

Minimalist Meaning, Internalist Interpretation

Paul M. Pietroski

This article offers a conception of semantics, and of what makes the human language faculty distinctive, based on five theses: Meanings are instructions to build concepts; concatenation calls for conjunction of monadic concepts; grammatical relations invoke certain thematic relations and a kind of existential closure; lexicalization is a partly creative process of abstraction; and meanings are internalistic properties of expressions. Each of these claims is defended elsewhere. The aim here is to connect them explicitly, and compare the result with alternatives, in the hope of providing a plausible conception of natural language meaning that coheres with Chomsky's minimalist program.

Keywords: concepts, internalism, lexicalization, meanings, semantics

1. Proposal: Meanings as Conjunctive *Begriffsplans*

This paper combines five theses, each discussed in more detail elsewhere (see Pietroski 2002 *et seq.*).

- (1) Meanings are instructions to build concepts.
- (2) Concatenation calls for conjunction of monadic concepts.
- (3) Grammatical relations invoke certain thematic relations and a restricted operation of existential closure (\exists -closure).
- (4) Lexicalization is a partly creative process of abstraction.
- (5) Meanings are internalistic properties of expressions.

These claims, explained below, are logically independent. But they cohere, inviting comparison with alternative quintets. The net result is a version of semantic internalism: Open-class lexical items are instructions to fetch monadic concepts that may have been abstracted in the course of acquisition; and the meaning of a phrase is an instruction to build a conjunctive monadic concept from fetchable elements, given a few relational/thematic concepts and an operation of existential closure. The position identified owes much to Chomsky's (1977, 1995a, 2000a) discussion of meaning and the general program of biolinguistics.

For helpful comments and discussion, my thanks to Cedric Boeckx, Susan Dwyer, Norbert Hornstein, Terje Lohndal, Dennis Ott, and two anonymous referees.



1.1. *Meanings are Instructions to Build Concepts*

For these purposes, I take it as given that humans have a language faculty that is in some respects distinctively human, and that the languages human children naturally acquire can be described as biologically implemented *procedures* that generate expressions. Chomsky (1986) calls these procedures, which correspond to stable states of the faculty, *I-languages*. By contrast, E-languages are sets of expressions; even for languages with endlessly many expressions, a single E-language might be determined by two or more generative procedures. The ‘I’/‘E’ distinction connotes the contrast between intensions (procedures, algorithms) and the extensions we characterize by appeal to procedures, as when we specify sets without listing their elements; cf. Church (1941).

Using this terminology, we can describe the Human Faculty of Language (HFL) as a biologically implemented capacity to acquire and use one or more I-languages that associate phonological instructions to articulatory/perceptual systems with semantic instructions to conceptual/intentional systems, by means of a constrained syntax; cf. Chomsky (1995b, 2000b). Abstracting from phonology, expressions of an I-language *are* semantic instructions.² These expressions have various grammatical properties of interest to syntacticians. But whatever additional properties they have, expressions of an I-language (henceforth, *i-expressions*) are presumably instructions to access certain mental representations and generate others, in accord with certain principles of lexicalization and composition. If only for simplicity, let’s say that the representations accessed and generated are concepts — or more specifically, ‘*i-concepts*’. I assume that concepts are *composable* mental representations; see e.g. Fodor (1987, 2003, 2008). But this is compatible with many proposals about how *i-concepts* are related to the representations that children lexicalize and the full range of representations available to human thinkers.

If some of the representations that children lexicalize are not conceptual — think of mental images, maps, and prototypes — then lexicalization must somehow associate these pre-lexical representations with concepts. But following Fodor and others, I suspect that children typically lexicalize concepts, many of which are not uniquely human. This allows for the hypothesis that lexicalization is fundamentally a matter of concept labeling, and that *i-concepts* just *are* lexicalized concepts; cf. Bloom (2000). Alternatively, one can hypothesize that lexicalized concepts are linked to formally distinct *i-concepts* that *are abstracted* in the course of lexicalization, and that *i-concepts* form a special subset of the concepts available to humans; see section 1.4 below. But whatever the details, the broad idea is one that many theorists should be able to accept: Lexical items are instructions to *fetch* concepts that meet certain conditions, while phrases are instructions to *combine* fetchable concepts in certain ways.

From this perspective, *i-expressions* are concept-construction-instructions (*Begriffsplans*) that reflect principles governing combination of *i-expressions* and

² Likewise, abstracting from semantics, expressions of an I-language are phonological instructions. The point is not to privilege semantics. And even if instructions to articulatory/perceptible systems appeared later, in terms of HFL’s evolutionary history, ontogeny may not recapitulate phylogeny in this respect.

'interfaces' between HFL and other aspects of human cognition. Correlatively, *i*-concepts are concepts that can interface with HFL: *i*-concepts are results of executing instructions (*i*-expressions) that are generated by I-languages; and for a given context, an *i*-concept can be expressed with an *i*-expression whose execution would create the *i*-concept.

Thesis (1) is to be understood, accordingly, as a claim about *i*-expressions.

(1) Meanings are instructions to build concepts.

If *i*-expressions pair phonological instructions with semantic instructions, then *i*-expression meanings can be described as instructions (generable via HFL) to fetch and/or combine *i*-concepts in certain ways. One can stipulate that meanings are not semantic instructions in this sense. But then it is hardly obvious that there are theories of meaning for I-languages, much less that such theories perspicuously describe the natural phenomenon of human linguistic understanding. Put another way, one can stipulate that *i*-expressions pair phonological instructions with *i*-meanings and hypothesize that theories of the natural phenomenon are theories of *i*-meanings. On this view, endorsed here, understanding an *i*-expression (perceiving its meaning) is a matter of recognizing that expression *as* a certain concept–construction–instruction. This proposal may be wrong. But it is less tendentious than it might initially appear. So let me distinguish (1) from some alternatives in the vicinity.

It is an ancient idea that lexical items *label* concepts, and that if a lexical item λ labels a concept *C* for a speaker *S*, then *S* can use λ to talk about whatever *S* thinks about by using *C*. I prefer to say that lexical items *are instructions to fetch* concepts, if only because (i) a polysemous item like *book* or *set* may be lexically linked to more than one fetchable concept, and (ii) as we'll see in section 1.4, lexical items may fetch concepts that are formally distinct from the concepts initially lexicalized. Though given enough caveats, one can say that lexical *i*-expressions (as used in contexts) indicate the concepts they normally fetch.

Some theorists prefer to say that words like *rabbit* and *table* indicate things like rabbits and tables, as opposed to concepts. In many cases, this is harmless. Predicates like *unicorn* and names like *Vulcan* present familiar worries. But one can maintain that words are 'normally' used to indicate 'real' things, given a suitably restricted notion of normal use, so long as one isn't too demanding about what counts as real; cf. Chomsky (2000a), Pietroski (2005a), Hinzen (2007). In saying that *i*-expressions are *Begriffsplans*, my point is not to deny that speakers often use words to talk about language-independent things. On the contrary, my suggestion is that *i*-expressions are often used in this way *because* they call for construction of *i*-concepts, which are often constituents of thoughts that are (somehow) about mind-independent things. Since *i*-expressions are also associated with phonological instructions, they can be used to communicate thoughts. So in suitably controlled contexts, intuitions about the truth or falsity of a thought communicated with an *i*-expression can serve as useful data points for theories of *i*-meanings.³

³ Likewise, I grant that speakers can use *i*-expressions to make assertions whose contents can

As these concessive remarks suggest, (1) is fully compatible with psychologized versions of Truth Conditional Semantics, according to which *i*-expressions are instructions to construct concepts that have Tarski-style satisfaction conditions.⁴ If semantic instructions are individuated in this externalistic fashion — if the instructions require construction of concepts with certain *truth-theoretic* properties — then thesis (5) is false.

(5) Meanings are internalistic properties of expressions.

But (1) can be combined with the hypothesis that each *i*-expression not only has a Tarskian satisfaction condition, it has that satisfaction condition essentially.

That said, (1) is also compatible with a claim that is more friendly to internalism about meaning/understanding: *i*-expressions are instructions to construct *i*-concepts, some of which can be refined and used (in contexts where truth matters) to form thoughts that count as true or false because they are (modulo some complications like vagueness) sufficiently like ‘Ideal Thoughts’ whose constituent concepts really do have Tarski-style satisfaction conditions. This claim, which remains agnostic about the nature of semantic instructions and concepts constructed, is compatible with (5). So while (1) does not presuppose (5), (1) may be offered as part of a package that includes (5).

In subsequent sections, it will be important to be clear about what (1) does and does not imply. So let me stress these points now. Thesis (1) is incompatible with ‘spare’ theories according to which *i*-expressions have *no* semantic properties other than Tarskian satisfaction conditions.⁵ But accepting (1), and describing

be characterized in terms of sets of possible worlds. When a suitably clear thought is asserted, we can talk about the worlds in which that thought is true; and for any suitably clear concept, we can talk about its extension at a world. But it does not follow that *i*-expressions have intensions, much less that the meaning of an *i*-expression is its intension. One can say instead that *i*-expressions are used to construct concepts that have (context-dependent) intensions; where these concepts can be combined to form thoughts whose contents are determined by the intensions of the constituent concepts.

⁴ See e.g. Larson & Segal (1995). This is not what Davidson (1967), Montague (1974), or Lewis (1972) proposed. But it is in the spirit of Harman (1970) and Partee (2004). A variant proposal is that *i*-expressions are instructions to construct concepts, and that the concepts constructed have satisfaction conditions, even though the instructions do not require this.

⁵ If we invent a Tarskian language in which ‘ Rx_1 ’ is satisfied by a sequence σ iff $\sigma(1)$ is a rabbit, then ‘ Rx_1 ’ can be used to talk about rabbits, even by speakers who do not associate ‘ Rx_1 ’ with any concept (much less a concept of rabbits as such). We might invent such a language, and stipulate that ‘ Rx_1 ’ has no other semantic properties, in order to guarantee that ‘ Rx_1 ’ has the same (concept-independent) interpretation for all users of the language. We can likewise stipulate that the logical constant ‘ p ’ is satisfied by σ iff $\sigma(p) = \text{Phosphorus}$, where $\sigma(p)$ is the element of σ corresponding to the p^{th} logical constant, and that ‘ h ’ is satisfied by σ iff $\sigma(h) = \text{Hesperus}$. Given the usual ancillary apparatus, it will follow that ‘ Rp ’ is true iff Phosphorus is a rabbit, and that ‘ Rh ’ is true iff Hesperus is a rabbit.

In general, expressions of the invented language can have (and be understood as having) truth-theoretic properties as their only semantic properties. Thesis (1) implies that I-languages do not have this spare Tarskian character. But to *define* a language whose semantics is concept-independent in this way, humans may need a prior I-language whose expressions are used as devices to fetch concepts that can be combined to form truth-evaluable thoughts.

the meanings of *i*-expressions as semantic instructions, leaves room for many conceptions of the relevant fulfillment conditions.

For example, a neo-Davidsonian might regard the untensed verb phrase *stab Caesar* as the following complex instruction: Fetch a singular concept of the individual Caesar; fetch a concept satisfied by ordered triples $\langle e, x, y \rangle$ such that e is a stabbing by x of y ; and saturate (the most internal variable of) the latter with the former, thereby forming a concept satisfied by ordered pairs $\langle e, x \rangle$ such that e is a stabbing by x of Caesar. On this view, the instruction *stab Caesar* is fulfilled — and in that sense, satisfied — by constructing any concept with the specified truth-theoretic profile. In which case, one might say that *stab Caesar* is itself satisfied by $\langle e, x \rangle$ iff e is a stabbing by x of Caesar. Likewise, one might say that *dog* is an instruction to fetch a concept satisfied by x iff x is a dog, and that the *i*-expression *dog* inherits this satisfaction condition.

There are, however, other coherent conceptions of semantic instructions and fulfillment. For each speaker, each lexical item might be an instruction to fetch a concept that has a certain ‘address’ in mental space — viz. the address of the concept that was lexicalized with the relevant phonological form (PF). Suppose that at least typically, a lexical item is acquired via some process in which: A PF is linked to a pre-lexical concept, perhaps initially by mere association, but eventually by pairing the concept with an address A such that (i) HFL can generate a lexical *i*-expression that links the PF to an instruction to fetch a concept paired with A , and (ii) the resulting *i*-expression can have certain additional features corresponding to grammatical idiosyncracies. If the concept lexicalized remains the only concept at address A , then fulfilling the instruction is always a matter of fetching that concept. But suppose that one or more additional concepts get assigned to that address. Perhaps a formally new concept is defined in terms of the concept lexicalized, and HFL does not assign a new address to a concept so formed. Or perhaps the cognitive processes underlying polysemy can result in a family of concepts having the same address, so far as HFL is concerned. In such cases, there may well be more than one way of fetching a concept from address A .⁶

Of course, the PF of *dog* could have been lexically linked to a concept of cats, while the PF of *cat* was linked to a concept of dogs. And we can imagine two individuals that differ only in this respect: In Oscar, the PFs are linked to concepts in this ‘inverted’ way; while in Felix, the PFs are linked to concepts in a ‘proper-English’ way. We can go on to say that Oscar’s I-language does not count as an idiolect of English. And it may well be that the normal way of acquiring an (I-language that counts as an) idiolect of English involves using the PF of *dog* to lexicalize a concept with which the lexicalizer can think about dogs. But even if the concept lexicalized has a doggish satisfaction condition, it does not follow

⁶ Suppose you tell me to fetch a box from a certain room. If I find four boxes in the room, I can fulfill your instruction by bringing back any one of the four. But if you also tell me to fetch a ball from another room, and put it in the box I fetch, then the sizes of the boxes/balls will impose constraints on which choices permit fulfillment of your tripartite instruction. In this context, see Pietroski (2005a) on the difference between *France is a hexagonal republic* and *France is a hexagonal*, and *France is a republic*; cf. Chomsky’s (1995a) discussion of *London* as a device for referring to different ‘things’ on different occasions of use.

that *dog* is an instruction to fetch a concept with that satisfaction condition, not even if the concept lexicalized is the concept fetched. An instruction to *fetch a concept from a given address* is not an instruction to *fetch a concept with a certain truth-theoretic property* — not even if every concept at the address has that truth-theoretic property.⁷

This last point will be important when we turn to thesis (5). But for now, the point is just that thesis (1) makes it possible to reject externalist conceptions of meaning while accepting many externalistic claims about concepts. Readers already inclined to agree, or not care, can skim through to section 1.2 below.

Let's grant that if a lexical item λ counts as an instance of the English word *dog*, then λ is lexically linked to a concept of dogs — i.e. a concept with which one can think about dogs, and think of them as such (cf. Fodor 2003). I am more skeptical of the following claim: λ is not an instance of the English word *dog* if λ is lexically linked to a concept that would apply to things of a kind biologically distinct from but superficially similar to the actual dogs. Yet even given this much externalism, along with the related idea that a concept of dogs needs to be a concept that is essentially tied to the biology of actual dogs, a lexical *i*-expression can count as an instance of the English word *dog* without being an instruction to fetch a concept of dogs. For an *i*-expression can be an instruction to fetch a concept from a certain address — an address that actually links the PF of *dog* to a concept of dogs, where being a dog is a biologically demanding condition — without being an instruction to fetch a concept of dogs, or even an instruction to fetch a concept that applies to dogs.

In terms of the much discussed 'twin earth' thought experiments, doppelgangers might execute the same address-focused instruction — fetch a concept from address A — yet thereby fetch concepts with different contents, at least if doppelgangers in different environments can have concepts with different contents.⁸ For just as we can imagine Oscar and Felix linking different concepts to the PF of *dog*, we can imagine Felix and Twin-Felix linking different concepts to the PF of *dog*, at least if Felix and Twin-Felix do indeed have different concepts. If the difference between dogs and cats is salient to Oscar and Felix, they may not link the PF of *dog* to the same address in mental space; in which case, Oscar and Felix may not have the very same I-language, even given an address-focused conception of semantic instructions. But if any differences between dogs and twin-dogs would go undetected by Felix/Twin-Felix, then presumably, these duplicates are type-identical with regard to how they link the PF of *dog* to an address in mental space. If this is correct, then so far as their shared HFL is concerned, the PF of *dog* is linked to the same address. In which case, Felix and Twin-Felix do have the same I-language, given an address-focused conception of

⁷ An instruction to fetch a box from a certain room differs from an instruction to fetch a red box, even if every box is both in that room and red.

⁸ For these purposes, we needn't worry about the differences between Putnam's (1975) scenarios — which need not correspond to possible worlds in Kripke's (1980) sense — and Burge's (1979) counterfactual situations that often hold the 'non-linguistic environment fixed. Extending Kaplan's (1988) notion of character to thought, Fodor (1987) famously appealed to 'narrow' contents; see Segal (2000) for helpful discussion. But whatever one says about concepts and *their* fulfillment conditions, the issues here concern *i*-expressions (hypothesized instructions to fetch and combine concepts).

semantic instructions.

Note that an address-focused instruction differs from any instruction to fetch a concept with a certain ‘narrow’ content, even if any concept fetched from that address indicates the same mapping from contexts to extensions. But my aim is not to eschew appeals to contents of any breadth. For all that I have said here, semantic instructions may have wide (but mind-dependent) contents that are individuated by the relevant addresses, or narrow contents, or both. More importantly, just as twins might execute the same lexical instruction, so they might execute the same phrasal instruction. If twins share the addresses A1 and A2, along with a binary operator ‘•’, each twin might fulfill the following instruction: Form an instance of ‘C1() • C2()’ by replacing ‘C1’ with a concept found at A1 and replacing ‘C2’ with a concept found at A2. If ‘•’ is a conjunction operator, which need not be invoked under a truth-theoretic guise, then twins who execute the instruction might form conjunctive concepts with different contents.⁹

On this view, two individuals can share an I-language and competently use the same lexical items to fetch concepts from the same addresses, even if (for whatever reason) the concepts fetched have *distinct* contents. My suspicion is that such cases are common, even without twin earth. It seems all too likely that speakers in the same conversation often share an *i*-expression (e.g. *water*), while severally using the expression to fetch concepts that differ extensionally in ways that are often but not always irrelevant — at least if we suppose that the concepts fetched via lexical items really do have extensions in contexts. Given two concepts that apply to the stuff in Lake Michigan, only one may apply to the stuff coming out of a certain tap; cf. Chomsky (2000a).

More generally, one shouldn’t insist that if two otherwise linguistically similar speakers each use a lexical item to fetch concept, a mere extensional difference in the concepts fetched guarantees that either (i) the lexical *instructions differ*, or (ii) each instruction is associated with the *same function* from contexts to extensions, given a substantive notion of context — e.g., a sequence of potential values for conceptual variables, as opposed to entire possible worlds or their ‘extension-determining’ aspects. Another option, considered but not developed by Putnam (1975), is to deny that meanings determine extensions in contexts.

The meaning of an *i*-expression may determine the general ‘shape’ of any concept constructed in accord with that semantic instruction, without determining a characteristic function (from contexts to contents) for any concept so constructed. One can invent a *Begriffsschrift* that employs lots of indices, and thereby confines the content variation in concepts constructed to values of variables; cf. Stanley (2000). But I-languages may care less about conceptual content. One can say, if one likes, that Twin-Felix fails to have an I-language that

⁹ This assumes that twin-concepts can have the same address, and hence that addresses can be individuated without reference to the satisfiers of addressed concepts. But if this is a substantive assumption, so is its negation. One can still say that each twin executes a demonstrative instruction, akin to “Fetch one of those,” with a distinctive ‘wide’ fulfillment condition (if each twin has his own concepts). My claim is not that it is incorrect to describe twin-earth cases externalistically, but rather, that theorists can — and perhaps should — describe such cases internalistically.

counts as an idiolect of English (and that his I-language fails to employ an English 'water'-address), while Felix fails to have an I-language that counts as an idiolect of Twinglish. But there may be no theoretically interesting distinction between an I-language that counts as an idiolect of English and an I-language that counts as idiolects of Twinglish. If the difference is simply a matter of which environment the implemented procedure inhabits, then for scientific purposes, there may be no difference English procedures and Twinglish procedures.

There certainly are substantive issues about concepts, and their relation to normative notions like truth, in the vicinity; see Burge (2005). And with or without such issues in mind, one can hypothesize that duplicates can have different I-languages, taking the relevant procedures (intensions) to be individuated externalistically. Chomsky (1986) did not stipulate that I-languages are individuated internalistically; his was a proposal about how to count natural languages for scientific purposes. In any case, Ludlow (forthcoming) offers a notion of ' Ψ -language' that is neutral in this respect, while preserving the idea that Ψ -languages are like I-languages in being psychologically implemented algorithms — as opposed to E-languages, which are sets (extensions) of expressions. Externalists can thus eschew appeal to E-languages, yet allow for duplicates who implement different Ψ -languages.

If the *only* constraints on semantic instructions are truth-theoretic, then *many* distinctions will be irrelevant to fulfillment. If the *i*-expression *groundhogs who like coriander* is satisfied by construction of any concept with the right truth-theoretic properties, and similarly for *woodchucks fond of cilantro*, then individuals who implement *very* different psychological processes may fulfill the same instruction. But one can individuate instructions finely — saying *both* that truth-theoretically identical phrases can be semantically distinct (internal constraints matter), *and* that twins can execute distinct instructions (external constraints matter). One cannot, however, stipulate that meanings are instructions, so individuated. For one cannot stipulate the nature of *i*-expressions.¹⁰

One can stipulate that instructions are *semantic* only if they are truth-theoretic. But then internalists can say that *i*-expressions are psemantic instructions, with silent 'p' for 'psycho' (cf. Katz & Fodor 1963), leaving it open whether or not psemantic instructions are individuated truth-theoretically. Expressions of an I-language may be meaningful, in the sense of being psemantic instructions, without being semantic instructions in the stipulated sense. If it is *also* stipulated that expressions are meaningful only if they are semantic instructions, one can regard it is an open question whether *i*-expressions are meaningful in this doubly stipulated sense. The substantive debates concern the nature of the instructions. And to repeat, thesis (1)

- (1) Meanings are instructions to build concepts.

¹⁰ One can introduce a technical notion of 'Meaning' according to which Meanings determine extensions, and describe a purpose P such that the following normative claim is plausible: other things equal, expressions of a language used for purpose P *ought* to have Meanings. But that is a very different project.

provides a relatively neutral way of framing these debates. But (1) can be combined with independently confirmable theses that are, when taken together, less neutral.

1.2. Concatenation Calls for Conjunction of Monadic Concepts

At the heart of compositional semantics lies a simple question: What is the significance of combining expressions? For invented languages, answers can be stipulated. But for I-languages (cf. Higginbotham [1985]), the questions are empirical: What is the significance of combining *i*-expressions? Or put another way, taking *i*-expressions to be *Begriffsplans*, what operations of concept combination do complex *i*-expressions invoke?

Any proposal has immediate consequences for the significance of combinable expressions, since combinables must be of the right types for purposes of combination. In part because of the justified influence of Frege (1879, 1892) and Montague (1974) — despite their own non-psychologistic (though not anti-psychological) projects — it became standard to assume that combination signifies function-application. In which case, for any pair of combined expressions, one must signify an element in the domain of the function signified by the other. This made appeals to type-shifting unavoidable, given that I-languages allow for complex predicates like *red ball*, along with simple predicates like *red* and *ball*. On this view, combination always signifies function-application, but with the caveat that combination may *also* signify a further operation of type adjustment for adjuncts. Put another way, the familiar idea was that while concatenation is a univocal instruction to apply a function to an argument, combination may have significance that goes beyond that of mere concatenation of expressions.

More recently, the trend has been towards views according to which *some but not all* cases of combination signify function-application (or equivalently, predicate-saturation), even within the tradition of transformational grammar.¹¹ Elsewhere, I have argued for a stronger hypothesis: Combination always signifies predicate-conjunction, but with the caveat that combination may *also* signify a further operation of type adjustment for arguments. In my view, while concatenation is a univocal instruction to conjoin monadic concepts, combination may have significance that goes beyond that of mere concatenation.

To clarify, it may help to consider some proposals about expression formation. We can characterize an operation COMBINE, such that if α and β are combinable expressions of an I-language, then COMBINE(α , β) is the result of combining them. This operation is somehow *asymmetric*, since phrases are headed; for example, combining a verb with a noun yields a verb phrase, not a mere verb–noun concatenation. But the operation COMBINE(α , β), so characterized, may or may not be *primitive*. It might be basic, like Chomsky's (1995b,

¹¹ See e.g. Higginbotham (1985), Larson & Segal (1995), and Heim & Kratzer (1998) — and in the tradition of categorial grammar, Steedman (2000). Of course, Davidson (1967b) had proposed logical forms involving both saturation and conjunction, at least for certain adverbially modified verb phrases.

2000b) operation $\text{MERGE}(\alpha, \beta)$. And there are several ways of being non-basic. The operation $\text{COMBINE}(\alpha, \beta)$ might turn out to be a disjunction of basic operations — e.g., $\text{ADJOIN}(\alpha, \beta)$ or $\text{SELECT}(\alpha, \beta)$ — reflecting two or more asymmetric ways of forming phrases. But another possibility is that while phrases are always formed in the same way, $\text{COMBINE}(\alpha, \beta)$ is a complex operation. In particular, phrase formation may be a process of *concatenating* expressions *and labeling* the result; see Hornstein (in press) and Hornstein & Pietroski (2008a, 2008b). So one can adopt the hypothesis below,

$$\begin{aligned}\text{COMBINE}(\alpha, \beta) &= \text{LABEL}[\text{CONCATENATE}(\alpha, \beta)] \\ \text{CONCATENATE}(\alpha, \beta) &= \alpha^{\wedge}\beta \\ \text{LABEL}[\alpha^{\wedge}\beta] &= [\alpha \beta]_{\alpha/\beta}\end{aligned}$$

with $\text{CONCATENATE}(\alpha, \beta)$ and $\text{LABEL}[\alpha^{\wedge}\beta]$ as the posited basic operations invoked by I-languages.

The subscript on $[\alpha \beta]_{\alpha/\beta}$ indicates that one of the concatenates is itself the label of the structured object, which can be identified with an ordered pair like $\{\{\alpha, \beta\}, \alpha\}$; cf. Chomsky (1995b). The source of phrasal asymmetry is then confined to the labeling operation, whose inputs are symmetric concatenations. Given any such view, one can go on to say that concatenation is an instruction to perform a certain binary operation on concepts, and that labeling is an instruction to perform some further operation on the concept associated with the non-dominant concatenate. Let me first illustrate this point from a traditional perspective, according to which concatenation signifies saturation, before turning to the idea that concatenation signifies conjunction.

One can hypothesize that the phrase $[\text{red ball}]_{\text{ball}}$ is the following instruction: Fetch a concept linked to *ball*, say $\text{BALL}(_)$, with ‘_’ as a variable ranging over individuals; fetch a concept linked to *red*, say $\text{RED}(_)$, and perform a ‘lifting’ operation to obtain the corresponding concept $\lambda X.X(X) \ \& \ \text{RED}(_)$; then saturate one of these concepts with the other, thereby obtaining a concept like $\text{BALL}(_) \ \& \ \text{RED}(_)$. One can describe the lifting operation as a reflex of the labeling: The label on $[\text{red ball}]_{\text{ball}}$ serves as an instruction to type-adjust the concept fetched via the other constituent; cf. Parsons (1970) and Kamp (1975). On this view, only *some* labels trigger this reflex. For the leading idea is that a phrase like $[\text{stab Caesar}]_{\text{stab}}$ in which a verb takes an argument, is an instruction to fetch concepts linked to the lexical constituents and then simply saturate one concept with the other (without any type-shifting). But asymmetry in the semantic effects of labeling is easily encoded.

Suppose that phrasal instructions are issued by *pairs* of expressions that have been concatenated *and labeled as a unit*, as opposed to the two concatenates taken separately. Think of $[\alpha \beta]_{\alpha}$ as a complex instruction that includes two sub-instructions that correspond to *label-relativized* concatenates: β relative to the label α , and α relative to the label α ; where one of these sub-instructions is simply α relative to itself. If phrasal instructions are issued by labeled concatenations, and CONCAT is the operation signified by concatenation, then one can say that the phrasal expression $[\alpha \beta]_{\alpha}$ is an instruction to apply CONCAT to the concepts obtained by executing two sub-instructions: β relative to α , and α relative to α ;

where relativizing an expression/instruction to itself makes no semantic difference.¹² The idea, which is independent of any particular hypothesis about CONCAT, is that the non-dominant concatenate β may need to be adjusted for purposes of combining with α . The *kind* of adjustment called for, if any, will depend on CONCAT and β . But one can hypothesize that CONCAT is the operation of saturation, and that predicate–adjunct combination (unlike predicate–argument combination) invokes type-lifting. For one can say that combining a predicate with a grammatical argument makes no difference to the instruction associated with the argument, while combining a predicate with an adjunct is an instruction to type-lift the adjunctive instruction.

If only for simplicity, let's assume that one way or another *Caesar* is marked as an argument, while *stab*, *ball*, and *red* are marked as predicates. And for ease of notation, let's underline arguments. (From a traditional perspective, one might think of underlining as an instruction to ignore phrasal relativization, thus precluding any substantive 'shifting' of the constituent semantic instruction.) Then the phrase [stab Caesar]_{stab} can be an instruction to perform saturation on concepts obtained by executing two vacuously relativized expressions/instructions: *stab* relative to *stab*, and Caesar relative to *stab*. That is, despite the phrasal label, [stab Caesar]_{stab} can be an instruction to perform saturation on concepts obtained by simply executing *stab* and *Caesar*. By contrast, [red ball]_{ball} can be an instruction to perform saturation on concepts obtained by executing two relativized expressions/instructions, one of which is *non-vacuously* relativized: *red* relative to *ball*, and *ball* relative to *ball*. That is, given the phrasal label, [red ball]_{ball} can be an instruction to perform saturation on (i) a concept obtained by executing *red* and type-lifting, and (ii) a concept obtained by simply executing *ball*.

Alternatively, one can invert the traditional perspective and adopt the following hypothesis: CONCAT is an operation of conjunction — more specifically, conjunction of monadic concepts; and predicate–argument combination, unlike predicate–adjunct combination, invokes a kind of type-adjustment. I discuss my specific neo-Davidsonian proposal, in terms of thematic roles, in section 1.3 below. But if combining a predicate with an adjunct makes no difference to the instruction associated with the adjunct, then [red ball]_{ball} can be an instruction to conjoin two monadic concepts, obtained by executing two vacuously relativized instructions: *red* relative to *ball*, and *ball* relative to *ball*. That is, despite the phrasal label, [red ball]_{ball} can be an instruction to conjoin concepts obtained by simply executing *red* and *ball*. And we can say that combining a predicate α with a grammatical argument β is an instruction to type-adjust the instruction issue by β itself. For given the phrasal label, [stab Caesar]_{stab} can be an instruction to conjoin a pair of monadic concepts obtained by executing two relativized expressions/instructions, one of which is *non-vacuously* relativized: *stab* relative to *stab*, and *Caesar* relative to *stab*. That is, [stab Caesar]_{stab} can be an instruction to conjoin concepts obtained by (i) simply executing *stab*, and (ii) executing *Caesar* and type-adjusting.

¹² This is a restricted version of the relativization that Higginbotham (1986) employs. For these purposes, I treat *ball*_N and *stab*_V as primitives. But see Hornstein & Pietroski (forthcoming) for more detailed discussion, and some implications of treating labels as formatives that can be combined with unlabeled lexical roots.

If it helps, think of underlining on this nontraditional view as an instruction to relativize *non-vacuously*: The ‘default’ effect of label-relativization is null, with a restricted kind of (thematic) relativization as the marked case. I cannot here defend this proposal in any detail. But in I-languages, adjunction seems to be open-ended (especially given relative clauses), while predicates seem to combine with at most three grammatical arguments.¹³ *Prima facie*, this casts doubt on the semantic tradition of treating adjuncts as the marked cases that call for special treatment. And if adjunction invokes monadic concept conjunction, as a (biologically implemented) recursive combination operation, one might wonder if I-languages *also* invoke a recursive operation of saturation in which a concept of adicity n combines with a semantic argument to yield a concept of adicity $n-1$. But the more important point here is that we can (at least conceptually) decompose the semantic effect of combining expressions into the semantic effects of concatenating and labeling. And since labeling need not be semantically vacuous, one can distinguish the semantic effect of combining expressions from the semantic effect of concatenating expressions — even if in some cases, this distinction is not semantically significant.

In short, each expression of the form $[\alpha \beta]_\alpha$ can be a ‘macro’ instruction: execute the sub-instructions α and β , obtaining concepts $C1$ and $C2$, respectively; then form the concept $\text{CONCAT}(C1, \text{LAB}[C2, \alpha])$, where $\text{LAB}[C2, \alpha]$ is the concept formed by subjecting $C2$ to the (perhaps vacuous) operation induced by the label α . Again, this general idea is compatible with various proposals about which operations CONCAT and LAB are. But suppose CONCAT is an operation of monadic concept conjunction, signified with ‘•’. Then $[\alpha \beta]_\alpha$ is an instruction to build a monadic concept of the following form: $C1(_) \bullet \text{LAB}[C2(_), \alpha]$; where $\text{LAB}[C2(_), \alpha]$ may be a complex thematic concept of the form $\exists x[C(x) \bullet \Theta(_, x)]$. Potential values of the variable ‘_’ include events, which can have individual participants, and individuals (which can be event participants). So concatenation can be an instruction to conjoin monadic concepts — and in this sense, conjunction can be the basic I-language mode of semantic combination — without each expression of the form $[\alpha \beta]_\alpha$ being an instruction to conjoin concepts formed by executing the sub-instructions α and β .

Strictly speaking, ‘•’ is a little more permissive than an operator that can be flanked *only* by monadic concepts. The idea is that (the conceptual operation indicated with) ‘•’ can be used to combine a monadic concept $C(x)$ with a formally dyadic concept $\Theta(_, x)$ — with ‘ x ’ ranging over potential participants of values of ‘_’ — so long as one variable in the resulting concept, $C(x) \bullet \Theta(_, x)$, is immediately closed to create a formally monadic concept. But one can hypothesize that this *slightly* relaxed operation of monadic concept conjunction is the one invoked by I-languages.¹⁴

¹³ In this context, see Pietroski (2005a) for an argument that *that*-clauses are semantically like adjuncts, and a review of the trend towards (i) analyzing ditransitive constructions in terms of a covert preposition (suggesting a maximum of *two* arguments) and (ii) analyzing many transitive constructions in terms of a covert light verb.

¹⁴ Here, I am also ignoring the difference between verbish instructions (which tend to be tensable) and nounish instructions (which tend to be indexable); see Hornstein & Pietroski (forthcoming) for discussion drawing on Baker (2005).

1.3. Grammatical Relations Invoke Thematic Relations and \exists -Closure

For present purposes, I take it as given that in I-languages, predicate-argument relations are at least often associated with thematic relations. Correspondingly, I assume, it is not *ad hoc* to say that an argument like *Caesar* in *stab Caesar* is somehow associated with a thematic concept like $\text{PATIENT}(_ \text{CAESAR})$; where this complex monadic concept applies to ‘events’ that have Caesar as their patient. If we can seriously entertain the more traditional idea that $[\text{red ball}]_{\text{ball}}$ is an instruction to form a concept like $\lambda X.X(_) \ \& \ \text{RED}(_)$, which gets saturated with a concept like $\text{BALL}(_)$, then we can seriously entertain a structurally similar idea: $[\text{stab Caesar}]_{\text{stab}}$ is an instruction to form a concept like $\text{PATIENT}(_ \text{CAESAR})$ and conjoin it with another monadic concept of events like $\text{STAB}(_)$; likewise, $[\text{Brutus} \text{stab Caesar}]_{\text{stab}}_{\text{stab}}$ is an instruction to add a conjunct like $\text{AGENT}(_ \text{BRUTUS})$.

I also assume that I-languages invoke a cognitive operation that is like existential closure in two respects. First, when the operation is applied to a monadic concept $C(_)$, it yields a complete thought of the form $\exists _ [C(_)]$; where a thought of this form is correct iff $C(_)$ applies to one or more things. More briefly, the operation converts $C(_)$ into a thought that is true iff $C(_)$ is not empty. Second, the operation can convert a formally dyadic concept of the form ‘ $C(x) \bullet \Theta(_, x)$ ’ — with values ‘ x ’ being potential participants of values of ‘ $_$ ’ — into a monadic concept of form ‘ $\exists x [C(x) \bullet \Theta(_, x)]$ ’; where a concept of this form applies to one or more potential values of ‘ $_$ ’ iff they are related, in the right thematic way, to one or more things that fall under the relevant monadic concept. But I do not assume that the closure operation invoked by I-languages, signified here with ‘ \exists ’, has the full power of existential closure to convert any concept of adicity n into a concept of adicity $n-1$ by binding any conceptual variable. For example, I do not assume that this operation can convert a tetradic concept of the form ‘ $R(x, y) \ \& \ S(w, x, z) \ \supset \ T(y, z, w)$ ’ into a triadic concept of the form ‘ $\exists x [R(x, y) \ \& \ S(w, x, z) \ \supset \ T(y, z, w)]$ ’. Correlatively, the cognitive effect of the posited closure operation need not be characterized (*à la* Tarski) in terms of the satisfaction conditions imposed by open sentences with arbitrarily many variables on sequences of arbitrary many potential values of variables.

In Pietroski (2005b, 2006, 2008a, 2008b), I show how these relatively modest resources can provide a compositional semantics that accommodates a wide range of constructions — including causative and ditransitive constructions, plural noun phrases, prepositions, negation, and quantificational determiners like *every* that take tensed clauses as their external arguments. This is not the place for an explicit fragment of a semantic theory. But to illustrate, suppose that internal and external arguments of predicates are relativized as such.

For any grammatical argument \dots , let *int- \dots* be the instruction issued by \dots when it appears as the internal argument of a predicate, and let *ext- \dots* be the instruction issued by \dots when it appears as the external argument of a predicate. (If each predicate takes at most two grammatical arguments, ‘*ext*’ can be replaced with ‘*~int*’, treating each external argument as the non-internal argument of its predicate.) And suppose that each argument is, by itself, an instruction to form a monadic concept. This concept may be complex and context sensitive, reflecting a

complex expression that includes a lexical noun along with a covert determiner and/or index. But for now, we can idealize, taking the relevant concepts to be atomic: CAESARIZER($_$) and BRUTUSIZER($_$); cf. Quine (1963), though see section 1.4 below. If grammatical arguments are instructions to build monadic concepts, it is easy to provide composition principles according to which the relativized expressions '*int-Caesar*' and '*ext-Brutus*' are instructions to build concepts of things with internal/external participants: $\exists x$ [CAESARIZER(x) • INTERNAL($_$, x)], $\exists x$ [BRUTUSIZER(x) • EXTERNAL($_$, x)]; where in each case, the variable introduced by the proper noun is \exists -closed.

Thus, [Brutus [stab Caesar]_{stab}]_{stab} can be an instruction to construct a concept like the following: $\exists x$ [BRUTUSIZER(x) • EXTERNAL($_$, x)] • [STAB($_$) • $\exists x$ [CAESARIZER(x) • INTERNAL($_$, x)]]]. And if *stab* is understood as an instruction to fetch a concept of things (actions) whose internal/external arguments are the patients/agents of those things, the formalistic notions INTERNAL($_$, x) and EXTERNAL($_$, x) can be replaced with more specific thematic contents: $\exists x$ [BRUTUSIZER(x) • AGENT($_$, x)] • [STAB($_$) • $\exists x$ [CAESARIZER(x) • PATIENT($_$, x)]]].¹⁵

Adding tense and a final existential closure can yield a complete thought, corresponding to *Brutus stabbed Caesar*: $\exists _$ [PAST($_$) • [$\exists x$ [BRUTUSIZER(x) • AGENT($_$, x)] • [STAB($_$) • $\exists x$ [CAESARIZER(x) • PATIENT($_$, x)]]]]]. Alternatively, the untensed clause can itself be the internal argument of a larger verb phrase as in *see Brutus stab Caesar*, as shown:

$$\text{SEE}(_) \bullet \exists x[\text{INTERNAL}(_, x) \bullet [\exists x'[\text{BRUTUSIZER}(x') \ \& \ \text{AGENT}(x, x')] \bullet [\text{STAB}(x) \bullet \exists x[\text{CAESARIZER}(x') \bullet \text{PATIENT}(x, x')]]]]]$$

And this construction can be treated on a par with *see a tree*:

$$\text{SEE}(_) \bullet \exists x[\text{INTERNAL}(_, x) \bullet \text{TREE}(x)]$$

Of course, one wants to see the details for other construction types. But for present purposes, I take it as premise that the requisite homework can be done, along the lines suggested in Pietroski (2005b, 2006, 2008a, 2008b): The Conjunctivist idea outlined in section 1.2 can be turned into detailed theory, by hypothesizing that (i) certain grammatical relations serve as instructions to introduce certain dyadic/thematic relations, and (ii) each cycle or 'phase' directs construction of a conjunctive concept whose main variable, introduced by the syntactic head, can be \exists -closed; where this may leave open exactly one thematically introduced variable, thereby allowing for subsequent conjunction with other monadic concepts. A more radical suggestion, not explored here, is that the requisite \exists -closures reflect interface conditions between HFL — which always generates instructions to construct monadic concepts, as opposed to complete sentences of type <t> — and cognitive systems whose representations have *correctness* conditions along with at least some *dyadic/thematic* constituents.

¹⁵ One can add that the internal/external participants of stabs are their agents/patients: $\forall e[\text{STAB}(e) \supset \forall x[\text{AGENT}(e, x) \equiv \text{EXTERNAL}(e, x)] \ \& \ \forall x[\text{PATIENT}(e, x) \equiv \text{INTERNAL}(e, x)]]$.

1.4. *Lexicalization is a Partly Creative Process of Abstraction*

I have been assuming that lexical items can be instructions to fetch monadic concepts, even with regard to words like *Brutus* and *stab*, which were presumably introduced in the course of lexicalizing non-monadic concepts like BRUTUS and STAB(*X*, *Y*) — or STAB(*X*, *Y*, *Z*), with ‘*z*’ as a variable for instruments, or STAB(*X*, *Y*, *E*), or STAB(*X*, *Y*, *Z*, *E*). I readily grant that humans and other animals have singular and polyadic concepts, and that such concepts are often lexicalized. But the concept lexicalized with an *i*-expression need not be the concept subsequently fetched with that expression. For lexicalization can be a process in which non-monadic concepts are paired with monadic analogs, even if the monadic analogs have to be abstracted from the concepts lexicalized.

Imagining mechanisms for such abstraction is not difficult. For illustration, suppose the concept lexicalized with *stab* is dyadic, with no event variable. Given STAB(*X*, *Y*), and enough logical apparatus to define new concepts, lexicalizers might be able to introduce a triadic concept STAB(*_*, *X*, *Y*): $\forall X \forall Y \{ \text{STAB}(X, Y) \equiv \exists _ [\text{STAB}(_ , X, Y)] \}$; cf. Davidson (1967b). Then STAB(*_*) might be introduced via thematic notions: $\forall X \forall Y \forall _ \{ \text{STAB}(_ , X, Y) \equiv \text{AGENT}(_ , X) \ \& \ \text{STAB}(_) \ \& \ \text{PATIENT}(_ , X) \}$; cf. Castañeda (1967) and Davidson (1985). Then *stab* can be an instruction to fetch STAB(*_*), with thematic notions introduced via grammatical arguments and/or prepositional phrases, which may be lexically mandatory or optional; see Pietroski (2008a, 2008b) for further details.

This kind of abstraction may require (non-recursive) cognitive resources that are not needed for Conjunctivist *composition* of *i*-concepts. But lexicalization is one thing, composition another. Or perhaps lexicalizers can use STAB(*X*, *Y*) to introduce STAB(*_*) more directly and in more restricted terms:

$$\exists X \exists Y [\text{STAB}(X, Y)] \equiv \exists _ \{ \exists X [\text{AGENT}(_ , X)] \bullet [\text{STAB}(_) \bullet \exists X [\text{PATIENT}(_ , X)]] \}$$

But whatever the details, the idea is that children can abstract a monadic concept of stabs from the presumably polyadic concept lexicalized with *stab*. And given an independent grip on the relevant thematic notions, one can allow for patientless (and/or agentless) stabs, at least as conceptual possibilities; cf. Parsons (1990). For even if each actual stab is a stabbing of something by someone, there may be no contradiction in the thought that some stab that lacks a patient also lacks an agent.

In general, a lexicalizable *n*-adic concept can be used to define a monadic analog, given *n* thematic relations. And there are various ways of reanalyzing singular concepts as monadic, even without appealing to predicates like *Brutusizer*, which apply to at most one thing. Suppose the name *Brutus* is introduced in the course of lexicalizing an atomic singular concept. Even if the pre-lexical concept is a simple ‘mental tag’ of type <e>, the name can be a complex expression used to fetch and combine two concepts — one demonstrative, and one metalinguistic (cf. Burge 1973 or Katz 1994). In many languages, proper nouns can and often must appear with an overt determiner or demonstrative, as in *the/that/our Brutus*, suggesting that the proper noun is (used to fetch a concept that is) of type <e, t>; see Longobardi (1994), Giannakidou & Stavrou (1999). If

names in English are structurally similar, with a covert functional element, the lexical proper noun *Brutus* can be analyzed as an instruction to fetch a concept of things called (with the sound of) *Brutus*. Independent evidence in favor of some such analysis is independent evidence that singular concepts are indeed paired with monadic analogs in the course of lexicalization; see Pietroski (2007), drawing on many others, for further discussion. But let me summarize, to this point, by returning to a more theoretically neutral idea.

Whatever one says about the significance of the $\text{COMBINE}(\alpha, \beta)$, expressions formed via this grammatical operation can be viewed as instructions to construct concepts of certain types. Unless one thinks that grammatical combination is effectively unrestricted, there will be many possible conceptual types such that *i*-expressions are never instructions to construct concepts of those types. On the assumption that $\text{COMBINE}(\alpha, \beta)$ signifies at most a few operations, and that lexical types are subject to nontrivial constraints of some kind, there will be some range of types such that each *i*-expression is an instruction to construct a concept of one of those types. All *i*-concepts will belong to this class of concepts.

Conjunctivism is a very restrictive thesis according to which all *i*-concepts are monadic. Theorists are free to adopt more permissive semantic theories, according to which *i*-concepts exhibit many types: $\langle e \rangle$, $\langle t \rangle$, $\langle e, t \rangle$, $\langle e, \langle e, t \rangle \rangle$; $\langle t, t \rangle$, $\langle t, \langle t, t \rangle \rangle$; $\langle \langle e, t \rangle, t \rangle$, $\langle \langle e, t \rangle, \langle \langle e, t \rangle, t \rangle \rangle$; $\langle \langle e, t \rangle, \langle e, t \rangle \rangle$, and perhaps others. But I know of no good reason for positing *i*-concepts of type $\langle \langle t, e \rangle, e \rangle$, and likewise for endlessly many other possible concepts definable in terms of $\langle e \rangle$ and $\langle t \rangle$. Ideally, one wants a characterization of the possible types for *i*-concepts that is empirically adequate without overgenerating. In this respect, Conjunctivists try to err on the side of positing too little: The idea is to start by supposing that all *i*-concepts are type $\langle e, t \rangle$, and then find out which facts can/cannot be accommodated in this fashion.¹⁶ But whatever one says about the space of *i*-concepts, the concepts fetched via lexical items have to be in this space, since $\text{COMBINE}(\alpha, \beta)$ only operates on combinable expressions that are used to fetch or construct *i*-concepts.

By contrast, there is no requirement that all concepts be *i*-concepts. Humans and other animals may have concepts that cannot be combined via the operation(s) signified via $\text{COMBINE}(\alpha, \beta)$. For example, concepts are often saturated. But it does not follow that some *i*-concepts saturate others. And on any view, some of the concepts that a child lexicalizes may fail to be of the right form for purposes of I-language combination. In which case, some and perhaps many pre-lexical concepts will have to be paired with *i*-concept analogs that can be fetched via the resulting words.

¹⁶ If all *i*-concepts are number neutral (see Schein 1993, 2002, 2006 and Pietroski 2005b, 2008a, 2008b), perhaps we should say that *i*-concepts are uniformly of type $\langle _ , t \rangle$; where judgments are of type $\langle t \rangle$, and a concept is of type $\langle _ , t \rangle$ if it applies to *one-or-more* things, allowing for concepts that are neither essentially singular (in that they apply to two or more things only distributively) nor essentially plural (in that they only apply, and only apply non-distributively, to two or more things). In which case, it may be that *no* pre-lexical concepts are *i*-concepts. For it may be that all pre-lexical concepts are either essentially singular or essentially plural. But let's set this complication aside.

Imposing an I-language on pre-lexical thought may thus require some conceptual ‘reformatting’. Frege (1879, 1884, 1892) envisioned a process of imposing a *Begriffsschrift* on pre-scientific thought, recognizing that this would require considerable reformatting of our natural ideas. By inventing modern logic, and offering some ‘fruitful definitions’ of key arithmetic notions, Frege provided a model of how such reformatting might proceed — at least in principle — for the special case of imposing a now familiar hierarchy of types, which correspond to concepts that can be combined via function-application. (See Horty (2007) for extended discussion.) Frege was not offering hypotheses about HFL; he was concerned with an idealized language of scientific thought, designed to reflect the mind-independent world, not the I-languages that children so readily acquire. But one can use the basic types $\langle e \rangle$ and $\langle t \rangle$ to define a space of ‘F-concepts’, and then characterize sub-regions of this space in terms of constraints corresponding to possible operations for combining F-concepts.

Frege’s operation of saturation imposes no constraint; for any F-concept, there is another such that they can be combined via saturation. But other operations, like monadic concept conjunction, only permit certain combinations of F-concepts. More generally, given one or more composition operations, let’s say that an F-concept of type $\langle \alpha \rangle$ is *licensed* iff for some type $\langle \beta \rangle$, an F-concept of type $\langle \alpha \rangle$ can be combined with an F-concept of type $\langle \beta \rangle$ given the operation(s). Correlatively, given one or more composition operations, we can at least imagine minds that go through a critical period in which many mental representations that are *not* licensed by those operations get used to abstract formally distinct but analytically related concepts that *are* licensed by those operations. One can hypothesize that human children go through some such critical period, because they have a language faculty, and that we humans thereby acquire a Fregean ‘second nature’ in the following sense: We acquire certain *i*-concepts, and thereby acquire some F-concepts we did not already have; where these formally new concepts can be combined via operations invoked by I-languages. Conjunctivists adopt the view that all *i*-concepts are monadic. In which case, the combination operations can be rather restrictive, but lexicalization needs to be a little creative, in ways that dovetail with the special role that certain dyadic/thematic concepts play in allowing I-languages to interface with pre-lexical thought.

1.5. *Meanings are Internalistic Properties of Expressions*

As already noted, one can say that meanings are instructions to create *i*-concepts, while maintaining that the semantic properties of expressions are individuated externalistically. Likewise, Conjunctivism is logically compatible with truth conditional semantics. One can say that theories of meaning for I-languages will take the form of Tarski-style theories of truth, with concatenation signifying conjunction. But truth conditional semantics faces considerable difficulties, even if one ignores biologically imposed constraints on composition. And in my view, it is especially implausible that *i*-expressions have Tarskian satisfaction conditions that compose in the fashion sketched above. In short, (1)–(4) may together make (5) plausible.

- (1) Meanings are instructions to build concepts.
- (2) Concatenation signifies predicate-conjunction.
- (3) Grammatical relations invoke thematic relations and \exists -closure.
- (4) Lexicalization is a partly creative process of abstraction.
- (5) Meanings are internalistic properties of expressions.

For if Conjunctivism is even roughly correct, it seems that the nature of semantic composition is determined by relatively simple computations, which permit construction of representations with formal properties that make them suited for ‘interfacing’ between HFL (the human faculty of language) and other aspects of human cognition. It would be amazing — in ways suggesting cosmic benevolence — if these simple computations also let us generate representations that reflect the language-independent world well enough to have compositionally determined satisfaction conditions. Frege and Tarski showed us, among other things, just how hard it is to *design* a language that has a truth-theoretic semantics, even when one can *stipulate* the operative composition principles. Especially given the much discussed paradoxes — involving, for example, self-reference and vagueness — a plausible externalism about *truth* may impose limits on the kinds of languages for which a truth-theoretic semantics can be given. One cannot just declare that an expression like *set of bald linguists who often use heterological words* has a compositionally determined satisfaction condition. But even setting such considerations aside, constraints on I-languages seem to be at best orthogonal to the requirements of compositional truth theories, and often at odds with the idea that there are Conjunctivist truth theories for I-languages.

To take a much discussed example, consider *The sky is blue*, which seems to be on a grammatical par with *The pie is round*. *Prima facie*, a truth conditional semantics will need to treat *sky* as a predicate satisfied by skies, one of which can be (in a context *c*) the relevant thing *x* such that the truth or falsity of the sentence (relative to *c*) depends on whether or not *x* is blue (in *c*). For now, set aside worries about what contexts need to be, and what it is to be blue (in a context). What is *x* supposed to be? Tarski could stipulate that the satisfiers of a given predicate were, for example, natural numbers; where these abstracta were antecedently well defined. Correspondingly, the right hand sides of metalinguistic claims like ‘*Pb* is true iff 2 is prime’ are couched in a theoretical idiom that is antecedently understood in terms of an explicit model of certain aspects of reality. Given certain notational conventions, the number two can be identified with a certain set; and we know, in a scientific way, what it is for a number to be prime. But what are skies, and what is it to identify one of them as the one said to be blue? Absent answers to these questions, it is hard to even begin evaluating the hypothesis that an *i*-expression like *The sky is blue* has a compositionally determined truth condition.

I don’t deny that we humans have one or more *sky-concepts*, along with various *color-concepts*. On the contrary, I think speakers can and do access such concepts in response to semantic instructions like *The sky is blue*. But even if one speculates that these concepts have satisfaction conditions, one need not burden semantic theorizing with the further speculation that *i*-expressions inherit these

satisfaction conditions. Recalling section 1.1, in any given I-language, the word *sky* may be an instruction to fetch a concept from a certain lexical address — as opposed to an instruction to fetch a concept with a certain satisfaction condition. If a child acquires an I-language that counts as an idiolect of English, then presumably, the child uses lexical *i*-expressions to fetch concepts that have been paired with phonological forms in ways that respect substantive constraints. But it doesn't follow that the child's *i*-expressions are (or that the child somehow takes her *i*-expressions to be) instructions to fetch concepts that meet these constraints. The instructions may be address-focused, even if more demanding conditions are met, *de facto*.

To be sure, theorists can inscribe 'axioms' and 'theorems' like the following: *x* satisfies *Brutus* iff *x* is a sky; *x* satisfies *blue* iff *x* is blue; *x* satisfies *blue sky* iff *x* is a sky and *x* is blue; etc. But absent a specific proposal about the potential values of the variable, this seems like an overly technical (and therefore misleading) way of saying that *blue* and *Brutus* are devices for fetching monadic concepts, and that *blue sky* is an instruction to conjoin concepts fetched via *blue* and *Brutus*. Of course, any one example is just that. But especially if concatenation uniformly signifies a simple operation like conjunction, truth conditional semanticists will continually face uncomfortable questions about values of the relevant variables.

One can be lulled into thinking that at least event variables are friendly to truth conditional semantics, since there are indeed events. But as Davidson (1985) himself noted, difficult questions of event individuation arise as soon as one takes such variables seriously in the context of a truth theory. (Are events individuated in terms of their spatiotemporal location, their causes/effects, their participants, or still other factors? Under what conditions are counterfactual claims about an event true?) And if we take events to be the values of event variables in theories of meaning for I-languages, without *stipulating* that these variable-values are also language-independent spatiotemporal particulars, then appealing to thematic relations quickly leads to the conclusion that events are distinguished in ways that reflect grammatical distinctions; see e.g. Schein (1993, 2002) and Tenny (1995). For example, no event of you facing me can be identical with an event of me facing you: Otherwise, a facing whose agent is *you* would be a facing whose agent is *me*; in which case, you'd be me. Or suppose there was an event of Jim Higginbotham drinking a (pint of) beer in a minute. It seems that qua value of an event variable, any such event must be distinct from any simultaneous event of Jim drinking beer for a minute, at least if *in a minute* and *for a minute* are truth-theoretic conjuncts of the relevant event descriptions. For if there is just *one* event of drinking that satisfies *both* conjuncts, then *prima facie*, *Jim drank beer in a minute* should be just as true as *Jim drank beer for a minute*.

To take another kind of example, *France is a hexagonal republic* is somehow weird in way that *France is hexagonal* and *France is a republic* is not. *Prima facie*, this asymmetry is at odds with the idea that a sentence of the form 'France is Φ ' is true iff the satisfier of 'France' satisfies ' Φ ', and that *hexagonal republic* is satisfied by *x* iff *x* satisfies both *hexagonal* and *republic*. I don't expect these simple illustrations to convince. Arguing against truth conditional semantics requires

discussion of many specific constructions, and the many potential replies.¹⁷ But it isn't hard to get a feel for the general Cartesian worry, to which any nativist should be sensitive — viz., that *i*-expressions are concept–construction-instructions whose basic architecture is determined by endogenous constraints, not the language-independent world; cf. Chomsky (2000a).

Humans are lucky to have HFL, a faculty that lets us acquire I-languages in conditions of limited experience, and then use *i*-expressions that are compositional in *some* sense that is compatible with both the nature of HFL and the concepts we lexicalize. If *i*-expressions are also compositional in a truth-theoretic sense, then the demands of HFL and infant psychology somehow conspire to yield I-languages that are relevantly like the languages that Frege and Tarski *designed* to be truth-theoretically compositional. That sounds, to my ear, like wishful thinking. I'll return to this point after a brief detour.

2. Speculation: Lexicalization as a Human Cognitive Tool

If my proposed account of semantic composition is roughly on the right track, it invites a non-standard conjecture about what makes HFL distinctively human.

Consider two theses, at least one of which is presumably true.

- (L) Humans have a special capacity to lexicalize mental representations.
- (C) Humans have a special capacity to combine mental representations.

Perhaps both are true. But other things equal, one doesn't want to posit two distinctively human capacities that somehow manage to interact in the right ways. Thesis (C) has a distinguished heritage; see Hauser, Chomsky & Fitch (2002) for a recent discussion. It can seem obvious that recursion is the key to human language. But as Chomsky (1957) famously discussed, there's recursion, and then there's recursion. A mere concatenator of atomic expressions can generate arbitrarily many complex expressions of the form $\alpha^{\wedge}\beta$; and non-human animals can surely concatenate at least some representations.

I suspect that other animals also have the basic capacities required to treat concatenations as Conjunctivist instructions: Do this, do that, and connect the results with something like an AND-gate. It seems obvious that nonhuman animals have many concepts — or if you prefer, pre-lexical mental representations — that are at least candidates for lexicalization. In terms of nonlinguistic cognitive capacities, our evolutionary cousins are quite impressive, especially compared with human infants. And it seems quite plausible that other social

¹⁷ See e.g. Stanley (2000, 2002) and Schein (forthcoming); cf. Pietroski (2005b, 2006b). This is *not* to deny the enormous value of extant work done within the framework of truth-conditional semantics. Nor is to suggest that a Conjunctivist theory will itself be adequate to capture the meanings of *i*-expressions. On the contrary, a Conjunctivist theory may capture only a semantic 'primal sketch' that is supplemented by other processes that are not compositional in the same way. As with vision, it is hard to say where one faculty/module ends, and interfaces to 'general cognition' begin. But in my view, talk of truth conditions reflects a mix of linguistic factors and other cognitive (or metaphysical) factors that are orthogonal to the study of HFL.

primates can (when suitably stimulated) come to have some thematic concepts corresponding to participation relations exhibited by events and agents/patients/places/times/etc.¹⁸ But humans have still more, on at least two fronts, even if we set aside the possibility that the concepts humans lexicalize are somehow especially lexicalizable.

First, as already noted, grammatical combination is asymmetric in a way that mere concatenation is not. Following Hornstein (in press), I have suggested that labeling is the source of this asymmetry; see also Boeckx (2008). Given labeled phrases, grammatical relations like 'being the internal/external argument of' can serve as devices for invoking thematic relations, and thereby distinguishing the significance of mere *concatenation* from the significance of I-language *combination*. But one might wonder: If labeling is all it takes to unleash the power of HFL from antecedently available animal capacities, why don't many species have analogs of HFL?

This leads to the second point, which is that the cognitive value of I-languages is limited if there are only a few lexical *i*-concepts. If only a few atomic concepts can be fetched, for purposes of creating complex concepts, recursively generating concept construction instructions won't do much good. So especially if concatenation signifies an elementary cognitive operation, which by itself only allows for simple cases of concept composition, it may be that lexicalization is the key new linguistic trick. Though correlatively, the utility of abstraction mechanisms that pair lexicalizable concepts with *i*-concepts may depend on a capacity to efficiently combine *i*-concepts via independently available operations. Perhaps some ancestral primates fortunately *connected* a capacity to label concatenations (or otherwise introduce some such source of grammatical asymmetry) with a capacity for formal abstraction that allowed for lexicalization (i.e. the creation of *i*-concepts, not mere pairs of signals with pre-lexical concepts).

Since this already seems like a lot to posit, in terms of distinguishing humans from other animals, I am wary of also positing a capacity to employ an operation like function-application (concept-saturation) recursively. By contrast, conjunction appears to be a simple and ubiquitous computation. Moreover, whatever we say about the significance of combination, we are faced with the empirical fact that human children naturally lexicalize with a vengeance; whereas other animals can 'only' learn to pair signals with concepts to a certain extent, given lots of explicit training. So returning to theses (L) and (C), we have independent evidence in favor of (L), making it undesirable to *also* posit (C). And other things equal, I would rather not say that other primates have a (perhaps unused) capacity to construe concatenations as recursive instructions to fetch and saturate polyadic concepts of arbitrary adicity.

Given that Conjunctivism seems to have a chance of yielding descriptive adequacy, I conclude that we should explore the possibility that lexicalization is a large part of the uniquely human aspect of HFL — and that *i*-concepts are monadic *because* conjoining monadic concepts is a common cognitive capacity that can be used in novel and fruitful ways, once existing concepts are paired with (monadic) *i*-concepts. From this perspective, I-languages impose a common

¹⁸ See Hurford (2007) for relevant and congenial discussion of various evolutionary issues.

and simple format on animal concepts, with the result that human *i*-concepts can be combined via simple operations. This may have real value in terms of computational efficiency.

It may also be occasionally useful for purposes of representing the world. But on the whole, one would predict what one actually sees when looking at I-languages in any detail: Many features (as in expressions like *bald linguists who like blue skies*) that seem like quirks, and often serious design flaws, if one assumes that expressions have compositionally determined Tarskian satisfaction conditions; and many features, like invocation of thematic roles, that make sense if one assumes that expressions are compositional in a more simple-minded, less world-directed way.

Again, I don't expect this compressed argument to convince. But the idea is that if Conjunctivism turns out to be the best account of semantic composition — and of how it came to pass that the biological world includes semantically compositional I-languages — then that is itself an argument against truth conditional semantics, given the plausible assumption that there are no *Conjunctivist* theories of *truth* for I-languages.

More generally, we should evaluate semantic externalism and other high-level theses about meaning in light of our best proposals concerning (i) the significance of the operation COMBINE, and (ii) how humans came to have a faculty in which this operation has that significance. Even if my proposal is entirely wrong in detail, I hope to have illustrated how a cluster of theses like (1)–(5) can hang together.

- (1) Meanings are instructions to build concepts.
- (2) Concatenation signifies predicate-conjunction.
- (3) Grammatical relations invoke thematic relations and \exists -closure.
- (4) Lexicalization is a partly creative process of abstraction.
- (5) Meanings are internalistic properties of expressions.

For we can and should ask how alternative clusters compare, if only to loosen the current grip of various semantic dogmas on our collective theoretical imagination.

References

- Bloom, Paul. 2000. *How Children Learn the Meanings of Words* (Learning, Development, and Conceptual Change). Cambridge, MA: MIT Press.
- Boeckx, Cedric. 2008. *Bare Syntax*. Oxford: Oxford University Press.
- Burge, Tyler. 1973. Reference and proper names. *Journal of Philosophy* 70, 425–39.
- Burge, Tyler. 1979. Individualism and the mental. *Midwest Studies* 4, 73–121.
- Burge, Tyler. 2005. *Truth, Thought, Reason: Essays on Frege*. Oxford: Oxford University Press.

- Chomsky, Noam. 1957. *Syntactic Structures*. The Hague: Mouton.
- Chomsky, Noam. 1977. *Essays on Form and Interpretation*. New York: North Holland.
- Chomsky, Noam. 1986. *Knowledge of Language*. New York: Praeger.
- Chomsky, Noam. 1995a. Language and nature. *Mind* 104, 1–61.
- Chomsky, Noam. 1995b. *The Minimalist Program*. Cambridge, MA: MIT Press.
- Chomsky, Noam. 2000a. *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, Noam 2000b. Minimalist inquiries: The framework. In Roger Martin, David Michaels & Juan Uriagereka (eds.), *Step by Step: Essays on Minimalist Syntax in Honor of Howard Lasnik*, 89–155. Cambridge, MA: MIT Press.
- Church, Alonzo 1941. *The Calculi of Lambda-Conversion*. Princeton, NJ: Princeton University Press.
- Castañeda, Hector-Neri, 1967. Comments. In Rescher (1967), 104–112.
- Davidson, Donald. 1967a. Truth and meaning. *Synthese* 17, 304–23.
- Davidson, Donald. 1967b. The logical form of action sentences. In Rescher (1967), 81–95.
- Davidson, Donald. 1985. Adverbs of action. In Bruce Vermazen & Merrill B. Hintikka (eds.), *Essays on Davidson: Actions and Events*, 230–241. Oxford: Clarendon Press.
- Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Explorations in Cognitive Science). Cambridge, MA: MIT Press.
- Fodor, Jerry A. 2003. *Hume Variations* (Lines of Thought). Cambridge: Blackwell.
- Fodor, Jerry A. 2008. *LOT II*. Oxford: Oxford University Press.
- Frege, Gottlob. 1879. *Begriffsschrift*. Halle: Louis Nebert. [English translation in van Heijenoort, Jean (ed.). 1967. *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931* (Source Books in the History of the Sciences). Cambridge, MA: Harvard University Press.]
- Frege, Gottlob. 1884. *Die Grundlagen der Arithmetik*. Breslau: Wilhelm Koebner. [English translation in Frege, Gottlob. 1974. *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number* [transl. by J.L. Austin]. Oxford: Basil Blackwell.]
- Frege, Gottlob. 1892. Function and concept. In Peter Geach & Max Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege*, 21–41. Oxford: Blackwell, 1980.
- Giannakidou, Anastasia & Melita Stavrou. 1999. Nominalization and ellipsis in the Greek DP. *The Linguistic Review* 16, 295–331.
- Harman, Gilbert. 1970. Deep structure as logical form. *Synthese* 21, 275–97.
- Hauser, Marc, Noam Chomsky & W. Tecumseh Fitch. 2002. The faculty of language: What is it, who has it, and how did it evolve? *Science* 298, 1569–1579.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in Generative Grammar* (Blackwell Textbooks in Linguistics). Oxford: Blackwell.
- Higginbotham, James. 1985. On semantics. *Linguistic Inquiry* 16, 547–93.
- Higginbotham, James. 1986. Linguistic theory and Davidson's program. In Ernest Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, 29–48. Oxford: Basil Blackwell.

- Hinzen, Wolfram 2007. *An Essay on Names and Truth*. Oxford: Oxford University Press.
- Hornstein, Norbert. In press. *A Theory of Syntax: Minimal Operations and Universal Grammar*. Cambridge: Cambridge University Press.
- Hornstein, Norbert & Paul Pietroski. Forthcoming. Basic operations. *Catalan Journal of Linguistics*.
- Horty, John. 2007. *Frege on Definitions: A Case Study of Semantic Content*. Oxford: Oxford University Press.
- Hurford, James R. 2007. *The Origins of Meaning*. Oxford: Oxford University Press.
- Kamp, Hans. 1975. Two theories about adjectives. In Edward L. Keenan (ed.), *Formal Semantics of Natural Languages*, 123–155. Cambridge: Cambridge University Press.
- Kaplan, David. 1989. Demonstratives. In Joseph Almog, John Perry & Howard K. Wettstein (eds.), *Themes from Kaplan*, 481–564. New York: Oxford University Press.
- Katz, Jerrold. 1994. Names without bearers. *Philosophical Review* 103, 1–39.
- Katz, Jerrold & Jerry A. Fodor 1963. The structure of a semantic theory. *Language* 39, 170–210.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Larson, Richard K. & Gabriel Segal. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.
- Lewis, David. 1972. General semantics. In Donald Davidson & Gilbert Harman (eds.), *Semantics of Natural Language*, 169–218. Dordrecht: Reidel.
- Longobardi, Giuseppe. 1994. Reference and proper names. *Linguistic Inquiry* 25, 609–665.
- Ludlow, Peter. Forthcoming. *The Philosophy of Generative Linguistics*. New York: Oxford University Press.
- Montague, Richard. 1974. *Formal Philosophy*. New Haven: Yale University Press.
- Parsons, Terence. 1970. Some problems concerning the logic of grammatical modifiers. *Synthese* 21, 320–334.
- Parsons, Terence. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics* (Current Studies in Linguistics 19). Cambridge, MA: MIT Press.
- Partee, Barbara H. 2004. *Compositionality in Formal Semantics*. Oxford: Blackwell.
- Pietroski, Paul. 2002. Function and concatenation. In Gerhard Preyer & Georg Peters (eds.), *Logical Form and Language*, 91–117. Oxford: Oxford University Press.
- Pietroski, Paul. 2003. Quantification and second-order monadicity. *Philosophical Perspectives* 17, 259–298.
- Pietroski, Paul. 2005a. Meaning before truth. In Gerhard Preyer & Georg Peters (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, 253–300. Oxford: Oxford University Press.
- Pietroski, Paul. 2005b. *Events and Semantic Architecture*. Oxford: Oxford University Press.
- Pietroski, Paul. 2006a. Interpreting concatenation and concatenates. *Philosophical Issues* 16, 221–45.
- Pietroski, Paul. 2006b. Character before content. In Judith Thomson & Alex Byrne

- (eds.), *Content and Modality: Themes from the Philosophy of Robert Stalnaker*, 34–60. Oxford: Oxford University Press.
- Pietroski, Paul. 2007. Systematicity via monadicity. *Croatian Journal of Philosophy* 7, 343–374.
- Pietroski, Paul. 2008a. Lexicalizing and Combining. Ms., University of Maryland, College Park.
- Pietroski, Paul. 2008b. Minimal Semantic Instructions. Ms., University of Maryland, College Park.
- Poeppel, David & David Embick. 2005. Defining the relation between linguistics and neuroscience. In Anne Cutler (ed.), *Twenty-First Century Psycholinguistics: Four Cornerstones*, 103–118. Mahwah, NJ: Lawrence Erlbaum.
- Putnam, Hilary. 1975. The meaning of 'Meaning'. In Keith Gunderson (ed.), *Language, Mind and Knowledge*, 131–193. Minneapolis: University of Minnesota Press.
- Rescher, Nicholas (ed.). 1967. *The Logic of Decision and Action*. Pittsburgh, PA: University of Pittsburgh Press.
- Schein, Barry. 1993. *Plurals and Events* (Current Studies in Linguistics 23). Cambridge, MA: MIT Press.
- Schein, Barry. 2002. Events and the semantic content of thematic relations. In Gerhard Preyer & Georg Peters (eds.), *Logical Form and Language*, 263–344. Oxford: Oxford University Press.
- Schein, Barry. 2006. Plurals. In Ernest Lepore & Barry Smith (eds.), *The Oxford Handbook in the Philosophy of Language*, 716–767. Oxford: Oxford University Press.
- Schein, Barry. Forthcoming. *Conjunction Reduction Redux*. Cambridge, MA: MIT Press.
- Segal, Gabriel. 2000. *A Slim Book about Narrow Content*. Cambridge, MA: MIT Press.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, Robert. 1999. *Context and Content*. Oxford: Oxford University Press.
- Stalnaker, Robert. 2003. *Ways a World Might Be*. Oxford: Oxford University Press.
- Stanley, Jason. 2000. Context and logical form. *Linguistics and Philosophy* 23, 391–424.
- Stanley, Jason. 2002. Making it articulated. *Mind and Language* 17, 149–68.
- Stanley, Jason & Zoltán Gendler Szabó. 2000. On quantifier domain restriction. *Mind and Language* 15, 219–261.
- Steedman, Mark. 2000. *The Syntactic Process* (Language, Speech, and Communication). Cambridge, Mass: MIT Press.
- Tenny, Carol L. 1994. *Aspectual Roles and the Syntax–Semantics Interface* (Studies in Linguistics and Philosophy 52). Dordrecht: Kluwer.

Paul Pietroski
University of Maryland
Department of Linguistics
1401 Marie Mount Hall
College Park, MD 20742
USA
pietro@umd.edu