced.

$$[\text{which}_2 \, [\ldots \, [\text{she}_{N1} \, [\text{chase}_V \, \cancel{\text{which}}_2]_V]_V\ldots]_S]_2$$

The embedded sentential expression can be treated as an instruction to construct the T-concept indicated below, which can be abbreviated as **1CHASE2**.

$$\uparrow•[\exists•[\text{EXTERNAL}(E, X), \text{FIRST}(X)], •[\text{CHASE}(E), \exists•[\text{INTERNAL}(E, X), \text{SECOND}(X)]]]]$$

Though instead of ignoring gender for simplicity, one can add that 'she' imposes a further constraint on external participants. The displaced wh-expression can also be treated as instruction to fetch a concept that applies a further restriction; 'who' plausibly adds a restriction to people. And crucially, one can treat the double-occurrence of the index as an instruction to invoke the Tarski trick, focusing on that index; cp. Heim and Kratzer (1998).

$$\text{SEM}([\text{which}_2 \, [\ldots]_S]_2) = +[\text{SEM}(\text{which}_2), \textbf{ADAPT}[\text{SEM}([\ldots]_S), 2]$$
$$= +[\text{SEM}(\text{fetch@which}), \textbf{TARSKI}\{2, \text{SEM}([\ldots]_S)\}$$

Recall that for any index i and T-concept **P**, **TARSKI**$\{i, \textbf{P}\}$ is the semantic concept below.

$$\exists\textbf{A}^*:\textbf{A}^*\approx_1\textbf{A}\{\text{ASSIGNS}(\textbf{A}^*, X, 1) \, \& \, \text{SATISFIES}(\textbf{A}^*, \textbf{P})\}$$

And the idea is that (the SEM of) 'which she chased' directs construction of a concept like •[$\text{ENTITY}(X)$, **TARSKI**$\{2, \textbf{1CHASE2}\}$], which applies to one or more things iff they were chased by whatever **A** assigns to the first index. Similarly, 'which chased her'

$$[\text{which}_1 \, [\ldots \, [\cancel{\text{which}}_1 \, [\text{chase}_V \, \text{her}_{N2}]_V]_V\ldots]_S]_1$$

can be analyzed as an instruction whose execution leads to construction of a concept that applies to one or more things iff they chased whatever **A** assigns to the second index.[16]

3.2 Quantification

We're finally in a position to describe the meanings of quantificational constructions like 'Every

cow arrived' and 'She chased every cow', by recasting the proposal in Pietroski (2005, 2006) explicitly in terms of instructions to build concepts, without appeal to truth values.

The central idea is simple: a determiner like 'every' fetches a number neutral concept concept of ordered pairs; where the ordered pair <x, y> can be identified with {x, {x, y}}, with x as its "external participant," and y as its "internal participant." More specifically, let's say that EVERY(O) applies to some ordered pairs iff: every one of their internal participants is one of their external participants; or put another way, (all of) their internals are among their externals. Likewise, MOST/THREE/SOME/NO(O) applies to some ordered pairs iff most/three/some/none of their internals are among their externals. And let's say that for any concept Φ(X), the concept MAX-Φ(X) applies to some things iff they are *the* things to which Φ(X) applies: MAX-Φ(X) ≡ ∀Y:Φ(Y)[AMONG(Y, X)].[17] Then •[EVERY(O), ∃•[INTERNAL(O, X), MAX-COW(X)]] applies to some ordered pairs iff their internals are the cows, and each of their internals is one of their externals.

We can say that from a semantic perspective, being an argument of a determiner differs slightly from being an argument of a verb, in that the former imposes a maximization condition.

D-INTERNALIZE:Φ(X) = **INTERNALIZE**:MAX-Φ(X)
➔ ∃•[INTERNAL(O, X), MAX-Φ(X)]

D-EXTERNALIZE:Φ(X) = **EXTERNALIZE**:MAX-Φ(X)
➔ ∃•[EXTERNAL(O, X), MAX-Φ(X)]

And given a concept Ψ(X) that applies to one or more things iff they arrived, the concept •[•[EVERY(O), ∃•[INTERNAL(O, X), MAX-COW(X)]], ∃•[EXTERNAL(O, X), MAX-Ψ(X)]] applies to one or more ordered pairs iff: their internals are the cows, their internals are among their externals, and their externals are the things that arrived. This concept applies to one or more things iff every cow arrived, assuming that ordered pairs exist if their participants/elements do. Likewise, •[•[EVERY(O), ∃•[INTERNAL(O, X), MAX-COW(X)]], ∃•[EXTERNAL(O, X), MAX-•[COW(X), Ψ(X)]] applies to one or more things iff every cow is a cow that arrived.

I mention the possibility of restricting the externals to cows that arrived because this may be relevant to the *conservativity* of determiners—see Barwise and Cooper (1981), Higginbotham and May (1981)—and the ways in which external arguments of determiners *differ* from relative clauses; see Pietroski (2005a) for further discussion. It is easy to construct a concept of those that arrived, given a suitable T-concept and quantification over assignment variants.

MAX-∃**A***:**A***≈$_1$ **A**{Assigns(**A***, X, **1**) &
Satisfies{**A***, ↑•[ARRIVED(E), ∃•[INTERNAL(E, X), FIRST(X)]]}}

This concept just is MAX-**TARSKI**{1, ↑•[ARRIVED(E), ∃•[INTERNAL(E, X), FIRST(X)]]}. But given restricted quantifiers, we can severely restrict the appeal to assignment variants. Let's say that for any assignments **A** and **A***, and any index *i*, **A***⊆$_i$ **A** iff: **A*** differs from **A** at most in that **A*** does not assign to *i* everything that **A** assigns to *i*; whatever **A** assigns to *i*, **A*** assigns *one or more but perhaps not all* of those things to *i*. Given an assignment that assigns (all and only) the cows to the first index, the concept indicated below is a concept of those cows that arrived.

MAX-∃**A***:**A***⊆$_1$ **A**{Assigns(**A***, X, 1) &
Satisfies{**A***, ↑•[ARRIVED(E), ∃•[INTERNAL(E, X), FIRST(X)]]}}

And we can define **REDUCEDTARSKI**{*i*, **P**} as follows.

∃**A***:**A***⊆$_i$ **A**{ASSIGNS(**A***, X, *i*) & SATISFIES(**A***, **P**)}

Let me conclude by showing how the constituents of a quantificational expression can be

instructions to build the requisite monadic concepts. As is standard within MP, I assume some version of the syntax shown below for 'She chased every cow'.

$$[[\text{every}_{D2}\ \text{cow}_N]_{D2}\ [\ldots\ [\text{she}_{D1}\ [\text{chase}_V\ \cancel{[\text{every}_{D2}\ \text{cow}_N]_{D2}}]_V]_V\ \ldots\text{s}]_{D2}]$$

For whatever reason—perhaps because 'every' needs an external argument—a copy of the indexed quantifier combines with the basic sentential expression, which then becomes the external argument of the quantifier. If original/lower copy is interpreted as an instruction to fetch the concept SECOND(X), perhaps because that is the only coherent interpretation available, then the embedded sentential expression is an instruction to construction a T-concept like **1CHASE2**. But the whole I-expression, headed by every$_{D2}$, is the following instruction:

$+[\text{SEM}([\text{every}_{D2}\ \text{cow}_N]_{D2}),\ \textbf{ADAPT}\{\text{SEM}([\ldots\ [\text{she}_{D1}\ [\text{chase}_V\ \cancel{[\text{every}_{D2}\ \text{cow}_N]_{D2}}]_V]_V\ \ldots\text{s}],\ D2\}.$

This is an instruction to conjoin the concepts obtained by executing two subinstructions: $+[\text{SEM}(\text{every}_{D2}),\ \textbf{ADAPT}\{\text{SEM}(\text{cow}_N),\ D2\}$ and $\textbf{ADAPT}\{\textbf{1CHASE2},\ D2\}$. The first subinstruction calls for conjunction of concepts obtained by (a) executing the indexed determiner instruction and (b) adapting a concept fetched with cow$_N$, in the way specified by classifying a noun as the internal argument of an indexed determiner. The second subinstruction calls for adapting **1CHASE2**, in the way specified by marking a sentential expression as the external argument of an indexed determiner. So one obvious hypothesis is given below.

SEM(every$_{D2}$) = fetch@every$_D$ ➔ EVERY(O)

ADAPT{SEM($\ldots_N$), D2} = **D-INTERNALIZE**:SEM($\ldots_N$)

ADAPT{SEM($\ldots_S$), D2} = **D-EXTERNALIZE**:**TARSKI**{2, SEM($\ldots_S$)}

This hypothesis has the desired consequences, assuming that every$_D$ fetches EVERY(O).

SEM([every$_{D2}$ cow$_N$]$_{D2}$) ➔ •[EVERY(O), ∃•[INTERNAL(O, X), MAX-COW(X)]]

ADAPT{SEM([… [she$_{D1}$ [chase$_V$ $\cancel{[\text{every}_{D2}\ \text{cow}_N]_{D2}}$]$_V$]$_V$ …s], D2} ➔

□•[EXTERNAL(O, X), MAX-**TARSKI**{2, **1CHASE2**}(X)]]

Conjoining the resulting concepts yields a concept of ordered pairs that meet three conditions: their internals are among their externals; their internals are the cows; and their externals are those things chased by whatever is assigned to the first index. And there are one or more such ordered pairs iff whatever is assigned to the first index chased every cow. So the corresponding T-concept can be the external argument of another determiner.

↑•[•[EVERY(O), ∃•[INTERNAL(O, X), MAX-COW(X)]],

∃•[EXTERNAL(O, X), MAX-**TARSKI**{2, **1CHASE2**}(X)]]

Alternatively, one can hypothesize that [every$_{D2}$ cow$_N$]$_{D2}$ requires that (all and only) the cows be assigned to the second index. Then one could replace appeal to **TARSKI**—in the rule for external arguments of determiners—with appeal to **REDUCEDTARSKI**.

ADAPT{SEM($\ldots_S$), D2}= **D-EXTERNALIZE**:**REDUCEDTARSKI**{2, SEM($\ldots_S$)}

There are various ways to build in the restriction. But one possibility is that the determiner itself is understood as a reflection of a restricted quantifier.

SEM(every$_{D2}$) = +[fetch@every$_D$, **D-INTERNALIZE**:SEM(2)]

➔ •[EVERY(O), ∃•[INTERNAL(O, X), MAX-SECOND(X)]]

This effectively treats the *index* as the internal argument of the determiner. So one might well look for additional syntax; cp. Larson (forthcoming). Then one might say either that the noun cow$_N$ is *also* understood as specifying the internal participants, with the consequence that the cows must be the things being tracked by the first index, or that adapting a noun to an *indexed* determiner just is a way of letting the index track the concept fetched with the noun.

$$\text{ADAPT}\{\text{SEM(cow}_N), \text{D2}\} = \text{ASSIGN}\{2, \text{MAX-COW(X)}\}$$

Any such account highlights the analogy between external arguments of determiners and relative clauses; cp. Heim and Kratzer (1998). But it does not treat the arguments of determiners as expressions of type $<e, t>$. Hence, the proposal here does not predict that relative clauses can be understood as external arguments of determiners. And indeed, 'Every cow which Aggie chased' has no sentential reading according to which every cow is such that Aggie chased it. But in a relative clause, the index of the displaced relativizer invokes the Tarski trick. The index of a displaced determiner phrase may do the same; or it may invoke a more restricted trick that does not require consideration of any *new* values of the variable in question. But in any case, we need not suppose that the external arguments of determiners are sentential expressions that combine with covert relativizers, given the option of invoking **TARSKI** or **REDUCEDTARSKI** as part of the hypothesized significance of being an external argument of a determiner.

On this view, certain aspects of phrasal syntax are correlated with significant adjustments of the concepts fetched or assembled via the constituent expressions. One can call this a kind of type-shifting even if there are no types to shift. But if a Conjunctivist semantics can handle quantificational constructions by appealing to simple operations like **INTERNALIZE** and **EXTERNALIZE**, given a maximalizing operator and **REDUCEDTARSKI**, then it is hard to argue that such constructions support appeals to saturation—as opposed to conjunction, **INTERNALIZE** and **EXTERNALIZE**—in the semantics of subsentential constructions.

4. Conclusion

An unsurprising pattern emerges from this exercise. If one adheres to the idea that combining expressions is fundamentally an instruction to construct *conjunctive* concepts, along with the idea that open class lexical items are instructions to fetch concepts with independent content, one is led to say that certain aspects of syntax and various functional items are instructions to *convert* fetchable/constructable concepts into concepts that can be systematically conjoined with others. Perhaps this is the *raison d'etre* of syntax that goes beyond mere recursive concatenation: grammatical relations, like being the internal/external argument of a verb or determiner, can carry a kind of significance that is intriguingly like the kind of significance that prepositions have. These old ideas can be combined in a Minimalist setting devoted to asking which conversion operations are required by a spare conception of the recursive composition operations that HFL can invoke in directing concept assembly. The list of operations surveyed here is surely both empirically inadequate, and yet already too rich. My aim has been to offer a specific proposal as one illustration of Minimalist thinking in semantics, guided by two thoughts: this kind of inquiry has been fruitful in studying I-language syntax; and the study of I-language semantics has the same target of inquiry if I-expressions are instructions to build concepts.[18]

References

Baker, M. 1997: Thematic Roles and Grammatical Categories. In L. Haegeman (ed.), *Elements of Grammar*, Dordrecht: Kluwer, pp. 73-137.

—2003: *Verbs, Nouns, and Adjectives*. Cambridge: Cambridge University Press.

Barwise, J. and Cooper, R. 1981: Generalized Quantifiers and Natural Language. *Linguistics and Philosophy* 4:159-219.

Boeckx, C. 2008: *Bare Syntax*. Oxford: Oxford University Press.

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

Borer, H. 2005: *Structuring Sense* (volumes I and II) Oxford: Oxford 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.

Chametsky, R. 1996: *A Theory of Phrase Markers and the Extended Base*. Albany: SUNY Press.

Chierchia, G. 1998: Reference to Kinds. *Natural Language Semantics* 6: 339–405, 1998.C

Chomsky, N. 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: *The Minimalist Program*. Cambridge, MA: MIT Press.

—2000a: Minimalist Inquiries. In R. Martin et.al. (eds.) *Step by Step:*
    Essays in Honor of Howard Lasnik (Cambridge, MA: MIT Press).

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

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

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

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

Di Sciullo, A. and Aguero, C. (eds.), 2010: *Biolinguistic Investigations.* Cambridge, MA: MIT.

Dowty, D. 1979: *Word Meaning and Montague Grammar*. Boston: Reidel.

—1991: Thematic Proto-roles and Argument Selection. *Language* 67:547-619.

Elbourne, P. 2005: *Situations and Individuals*. Cambridge, MA: MIT.

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).

Fodor, J. 1975: *The Language of Thought*. New York: Crowell.

—1986: *Psychosemantics*. Cambridge, MA: MIT.

—2003: *Hume Variations*. Oxford: OUP.

Fox, D. 1999: *Economy and Semantic Interpretation.* Cambridge, MA: MIT Press.

Frege, G. 1884: Die Grundlagen der Arithmetik. Breslau: Wilhelm Koebner. English translation
    in J. L. Austin (trans.), *The Foundations of Arithmetic* (Oxford: Basil Blackwell, 1974).

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

Gillon, B. 1987: Readings of Plural Noun Phrases in English.
    Linguistics and Philosophy 102:199-219.

Halle, M. and Marantz, A. 1993: Distributed Morphology and the Pieces of Inflection.
    In K. Hale and S. Jay Keyser (eds.), *The View from Building 20,*
    Essays in Honor of Sylvain Bromberger, (Cambridge, MA: MIT Press).

Hauser, M., Chomsky, N., and Fitch, T. 2002: The faculty of language: what is it,
    who has it, and how did it evolve? *Science* 298: 1569-1579.

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

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

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

Hornstein, N. 2009: *A Theory of Syntax: Minimal Operations and UG*. Cambridge: CUP.

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.

Hurford, J. 2007: *The Origins of Meaning*. Oxford: OUP.

Jackendoff, R. 1990: *Semantic Structures*. Cambridge, MA: MIT Press.

—2002: *Foundations of Language*. Oxford: Oxford University Press.

Jacobson, P. 1999: Variable Free Semantics. *Linguistics and Philosophy* 22:117-84.

Kamp, H. 1975: Two Theories About Adjectives. In E. Keenan (ed.), *Formal Semantics of Natural Languages*, Cambridge University Press.

Katz, J. and Fodor, J. 1963: The Structure of a Semantic Theory. *Language* 39: 170-210.

Katz, J. 1994: Names Without Bearers. *Philosophical Review* 103: 1-39.

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.

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

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).

Longobardi, G. 1994: Reference and Proper Names, *Linguistic Inquiry* 25: 609-65.

Ludlow, P. 2002: Natural Logic and LF. In In Preyer and Peter (2002).

Marantz, A. 1984: *On the Nature of Grammatical Relations*. Cambridge, MA: MIT Press.

Marr, D. 1982: *Vision*. San Francisco: Freeman and Company.

May, R. 1985: *Logical Form: Its Structure and Derivation*. Cambridge, MA: MIT Press.

Montague, R. 1974: *Formal Philosophy*. New Haven: Yale University Press.

Parsons, Terence. 1970: Some problems concerning the logic of grammatical modifiers. *Synthese* 21:320-334.

—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.

Partee, B. and Rooth, M. 1983: Generalized Conjunction and Type Ambiguity. In R. Bäuerle *et. al.* (eds.), *Meaning, Use, and Interpretation of Language* (Berlin: de Gruyter).

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).

—2006a: Interpreting Concatenation and Concatenates. *Philosophical Issues* 16: 221-45.

—2008: Minimalist Meaning, Internalist Interpretation. *Biolinguistics* 4: 317-41.

—2010: Concepts, Meanings, and Truth: First Nature, Second Nature, and Hard Work. *Mind and Language* (to appear).

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: Oxford University Press.

Putnam. H. 1975: The Meaning of 'Meaning'. In K. Gunderson (ed.),
    Language, Mind and Knowledge (University of Minnesota Press).

Pylyshyn, Z. 1997: Things and Places. Cambridge, MA: MIT Press.

Ramchand, G. 2008: Verb Meaning and the Lexicon: a first phase syntax. Cambridge: CUP.

Quine, W. 1963: On What There Is. In *From a Logical Point of View*.
    New York: Harper and Row.

Salmon, N. 2006: A Theory of Bondage. *Philosophical Review* 115:415-448.

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.

Schwartzschild, R. (1996): *Pluralities*. Dordecht: Kluwer.

Strawson, P. 1950: On Referring. *Mind* 59: 320-44.

Svenonius, P. forthcoming: Spatial P in English. In G. Cinque and L. Rizzi, eds,
    The Cartography of Syntactic Structures, vol. 6 (Oxford: Oxford University Press).

Szabo, Z. 2000: Compositionality as supervenience. *Linguistics and Philosophy* 23: 475–505.

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).

Uriagereka, J. (1999). Multiple Spell-out. In S. Epstein and N. Hornstein (eds.),
    Working Minimalism (Cambridge, MA: MIT Press).

Vendler, Z. 1959: Verbs and Times. *Philosophical Review*: 66: 143-60.

[1] I understand MP broadly, not merely as an attempt to simplify extant conceptions of syntax; see note 2. But this paper is not a review of the valuable literature that bears on attempts to simplify accounts of the "syntax-semantics interface;" see, for various examples, Fox (1999), Borer (2005), Jackendoff (2002), Ramchand (2008). The focus here is on composition operations; cp. Hornstein and Pietroski (2009). Pietroski (2010) offers independent arguments for the view on offer, while exploring the implications for truth and the concepts that interface with HFL.

[2] For present purposes, I take it as given that humans have a faculty of language. But other things equal, one wants to posit as little as possible—especially in terms of distinctively human capacities—in order to describe and explain the linguistic metamorphosis that children undergo; cp. Hauser, Chomsky, and Fitch (2002), Hurford (2007). This bolsters the general methodological motivation, already strong, to simplify descriptions of the states of linguistic competence that children acquire; cp. Hornstein (2009). If such competence includes knowing which meanings a given PHON can have (see note 3), then in evaluating attempts to simplify any other aspects of competence, we must consider implications for the semantic properties of expressions (cp. Hornstein and Pietroski [2009]) and representations that interface with HFL in ways that let humans use this faculty as we do. Chomsky (1995) argued, in particular, that the expressions generated by HFL just *are* PHON-SEM pairs. My proposal does not require this very spare conception of expressions. But if expressions have further ("purely syntactic") properties, that only amplifies the motivations for a spare conception of how SEMs are related to concepts.

[3] There are, however, many ways in which speakers *don't* compute interpretations. This is one moral of many "poverty of stimulus" arguments, based on observations concerning (i) which sentences *imply* which, and (ii) logically possible interpretations that certain word-strings *cannot* support. See, e.g., Higginbotham (1985), drawing on Chomsky (1965). For reviews of some relevant psycholinguistic work, see Crain and Pietroski (2001).

[4] Cp. Marr (1982), Evans (1981), Peacocke (1986), Davies (1987), Pietroski *et.al.* (2009). Given that implementation matters, it seems obvious that explanations in this domain can and should be framed within a "biolinguistic" framework; see DiSciullo and Aguero (forthcoming). Correlatively, we don't merely want theories that respect generic compositionality principles like the following: the meaning of expression α is determined by α's syntactic structure and the meanings of α's constituents. If the actual composition operations reflect innate aspects of human cognition, generic principles will be respected by languages that no child could acquire. In this sense, mere compositionality is multiply realizable (see Szabo 2000), raising the question of how it is realized in human I-languages; cp. Hurford (2007).

[5] Cp. Katz and Fodor (1963). One can still say that each concept has an extension in each context, and that in this sense, I-expressions link sounds to extensions. But if a lexical item L is polysemously linked to more than one concept, then an instruction to fetch a concept linked to L is fulfilled by fetching any concept linked to L—much as an instruction to fetch a rabbit from a

room with rabbits is fulfilled by fetching any rabbit from the room. Though I have nothing to say about where polysemy ends and homophony begins.

[6] Perhaps some I-languages count as idiolects of English only if they are adequate tools for communication among certain people (including us). In which case, some I-languages may so count only if their lexical items are used to fetch concepts that are "extensionally similar" in roughly the following sense: there is suitable overlap with regard to what the relevant concepts apply to; and for purposes of communication, disparities can be resolved or ignored.

[7] This would provide at least the start of an explanation for why 'France is a hexagonal republic' is defective in a way that 'France is hexagonal, and France is a republic' is not. See Pietroski (2005b), drawing heavily on Chomsky (1975, 2000b).

[8] See, e.g, Link (1983), Schwartzschild (1996). Letting 'π' range over plural entities and '$\in$' have its usual meaning: $\forall\pi\forall\pi^*\{(\pi = \pi^*) \equiv \forall x[(x \in \pi) \equiv (x \in \pi^*)]\}$.

[9] I grant that adverting to lattices, with basic entities as terminal nodes, can be illuminating in various ways; see e.g., Link (1983), Schwartzschild (1996), and Chierchia (1998). But instead of interpreting each nonterminal node as a potential assignment of exactly one entity with elements to plural variable, one can interpret each such node as a potential assignment of more than one entity to a number-neutal variable; see Pietroski (2006). More speculatively, one might hope to accommodate mass nouns like 'water' and 'wood' in terms of a variable that is neutral as between one-or-more things or "some stuff," with 'chop (some) wood' as an instruction to build •[CHOP(X), ∃•[PATIENT(E, X), WOOD(X)]]; where WOOD(X) applies, mass/count-neutrally, to (any sample of) wood. One could then distinguish PIZZA(X) from •[PIZZA(X), COUNTABLE(X)], or •[PIZZA(X), COUNTABLE-AS-(X, PIZZA)], to distinguish 'ate some pizza' from 'ate a pizza'; cp. Gillon (1987), Chierchia (1998).

[10] Or if you prefer, for any one or more things: the concept [↑Φ(X)](Z) applies to them iff Φ(X) applies to one or more things; and [↓Φ(X)](Z) applies to them iff Φ(X) applies to nothing. But omitting the extra brackets and variable position turns out to be at least as perspicuous.

[11] See Ludlow (2002) for discussion in the context of the "natural logic" tradition as updated by modern conceptions of grammar, with particular attention to negative polarity facts.

[12] Partee (2006) raises the same kind of question, against a different background, though with less suggestion that a typology-free semantics might work.

[13] Think of an assignment as assigning one or more things to the free conceptual variable 'X' and one or more things to each index in the SEM. Other dependencies on assignments can be encoded in familiar ways, modulo number-neutrality. Larson and Segal's (1995) treatment is especially friendly to any Tarski-inspired theory.

[14] Nothing hangs on labeling names with 'N' and ignoring an internal structure, as opposed to $[\varnothing_D\ \text{Bessie}_N]_D$, with a covert determiner and a lexical proper noun.

[15] For example, CHASE might be marked as an "action" concept in a system that (for purposes of interfacing with SEMs) represents the agents/patients of actions as their external/internal participants; see Pietroski (2005a 2008) for related discussion drawing on Baker (1997).

[16] As Heim and Kratzer's (1998) system nicely highlights, even if one appeals to saturation as a composition operation, one still needs to posit Tarskian abstraction—often encoded as lambda abstraction (after Church [1941])—as a distinct operation. And this so, even given a Fregean typology. Suppose we treat indices and traces as constituents, as in $[2\ [\text{she}_1\ \text{chased}_V\ t_2]_S]$, with the embedded sentence as an expression of type $\langle t \rangle$ and the larger expression as of type $\langle e, t \rangle$. From an I-language perspective, one can say (modulo tense and gender) that relative to any assignment **A**: the concept formed by executing $[\text{she}_1\ \text{chased}\ t_2]_S$ denotes truth iff whatever **A** assigns to 1 chased whatever **A** assigns to 2; and correlatively, the concept formed be executing $[2\ [\text{she}_1\ \text{chased}_V\ t_2]_S]$ applies to x iff whatever **A** assigns to 1 chased x. But the idea *isn't and can't be* that the index denotes a function-in-extension of type $\langle t, \langle e, t \rangle\rangle$, which maps the truth value of $[\text{she}_1\ \text{chased}\ t_2]$ onto a function of type $\langle e, t \rangle$. Rather, '2' has to indicate a hypothesized (syncategorematic) instruction to convert a *representation* of one sort into a *representation* of another sort. Heim and Kratzer's third composition rule, in addition to rules for saturation and conjunction, makes this vivid. The attractive idea is that the higher copy of the lower index triggers quantification over assignment variants (taking assignments to be Tarskian sequences):

$$\|[2\!\wedge\![\text{she}_1\ \text{chased}_V\ t_2]_S]\|^{\mathbf{A}} =$$
$$\lambda x.\ \text{T iff } \exists \mathbf{A'}\!:\!\mathbf{A'}\!\approx_2\!\mathbf{A}[(x = \mathbf{A'}2)\ \&\ \|[\text{she}_1\ \text{chased}_V\ t_2]_S]\|^{\mathbf{A'}} = \text{T}.$$

This has the desired result, taking the lambda-expression to be a theorist's representation of the hypothesized *concept* obtained in two stages: execute the sentential instruction, obtaining a concept that is doubly sequence-sensitive, and modify the resulting concept as directed by '2'. One can remain agnostic about the detailed forms of the concepts constructed. And from an E-language perspective, one can take the lambda-expression to be (only) a theorist's representation of the hypothesized *satisfaction condition*. But from an I-language perspective, the goal is to say how competent speakers represent the alleged satisfaction condition. And while theorists can abbreviate—as in $\|[2\!\wedge\![\text{she}_1\ \text{chased}_V\ t_2]_S]\|^{\mathbf{A}} = \lambda x.\text{CHASE}(\mathbf{1, X})$—we should remember that the corresponding psychological hypothesis presupposes some version of the Tarski trick. I just want to make such appeal explicit, so that we can ask what *other* mental machinery we need to posit in accounts of how I-concepts are constructed. For many purposes, it is fine to use a notation that effectively mixes appeals to saturation and abstraction. But this makes it harder, though by no means impossible, to see which operation does what work where.

[17] So given some but not all of the things a concept applies to, the "maximized" concept does not apply to them. The number-neutral 'MAX-Φ(X)' can be cashed out with a first-order variable: $\forall x[X(x) \equiv \Phi(x)]$. But this says that for each domain entity, it is one of the one or more Xs iff it meets a certain condition. It *doesn't* say that there is a set **s** such that for each domain entity, it is an element of **s** iff it meets the condition: $\forall x[(x \in \mathbf{s}) \equiv \Phi(x)]$. Suppose the domain entities are all

and only the Zermelo-Frankl (ZF) sets. Then there *are* one or more entities (viz., the ZF sets) such that each entity is one of them iff it is nonselfelemental; but there is no set whose elements are these entities. And a concept of "being among" (or inclusion) could be used to introduce a concept of ordered pairs: EVERY(O) $\equiv$ ιY:INTERNAL(O, Y){ιY:EXTERNAL(O, X)[AMONG(Y, X)]}; or in first-order/singular terms, $\forall$o:$\exists$x[Oo & Internal(o, x)]{$\exists$p[Op & External(p, x)].