

Overture (*Draft: 19Nov*)

Human children naturally acquire languages that are semantically unbounded. Each of these languages has endlessly many expressions, which exhibit endlessly many meanings. But what are these meanings? We can say, rightly enough, that to understand an expression is to discern its meaning. This suggests that the meaning of an expression is something—perhaps a property or part of the expression—discerned when the expression is understood. Yet the question remains. What does one discern when one discerns the meaning of an expression? We can also frame the issue in terms of acquisition. What does one come to know when one comes to know the meaning of an expression? Or in simpler terms, what are meanings?

To understand the question, it may help to consider some possible answers. Perhaps meanings are concepts, or rules governing the use of expressions in discourse, or functions from contexts to satisfaction conditions, or ways of thinking about such functions, or figments of philosophical imagination, or grommets. In my view, each of these proposals is wrong. The last is obviously wrong, and thus useful for illustration. Meanings are not grommets; for example, the meanings of ‘eager’ and ‘easy’ are not. Whatever one discerns in discerning the meaning of the word ‘eager’, and whatever one comes to know in coming to know what ‘easy’ means, it is not a grommet. The meanings of ‘please’, ‘if’, and ‘laryngitis’ also fail to be grommets, as do the meanings of phrases like ‘eager to please’ and ‘easy to please if you have laryngitis’. And so on. Likewise, I claim, meanings are neither figments nor functions, neither concepts nor rules governing the use of expressions in discourse. But there is something right about each of these suggestions, whereas the grommet theory is just plain wrong.

Proposals that fall under the heading of “Truth Conditional Semantics,” following seminal work by Davidson (196x, 196x) and Montague (197x), require special attention because of their various forms and explanatory virtues. But the basic idea is that words and phrases have the properties ascribed to them by certain theories of truth, which remain only partially specified, and that understanding an expression is a matter of recognizing it as having the relevant truth-theoretic property. In chapter one, I focus on “Davidson’s Conjecture.”

For each naturally acquirable human language **L**, there is a truth theory that can serve as the core of an adequate theory of meaning for **L**.

While this hypothesis proved fruitful, I think it is also wrong. But criticism is not my chief aim. The more important goal is to articulate and defend an alternative to Truth Conditional Semantics that preserves the insights that have been gained by taking Davidson’s Conjecture seriously.

Much of what I say is explicit or implicit in Chomsky (1965, 197x, 1986, 1995, 2000a), who also offered a fruitful idea about the languages that human children naturally acquire.

Such languages, in so far as we can provide theories of them, are stable states of a mental faculty that operates in accordance with constraints that reflect innately determined aspects of human biology.

Like many others, I want to unify this well-supported hypothesis—also reviewed in chapter one—with what is right about Truth Conditional Semantics and the accounts of specific constructions it has inspired. Semanticists have learned much by trying to construct theories of truth for natural human languages. But unification is rarely if ever a matter of simply combining extant theories. And in my view, the following claim is implausible, despite attempts to defend it.

For each stable state **L** of the human language faculty, there is a truth theory that can serve as the core of an adequate theory of meaning for **L**.

Many theorists would reject this last claim, which concerns a particular and perhaps modular aspect of human psychology. Truth and falsity can depend, in open-ended ways that seem to resist description in terms appropriate to theories of a biologically instantiated faculty, on sundry contextual factors deemed relevant by reasonable communicators. That is one way of expressing a Wittgensteinian/Austinian point of agreement between Chomsky and Davidson. But this point also encapsulates a real tension between Davidson's Conjecture and Chomsky's mentalistic conception of human languages. I'll argue that we can't have both, and that we should preserve the latter by replacing Davidson's Conjecture with an "internalist" alternative, according to which meanings do not determine truth (reference, or satisfaction) conditions.

To advertise, I think meanings are instructions to construct concepts of a special sort. The interest and plausibility of this claim depend on how it is spelled out. But that is true of any equally compact and remotely plausible proposal about what meanings are. More controversially, I claim that meanings are instructions—generated by the human language faculty—to construct *systematically conjoinable monadic* concepts from available mental resources that include many concepts of varying adicity. Meanings are still features of expressions, in my view, since words and phrases are pairs of semantic and phonological instructions. Abstracting from phonology, words and phrases *are* meanings: expressions are instructions to construct monadic concepts. Since that proposal may not seem remotely plausible, pages follow.

In the rest of this chapter, I provide an overview of the book by discussing some issues about human cognition that will be addressed via the idea that meanings are concept construction instructions. One reason for starting this way is that many rational primates do not find the study of linguistic meaning intrinsically interesting. I sympathize. Moreover, while justifying theories in this domain requires attention to details that reflect distinctively human linguistic talents, theorists can be motivated by the hunch that meanings connect such talents with older aspects of cognition that humans share with other animals. I share this hunch. For readers who prefer less stage setting, section one may be enough. Subsequent sections can be skimmed, and perhaps used later as a summary of main points. For specialists, the following paragraph may suffice.

Meanings combine in ways that permit only certain kinds of expressions. When a concept C is lexicalized, a linguistic signal must be linked to a concept with a "legible format," even if this requires creation of a formally distinct but analytically related concept of a legible type. I think a lexical meaning just is an instruction to fetch a concept with a certain formal character, while a phrasal meaning just is an instruction to fetch and combine concepts, in a way that is implemented by human biology. To this point, many theorists could agree. But in my view, predicate-conjunction is the only semantic composition operation; see Pietroski (2005, 2006). Each lexical item is of the type exhibited by monadic predicates, and phrases are understood as conjunctions of such predicates. In which case, lexical meanings are instructions to fetch monadic concepts, and phrasal meanings are instructions to conjoin fetchable concepts. Such instructions can be *modeled* truth-theoretically, but only if we idealize away from two ubiquitous features of natural language: a single word can be linked, polysemously, to multiple concepts; and when a sentence is uttered, the truth or falsity of the utterance often depends on aspects of the communicative situation not indexed with any element of the sentence. From this perspective, linguistic expressions are not representations that have truth theoretic properties. Meaningful expressions are instructions to construct concepts that can be used in thought.

1. Meanings Fit Together

Children have finite cognitive resources and limited experience. Yet they acquire languages that are semantically unbounded. This raises questions. Claims about meaning should be compatible with plausible answers. And that turns out to be a real constraint on such claims. In this section, I rehearse the first of two important morals: meanings compose, and they somehow relate words to concepts; see Fodor (ouvre). Meanings *fit* in ways that lots of things don't. So if we want to know what meanings are, perhaps we should ask how and what they fit.¹

In acquiring a language, one acquires a capacity to understand expressions of that language. Understanding the expressions of a spoken language is, to a first approximation, a matter of associating certain sounds with certain meanings. Indeed, a child who acquires English acquires a capacity to associate boundlessly many sounds with meanings. This suggests, and perhaps requires, that the languages children naturally acquire are semantically compositional in the following sense: each such language has finitely many lexical (atomic) expressions, each with a lexical meaning, along with finitely many modes of combination; and the meaning of any nonlexical (complex) expression Σ is determined by the meanings of lexical constituents of Σ , given how those constituents are combined in Σ . One expects languages like English to respect even more demanding composition constraints, since humans have *very* finite cognitive powers, and childhood experience is *very* limited.² It is, however, hard to formulate strong principles in a way that avoids counterexamples. The weak principle has the virtue of maximizing plausibility.

It is also suggested, and perhaps required, by the semantic systematicity of languages like English. The meaning of 'lawyer from London who needed a doctor in Boston' is related, in a certain way, to the meanings of its constituents. The meaning of 'doctor from Boston who needed a lawyer in London' is related, in a similar way, to the meanings of its constituents. And while someone can understand 'doctor from Boston' without understanding 'lawyer in London', say because of a failure to understand 'lawyer' or 'London', anyone who understands 'lawyer from London who needed a doctor in Boston' can understand 'doctor from Boston who needed a lawyer in London'; see, e.g., Fodor and Pylyshyn (199x), Fodor and Lepore (200x). Such systematicities further confirm that languages like English are semantically compositional.

I mention this familiar point not to alleviate doubt, but to highlight an explanandum: *why/how* are such languages semantically compositional? As we'll see, plausible answers constrain the space of plausible conceptions of meaning. But first, let me clarify the question, in terms of two toy languages that are semantically compositional for different reasons.

Both languages have 2600 atomic symbols—'A', 'B', 'C', ... , and 'Z'—each a monadic predicate. The predicates are satisfied by armadillos, blue things, crocodiles, ... , and zesty things. In both languages, there is exactly one way of combining symbols: if ' α ' and ' β ' are symbols, so is ' $[\alpha \beta]$ '. The complex symbols thus include ' $[A B]$ ', ' $[B Y]$ ', ' $[[A B] Z]$ ', etc. In the first language, Andish, combination indicates predicate-conjunction. But in the second, Orish, combination indicates predicate-disjunction. In Andish, ' $[A B]$ ' is a predicate satisfied by an entity e iff e is an armadillo *and* e is blue; ' $[[A B] Z]$ ' is satisfied by e iff (e is an armadillo *and* e is blue) *and* e is zesty; etc. In Orish, ' $[A B]$ ' is satisfied by e iff e is an armadillo *or* e is blue; ' $[[A B] Z]$ ' is satisfied by e iff (e is an armadillo *or* e is blue) *or* e is zesty; etc. That's it. These satisfaction conditions are the only meanings the symbols have. Both languages are semantically compositional, trivially. But this compositionality is "realized" differently in the two languages.³

We can imagine a contrasting language, Calendrish, in which combination indicates predicate-conjunction on odd days and predicate-disjunction on even days. So there may be a sense in which Andish and Orish are both semantically compositional “because” in each language, the meaning of each complex symbol is determined by the meanings of constituent atomic symbols and how those symbols are combined *independent of the date*. Still, there is a more interesting sense in which Andish is semantically compositional because its sole mode of combination signifies predicate-conjunction, while Orish is semantically compositional because its sole mode of combination signifies predicate-disjunction. Imagine a theorist who knew nothing about how these languages worked, but was told merely that Andish was semantically compositional and regularly acquired by children. Such a theorist might ask questions.

Why does Andish have this character? Whatever the reason, is Orish the same in this respect? If not, is Orish semantically compositional in another way?

If so, is Orish equally acquirable by children who can acquire Andish? And if not, is that because Orish is not semantically compositional in the way that Andish is?

If so, could it be that children can only acquire languages that are semantically compositional in certain ways? And could that be because only certain modes of composition are implementable by the biology that lets children acquire Andish?

Such questions would call for substantive proposals about why (or how) the meanings of complex Andish symbols are determined by the meanings of constituent atomic symbols.

Let’s return now to the real question of why languages like English are semantically compositional.⁴ And by way of eliminating the vague phrase ‘like English’, let me offer some terminological suggestions that will prove useful below. Say that a *language* associates signals of some kind with interpretations of some kind. This allows for bee languages, invented formal languages, and languages of thought, along with English, ASL, and Japanese. It also leaves room for various conceptions of what languages are. Many “things” can be said to associate signals with interpretations: sets of signal-interpretation pairs; rules or algorithms of an appropriate sort; mechanisms that implement such algorithms; etc. But put these nice distinctions aside for now. Let a Language be a language with overt signals—like the sounds of English, or signs of ASL—that is acquirable by any biologically normal human child in conditions of experience that are not atypical for members of our species. In short, Languages are languages that kids are apt to acquire. (Capitalization indicates a restricted and perhaps technical variant of an ordinary term.) To review: I assume that Languages are semantically compositional, and hence that expressions of a Language have meanings; we’re heading for the claim that these meanings somehow fit together and fit concepts; and this is to motivate, or at least explain, the idea that these meanings are instructions to construct *systematically conjoinable monadic* concepts.

I suspect that Languages employ only a few modes of combination. Initially, it might seem otherwise, since we describe such modes in many ways: combining a verb like ‘stab’ with a name like ‘Caesar’; combining a name with a phrase like ‘stab Caesar’; combining a tensed verb like ‘saw’ with a clause like ‘Brutus stab Caesar’; combining a determiner like ‘every’ with a noun like ‘senator’; combining however one does to form the noun phrase ‘senator who saw Brutus stab Caesar’; etc. But there are underlying generalizations, across “construction types,” that can be described with theories (grammars) that describe Languages economically. It is an empirical question how deep these generalizations go. Answering requires an act of balancing

various kinds of descriptive adequacy and theoretical illumination. But if explanation is the goal, one looks for simplicity, if only to find out whether and to what degree the world will cooperate. And in my view, developments in the study of transformational grammar have been successful enough to warrant optimism in this regard.⁵

Moreover, *if* there are more than a few grammatically independent modes of combination across Languages, some of the relevant grammatical distinctions may be semantically irrelevant. For example, many accounts of semantic composition do not distinguish adjectival and adverbial modification, at least for simple cases like ‘red ball’ and ‘ran quickly’. Many accounts also unify treatments of verb-object and determiner-noun combinations, as in ‘saw Bill’ and ‘every senator’. Even if these grammatical distinctions are not theoretically superficial, they may be insignificant, in that one semantic operation applies across different constructions. Predicate-conjunction may apply when a nounish or verbish (entity or event) predicate combines with an adjectival or adverbial modifier, while predicate-saturation applies when a verb or determiner (first-order or second-order predicate) combines with a noun. Such theories may be wrong. But available proposals are good enough, in my view, to make the following hypothesis plausible: in Languages, there are only a few semantically significant modes of combination, corresponding to at most a few—perhaps just one or two—operations that relate meanings to meanings in certain ways; and Languages are semantically compositional *because* the meanings of complex expressions are determined by the meanings of simpler expressions in *these* ways.⁶

In which case, meanings don’t merely compose, they compose in specific ways that might be described with a formula or two. I return to this point from several angles. (If it already seems anodyne, so much the better.) With regard to explaining semantic compositionality, the idea is that Languages are relevantly like Andish and Orish, though more complicated. Whatever lexical meanings are, they can be combined in certain ways. And whatever the meaning of a complex expression is, it is a “result” of combining simpler meanings in a certain way. So if we can figure out *how* meanings compose, we might thereby come to know something substantial about what meanings are. For meanings, lexical or not, can only be what modes of composition let them be.

Let me illustrate with a fable that will soon become a proposal. Suppose that our imagined theorist—call him Harold—somehow learns that in Andish, predicate-conjunction is the sole semantic composition operation. This discovery might lead to rejoicing,

Now I know why Andish is semantically compositional! I also know that every expression of Andish is a monadic predicate! For no other expressions could be semantically combined in the way that Andish expressions are combined.

followed by more sober reflection.

So the meanings of Andish expressions are all of a type such that meanings of this type are predicatively conjoinable. But what are these meanings? They might be concepts, or satisfaction conditions, or things as yet undreamt of. Yet whatever they are, I can say that expressions of Andish have meanings of type $\langle e, t \rangle$; where an entity designator has a meaning of type $\langle e \rangle$, and a truth-evaluable sentence has a meaning of type $\langle t \rangle$. In short, each Andish expression has a meaning that may be combined with the meaning of a singular term to yield a sentential meaning.

The next thought threatens to spoil all.

Yet if each expression is so understood, then Andish has neither singular terms nor sentences. And how could that be so? I have seen native speakers refer to and make diverse claims concerning armadillos, crocodiles, and other entities.

But our Harold, leary of sacrificing a generalization to accommodate observations, conjectures.

Still, perhaps Andish *itself* has only expressions of type $\langle e, t \rangle$. When speakers seem to be using sentences, maybe they are using predicates *in combination with* a conceptual operator like existential closure. And when speakers seem to be using singular terms, they may be using predicates along with other operators. Russell (19xx), Davidson (196x), and Kaplan (198x) may have been onto something.

Conjecture in hand, Harold plans its defense.

I will look for symptoms of covert operators, and argue that children do not *learn* how meanings fit together. For if children presuppose that words are predicates to be conjoined, nature might thereby lead them to perceive discourse as the use of overt expressions *in concert with* some independently available concepts. I will be accused of ascribing mental languages to speechless animals, and rightly so. For if Andish expressions have no constituents of type $\langle e \rangle$ or $\langle t \rangle$, then the judgments expressed do. So I will confess in advance, directing critics to read Fodor (19xx) and Gallistel (19xx), references therein, and so on. This will at least buy me time.

The methodological moral is bland. One can try to articulate and defend a thesis about meaning by starting with a hypothesis about semantic composition, and then exploring its consequences for lexical meanings and the relation of meaning to thought. That's what I will do. More controversially, I will argue that Languages are fundamentally like Andish. In the next section, I try to make this claim (about how meanings fit together) seem less crazy and worth defending. In sections three and four, I return to the questions of what meanings fit, and how they are related to concepts. Providing a few details will then bring us back to the idea that meanings are instructions to construct conjunctive monadic concepts from prelexical mental resources.

One might wonder what any of this has to do with Truth Conditional Semantics. But if (i) meanings compose in a simple Conjunctivist fashion, and (ii) Truth Conditional Semantics is correct, then (iii) expressions of a Language have truth-theoretic properties that compose in a simple Conjunctivist fashion. So if (i) is plausible, but (iii) is not, that tells against (ii). I plan to press this argument. Hence, I need to motivate (i) without presupposing that (ii) is false. So some of this book is, despite its title, devoted to defending a Conjunctivist semantics that seems—and initially is—*friendly* to the idea that sentences have compositionally determined truth conditions. But later, I will argue that sentential meanings do *not* determine (functions from contexts to) truth conditions, and that this negative thesis is especially plausible if meanings compose in a simple Conjunctivist fashion. A side benefit of this Trojan Horse strategy is that readers unpersuaded by the arguments against (ii) may still find something of value here. Some may be persuaded to take (iii) seriously, given the arguments in favor of (i). Others may reject (iii), retain (ii), and at least be bothered by the arguments in favor of (i). My own views are not firm. In the end, I see the conditional—if (i) and (ii), then (iii)—as a reason for rejecting (ii). But I don't pretend that this response is dictated by currently available evidence.

2. Proposal: Meanings Conjoin, Almost like Thoughts

Many semanticists assume that combining expressions, in a Language, often signifies predicate-conjunction.⁷ Simple modifications of nouns, as in ‘red ball’ and ‘ball that fell’, are the most obvious motivating examples; though as we’ll see, there are many others. This invites a stronger hypothesis, defended in Pietroski (2002, 200x, 2005, 2006): predicate-conjunction is the *only* semantic composition operation; appearances to the contrary are due to (interactions of meanings with) *other* cognitive factors that are germane to truth or communication. On this view, all complex expressions are conjunctions of monadic predicates, even determiner phrases like ‘every senator’. As with any minimizing/reductive proposal, the immediate challenges concern basic descriptive adequacy. It needs showing that one can, without *ad hoc* devices, redescribe the elementary facts and initial generalizations that have led semanticists to posit other modes of combination. But formally, this is doable homework; see chapter two. Illustrations follow.

The harder and more interesting task is to motivate such redescription, in ways that go beyond mere appeal to simplicity, by using the sparer conception of semantic composition to explain some otherwise puzzling facts. This is the burden of chapter three. Though in a broader sense, the whole of this book together with Pietroski (2005) is an attempt to motivate “Conjunctivism” and its implications. Other accounts of semantic composition also require motivation, given an alternative. But even when comparing more familiar views with Conjunctivism, my primary aim is to argue for the latter, not to criticize the former.

If combining expressions of a Language always signifies predicate-conjunction, there should be many symptoms of this. And some of these symptoms, if found, would be surprising given more permissive conceptions of semantic composition. For example, if Conjunctivism is correct, Languages should abhor expressions of type $\langle e \rangle$ and $\langle t \rangle$ *at least for purposes for phrasal combination*. The italicized qualifier is required on my view, which is not maximally reductive. But even the qualified claim is strong enough to warrant comment.

The idea of treating apparent entity-designators as misleading expressions, of a type other than $\langle e \rangle$, is well known from Russell (191x), Quine (195x), Montague (197x), and many others. For now, I simply report that there are many ways of analyzing a name like ‘Socrates’ as a monadic predicate, of type $\langle e, t \rangle$, that applies in a given conversation to the intended bearer of the name; section four provides a sketch. Chapter three develops arguments from Burge (197x) and others in favor of such analyses. And if names like ‘Socrates’ are *not* lexical items of type $\langle e \rangle$, this raises the question of why not. Why are the relevant name-sounds *not* paired with (singular) concepts in ways that would yield grammatically atomic singular terms?

In this context, I also discuss the paucity of “supratransitive” verbs that can combine with three or more grammatical arguments to form a complete sentence. There are many supradyadic concepts that can combine with three or more singular concepts to form a complete thought. (Think of the concepts BETWEEN and SELL.) But in Languages, there are few if any lexical verbs of type $\langle e, \langle e, \langle e, t \rangle \rangle \rangle$, $\langle e, \langle e, \langle e, \langle e, t \rangle \rangle \rangle \rangle$, etc. Why? Moreover, there are independent reasons for analyzing ditransitive verbs and at least many transitive verbs—like ‘give’, ‘kick’, ‘sell’ and ‘kill’—as semantically monadic predicates. This raises the general question of why children *don’t* acquire far more lexical items that are *not* of type $\langle e, t \rangle$. My answer will not surprise: meanings can only be what semantic composition operations let them be; and in Languages, the only such operation is predicate-conjunction.

Conjunctivism predicts that words impose a “monadic format” on concepts, and hence that nonmonadic concepts are regularly lexicalized with expressions of type $\langle e, t \rangle$ —with consequent need for some devices that can be used to indicate relationality and singularity. This highlights the question of why relations are so often expressed, in Languages, with functional (“closed class”) expressions like prepositions, and why singular concepts are expressed with *indices* as opposed to lexical labels for nameable things. These features of Languages are not expected, given composition operations that permit lexical singular terms and polyadic predicates. But given Conjunctivism, such features make sense. Yet even if entity designators are misleading, and verbs tend to be semantically monadic, one might balk at the idea that a complex expression is misleading if it seems to combine expressions at least one of which is of type $\langle t \rangle$.

Consider the negation and connective in ‘Socrates is *not* mortal *or* Socrates is mortal’. How could this sentence be understood, as the truism it is, with no constituent of type $\langle t \rangle$? I return to connectives, which are easy to deal with, but will say something here about negation. (Specialists may be curious/skeptical.) If you can wait until chapter two for a more relaxed and systematic presentation, or just don’t care, the next several paragraphs can be skimmed.

Suppose it could be shown that “atomic sentences,” without negation or connectives, all exhibit the logical form ‘ $\exists\Phi$ ’; where ‘ Φ ’ is a monadic predicate, typically conjunctive, and ‘ \exists ’ is a kind of existential quantifier. To keep things simple—and for now, fully compatible with Truth Conditional Semantics—suppose that ‘Socrates is mortal’ has the logical form ‘ $\exists[S\bullet M]$ ’; where ‘ \bullet ’ signifies predicate-conjunction, ‘ M ’ is satisfied by each mortal and nothing else, and Socrates is the unique satisfier of ‘ S ’. If it helps, think of ‘Socrates is mortal’ as having the (idealized) *grammatical* form ‘ $S\wedge M$ ’, in which predicates are *combined*; and suppose that a grammatical form can be mapped, in accord with severe constraints, to a *logical* form in which predicates are *conjoined*.⁸ But if ‘S(ocrates)’ is formally of type $\langle e, t \rangle$, as opposed to $\langle e \rangle$, let’s not rush into assuming that ‘ $\exists[S\bullet M]$ ’ is of type $\langle t \rangle$ as opposed to $\langle t, t \rangle$. Logical forms can be predicative.

We can define ‘ \exists ’ so that ‘ $\exists[S\bullet M]$ ’ is a *predicate of truth values* satisfied by truth (\top) iff ‘ $S\bullet M$ ’ is satisfied by something. As Tarski (193x) taught us, truth can be viewed as a special case of satisfaction. But any expression with a satisfaction condition, satisfied by things of type α , can be viewed as a predicate of type $\langle \alpha, t \rangle$. (That is why there are so many ways of analyzing names as predicates.) So a sentence *can* be treated as an expression of type $\langle t, t \rangle$.⁹ And from this perspective, negation looks like modification—like ‘red’ in ‘red ball’.

Assume exactly two truth values, \top and \perp . Define ‘ \neg ’ as a predicate, of type $\langle t, t \rangle$, satisfied by and only by \perp . Then ‘ $\neg\exists[S\bullet M]$ ’ is satisfied by something iff that thing satisfies both ‘ \neg ’ and ‘ $\exists[S\bullet M]$ ’. Hence, ‘ $\neg\exists[S\bullet M]$ ’ is satisfied by \perp or by nothing, depending on whether or not \perp satisfies ‘ $\exists[S\bullet M]$ ’; and ‘ $\exists[S\bullet M]$ ’ is satisfied by \perp iff ‘ $S\bullet M$ ’ is not satisfied by anything. It follows that ‘ $\exists[\neg\exists[S\bullet M]]$ ’ is satisfied by \top iff ‘ $S\bullet M$ ’ is not satisfied by anything—just as if ‘ \exists ’ signified satisfaction of the attached predicate, while ‘ $\exists[\neg\exists]$ ’ signified lack of satisfaction. The key steps in that reasoning are repeated below.

\top satisfies ‘ $\exists[\neg\exists[S\bullet M]]$ ’ \equiv something satisfies ‘ $\neg\exists[S\bullet M]$ ’
something satisfies ‘ $\neg\exists[S\bullet M]$ ’ \equiv \perp satisfies ‘ $\exists[S\bullet M]$ ’
 \perp satisfies ‘ $\exists[S\bullet M]$ ’ \equiv nothing satisfies ‘ $S\bullet M$ ’

Following Tarski, we can think of ‘ $\exists[\Phi]$ ’ as a predicate satisfied by everything or nothing, depending on whether or not ‘ Φ ’ is satisfied by something. So sentential negation can be

analyzed as a monadic predicate (satisfied only by \perp) that conjoins with a sentential predicate, followed by a kind of existential closure—much as a name can be analyzed as a monadic predicate (satisfied by only one thing) that can be conjoined with another, followed by the same kind of closure. This requires two “closure cycles” for ‘Socrates is not mortal’, and sentential scope for the negation. But that is not crazy, even if turns out to be wrong.

Chapter two includes Conjunctivist analyses of other “sentence involving” expressions. Relative clauses, connectives like ‘or’, and determiners like ‘every’ can be analyzed as monadic predicates. This requires nothing more exotic than existential quantification over ordered pairs of entities and truth values, given an independently motivated tolerance for plural variables as in Boolos (1998); where a plural variable has more than one value relative to each assignment of values to variables. Teaser: determiners are plural monadic predicates satisfied by *ordered pairs* of the form $\langle e, \top \rangle$ or $\langle e, \perp \rangle$, and this explains the *conservativity* of determiners. Promissory note: the appeal to truth values turns out to be dispensable convenience.

It will be noticed that ‘ \exists ’ is not of type $\langle \langle e, \mathbf{t} \rangle, \mathbf{t} \rangle$, like Frege’s (1879) quantifiers. Rather, ‘ \exists ’ maps predicates (of entities or truth values) to predicates of truth values—as if one cognitive system that recursively conjoins monadic predicates, via ‘•’, regularly interfaces with another system that traffics in representations that can be evaluated in a binary fashion. But we can stipulate that ‘ $\exists e$ ’ is of type $\langle \langle e, \mathbf{t} \rangle, \mathbf{t} \rangle$, and that expressions of type $\langle \mathbf{t} \rangle$ include ‘ $\exists e[M(e)]$ ’, ‘ $\sim \exists e[M(e)]$ ’, and ‘ $\exists e[S(e) \& M(e)]$ ’; where ‘&’ and ‘ \sim ’ signify (propositional) conjunction and negation.¹⁰ This useful for purposes of explicitly formulating a metalanguage in which to offer hypotheses about Languages. But definition cannot ensure that any expression of a Language is of the same type as a *thought* of type $\langle \mathbf{t} \rangle$, just as definition cannot ensure that any expression of a Language is of the same type as a *singular concept* of type $\langle e \rangle$. As discussed in more detail below, it is an empirical question how expression-types are related to concept-types.

Let me stress: I am not denying Frege’s (189x) insight that constituents of *thoughts* (*Gedanken*) exhibit a saturated/unsaturated distinction; and I agree that thoughts themselves belong on the left side of this divide. Grant that thoughts and singular concepts are “representationally complete” in some way that predicative concepts, which map complete representations to thoughts, are not. Grant also that this important distinction is usefully marked by saying that thoughts are of type $\langle \mathbf{t} \rangle$, that singular concepts are of type $\langle e \rangle$, and that all “unsaturated” thought constituents are of some higher type defined in terms of the basic types $\langle e \rangle$ and $\langle \mathbf{t} \rangle$. None of this precludes the hypothesis that each expression of a *Language* is understood as a monadic predicate of type $\langle e, \mathbf{t} \rangle$. On the contrary, we can use the standard metalanguage to formulate the restrictive hypothesis that even names and sentential expressions of a Language are understood as expressions at an initial level of a *much* larger type hierarchy.

That’s enough technicality for now. Let’s go back to the hypothesis that all sentences exhibit the logical form ‘ $\exists[\Phi]$ ’; where ‘ Φ ’ is a monadic predicate, typically conjunctive. Repeating the example, suppose that ‘Socrates is mortal’ has the logical form ‘ $\exists[S \bullet M]$ ’, deferring further fuss about exactly what kind of existential closure ‘ \exists ’ provides. It will do, for the moment, if ‘ $\exists[S \bullet M]$ ’ just means that something is both Socrates and mortal. For if this illustrates the general case, then a sentence is more like a discourse—in which a series of thoughts may be presented for some purpose—than one might have thought. I end this section with some brief remarks on this score, returning to the point in chapter three.

A sentence is not a mere list of words, much as a judgment is not a mere juxtaposition of ideas, and a fact is not a mere concatenation of things; see Frege (189x), Kant (18xx), Plato (xxxx). Similarly, a discourse is not a mere list of sentences, or thoughts, or bits of information. At least in typical cases, a discourse has a conjunctive unity. To be sure, a discourse like ‘A man arrived. He sang. He danced. He left.’ is not equivalent to a mere conjunction of independent propositions: $\exists e[\text{Man}(e) \ \& \ \text{Arrived}(e)] \ \& \ \exists e[\text{DemonstratedMale}(e) \ \& \ \text{Sang}(e)] \ \& \ \dots \ \& \ \text{Left}(e)$; see Kamp (197x), Heim (197x), and much work since. But the discourse implies the first conjunct *and something more*. Likewise, the sentence ‘Socrates is mortal’ implies more than ‘Socrates is something, and something is mortal’. Yet each conjunct is implied.

Inference is a kind of discourse. And as medieval logicians noted, the default direction of inference is from ‘red ipod’ to ‘ipod’, as in ‘That’s a red ipod, and hence an ipod’. It takes extra work, like negation, to go the other way, as in ‘That’s not an ipod, and hence not a red ipod’. Davidson (196x) extended the point to adverbial modification by positing event variables: if ‘Caesar died yesterday’ means that there was an event of Caesar dying *and* it occurred yesterday, then there are important analogies between ‘red ipod’ and ‘died yesterday’. Following Parsons (1990), Schein (1993), and others, I think the analogies extend beyond adjectives and adverbs to grammatical *arguments*: ‘Brutus killed Caesar yesterday while senators shouted’ means roughly that there was a killing—an event that ended in a death—such that Brutus did it, it befell Caesar, it happened yesterday, and it happened while senators shouted. From this perspective, an action report is a mini-discourse on who did what amidst the rest of the passing show.

A discourse may also turn abstract. It may start with ordinary reference to people and aspects of their environment, but shift to quantification over less obvious “things.” And to advertise a larger theme, not all abstraction is in the service of getting at Reality. Lexicalization and composition often saddle speakers with abstraction, but I think that Languages are aimed at the concepts we use to think about things, not the things we think about. As discussed in chapters four and five, the variables in logical forms start looking pretty weird if one regards them as ranging over a mind-independent domain, as if ordinary sentences were parts of a scientific discourse. Consider ‘We enjoyed the deep blue of the sky’. Now ask which things plausibly and compositionally satisfy ‘sky’, ‘blue’, ‘deep’, ‘deep blue’, ‘deep blue of the sky’, ‘enjoyed’ and ‘enjoyed the deep blue of the sky’? On the other hand, Linguistic variables seem quite natural if expressions are compositionally determined instructions to construct concepts. Or so I’ll argue.

From this perspective, Linguistic abstraction is not a reflection of Reality, but a response to demands imposed by simplicity of composition. If combining lexical items always signifies predicate-conjunction, then lexical items have to cooperate by imposing a monadic format on nonmonadic concepts. That is the main point of chapters two and three. In some cases, we may be fortunate, in that internal demands lead to expressions that “fit the world” in addition to fitting together. But in general, we are lucky if our *concepts* fit the world in ways that can—with work, in contexts—lead to thoughts that have genuine truth conditions and are occasionally true. We are very lucky to have Languages, but Linguistic expressions do not fit each-other-and-the-world in the ways required for truth evaluability. That is the main point of chapters four and five.

There was no guarantee that nature would give rise to truth-evaluability at all. And it is hard enough to explain how humans can have *thoughts* with the logical structure required for genuine truth conditions, even without requiring that such thoughts (or their truth conditions) be

compositional in the same way(s) that Linguistic expressions are. Humans may use meaningful expressions to construct concepts that can combine in ways that meanings cannot. I want to explore this possibility, in part to see if appeal to Conjunctivist meanings can help explain why human cognition is so unified, and how biology arranged for a Fodorian Mentalese so much like Frege's *Begriffsschrift*. But now things have gotten very abstract, and too far ahead.

Returning to the theme of this section, I mentioned some analogies between sentences and discourse, in suggesting that meanings conjoin *almost* like thoughts do. Thanks to Tarski, we can coherently formulate the hypothesis that Languages are a lot like Andish. But my claim is not that meanings compose as concepts do. On the contrary, concepts have various adicities, and rendering these distinctions invisible is one important job for lexicalization. Explaining this last claim, which may seem odd, is the next task. We are also sneaking up on Davidson's Conjecture. For one might wonder: is there, for each Language, a theory of *truth* with predicate-conjunction as the only mode of semantic composition? And if Conjunctivism is correct, then as we're about to see, lexical meanings are not friendly to Truth Conditional Semantics.

3. Meanings Fit Concepts

Meanings can only be what composition operations let them be. Meanings also relate, somehow, to concepts. But meaning-concept relations can only be what composition operations let them be. So different conceptions of semantic composition will often have different implications for meaning-concept relations. And we can try to find out which implications, if any, are correct.

It is useful, at least for expository purposes, to describe possible bundles of semantic composition operations in terms of the lexical meanings they allow. (For simplicity, let's assume that these operations are all binary, in the sense of mapping pairs of meanings to meanings.) At one end of the continuum are limiting case "bundles" consisting of just *one* operation, like predicate-conjunction, that is geared to a *single* semantic type: for each such operation **O**, there is a type $\langle\alpha\rangle$ such that **O** makes it possible to combine two meanings of type $\langle\alpha\rangle$, yielding a meaning of type $\langle\alpha\rangle$; and that's the only kind of combination that **O** supports. Another kind of operation licenses combination of two specific meaning types, with a further distinction being whether or not such combination yields novel meaning types. We might imagine three languages governed, respectively, by the three principles below.

$\langle\text{red}\rangle^{\wedge}\langle\text{red}\rangle \implies \langle\text{red}\rangle$ $\langle\text{red}\rangle^{\wedge}\langle\text{blue}\rangle \implies \langle\text{red}\rangle$ $\langle\text{red}\rangle^{\wedge}\langle\text{blue}\rangle \implies \langle\text{green}\rangle$

And of course, one can imagine languages with two or more such composition principles. But in any such language, the composition operations impose severe constraints on the kinds of meanings that lexical items can exhibit. In terms of acquisition, if one knows the composition operations, one knows what kind(s) of meanings lexical items need to have, in order to compose.

Other operations, like function-application, are more flexible.¹¹ For *any* pair of meanings of the type $\langle\alpha\rangle$ and $\langle\alpha, \beta\rangle$, function-application maps that pair of meanings onto a meaning of type $\langle\beta\rangle$. In one sense, this operation adds nothing new. For while a complex expression can be of a type not exhibited by any lexical expressions—there may be words of type $\langle e \rangle$ and $\langle e, t \rangle$, but none of type $\langle t \rangle$ —the meaning of a complex expression has to be *simpler*, in type, than one of its constituents. Put another way, function-application is a "type reducing" operation. Such an operation imposes no specific constraints on what lexical meanings can be, though it may impose constraints on a lexicon. If $\langle\alpha\rangle$ is a basic type, not of the form $\langle\beta, \gamma\rangle$, then composition by function-application requires that lexical meanings of type $\langle\alpha\rangle$ be accompanied by lexical

meanings of type $\langle\alpha, \beta\rangle$, for at least one type $\langle\beta\rangle$. Likewise, meanings of type $\langle\alpha, \beta\rangle$ must be accompanied by meanings of type $\langle\alpha\rangle$ or $\langle\langle\alpha, \beta\rangle, \gamma\rangle$ for some type $\langle\gamma\rangle$. But by itself, this says nothing about the kinds of meanings that lexical items can (or cannot) exhibit.

If $\langle e\rangle$ and $\langle t\rangle$ are expression types in a language whose sole composition operation is function-application, then the language has lexical types that *determine* $\langle e\rangle$ and $\langle t\rangle$. But nothing follows about what the lexical types are. They might be $\langle red, e\rangle$, $\langle red\rangle$, $\langle blue, t\rangle$, and $\langle blue\rangle$. Instances of $\langle\alpha, e\rangle$, $\langle\alpha\rangle$, $\langle\beta, t\rangle$, and $\langle\beta\rangle$ can generate expressions of types $\langle e\rangle$ and $\langle t\rangle$; where α and β can be complex types, like $\langle\langle t, e\rangle, \langle t, e\rangle\rangle$ and $\langle\langle e, t\rangle, \langle e, t\rangle\rangle$. Similarly, if you know there are complex expressions of type $\langle t\rangle$ with lexical constituents of type $\langle e\rangle$, you may suspect that $\langle e, t\rangle$ is a lexical type. But composition can be achieved via other types, which need not be lexical, of the form $\langle e, \alpha\rangle$, $\langle\alpha, \beta\rangle$, and $\langle\beta, t\rangle$.

One can list the permissible lexical types, and stipulate that there are no others. But then the constraints come entirely from the list, with the operation (function-application) simply imposing global constraints of “fit,” ruling out lexicons whose types do not include pairs of the form $\langle\alpha\rangle$ and $\langle\alpha, \beta\rangle$. Alternatively, a language might allow for endlessly many types, specified recursively, as in the two examples below.

$\langle e\rangle$ and $\langle t\rangle$ are types; and
 if $\langle\alpha\rangle$ and $\langle\beta\rangle$ are types, so is $\langle\alpha, \beta\rangle$
 $\langle red\rangle$, $\langle blue\rangle$, and $\langle green\rangle$ are types; and
 if $\langle\alpha\rangle$, $\langle\beta\rangle$ and $\langle\gamma\rangle$ are types, so is $\langle\alpha, \langle\beta, \gamma\rangle\rangle$

Such conditions can be supplemented with a clause implying that there are no other types. In which case, the “function-application rooted in $\langle e\rangle$ and $\langle t\rangle$ ” (FRET) conception of permissible types rules out endlessly many types— $\langle red\rangle$, $\langle red, blue\rangle$, etc.—and permits endlessly many: $\langle e, t\rangle$, $\langle e, \langle e, t\rangle\rangle$, ... , $\langle t, t\rangle$, $\langle t, \langle t, t\rangle\rangle$, ... , $\langle e, e\rangle$, $\langle e, \langle e, e\rangle\rangle$, ... , $\langle e, \langle\langle t, e\rangle, \langle t, e\rangle\rangle\rangle$, etc.

The FRET conception is more descriptively powerful than the predicate-conjunction (PC) conception, since the latter only permits the type $\langle e, t\rangle$. (By contrast, the “function-application rooted in $\langle red\rangle$, $\langle green\rangle$, and $\langle blue\rangle$ ” conception is orthogonal, since the endlessly many types it permits do not include $\langle e, t\rangle$.) Putting this obvious point in another way, the FRET conception is based on a particular choice of basic types, and the PC conception follows from the nature of the operation. In themselves, these implications are neither good or bad. If the goal is to explain features of Languages, one might worry about the descriptive power of FRET, especially if it is part of a still more permissive conception of semantic composition that also includes PC. But here, my aim is simply to compare two conceptions of lexicalization, with the aim of making their predictions clear. So I return to the conceptions after some remarks about lexicalization.

I assume that at least typically, for a child acquiring a spoken Language, lexicalization is a process in which a word-sound gets linked to one or more independent mental representations. We humans may occasionally invent words, like ‘brillig’, and get them into a Language without ever associating them (in any stable way) with ways of thinking about things. Perhaps one can acquire ‘quark’ or ‘eigenvalue’ while waiting for an independent representation that never comes. But these are presumably special cases. When a child acquires the word ‘dog’, upon hearing it pronounced repeatedly in the presence of a dog that has captured her interest, the child is representing the dog. If another child associates the sound of ‘dog’ with a representation of a chair, or with nothing in particular, he does not thereby acquire the word ‘dog’ with its English

meaning. Acquiring this word requires, at the very least, associating the relevant sound with a dog-representation of the right sort. It's not enough to represent a salient dog as Fido, or as something brown, or as brown-furry-nice-woof-scary. The lexicalizer has to represent a dog *as* a dog, or something very like that, whatever that amounts to.

I have nothing useful to say about what the right sort of representation is. I will simply say, idealizing in the usual way, that children lexicalize *concepts*. As this suggests, lexicalization is something that happens to concepts.¹² I follow Fodor (197x, 198x, 199x, 200x) and others in taking concepts to be symbols of one or more mental languages. But let me stress 'one *or more*'. I do not assume that the concepts a child lexicalizes are, prior to lexicalization, elements of a *single* systematic language of thought. Some of the concepts a child will soon lexicalize may be cognitively isolated from others; see Spelke (200x), Carruthers (200x). Prelexicalized concepts may not yet satisfy Evans' (198x) Generality Constraint, discussed in chapter three. I assume that prelexicalized concepts can be combined to form thoughts, at least in the sense that each concept belongs to a family of concepts that can be combined. But let's not prejudge whether the concepts of a child acquiring a Language already form one big unified family, or multiple families that may sometimes overlap and communicate via intermediaries. (The contrast matters if *truth*-evaluable thoughts must be composed of concepts that satisfy a Generality Constraint.)

In any case, concepts surely compose in ways that go beyond predicate-conjunction. Humans, including small ones, have concepts of diverse types. Just focusing on formal distinctions, we have singular concepts and predicative concepts of several kinds: lots of monadic and dyadic concepts, like DOG() and BIGGER(,); at least many triadic concepts like BETWEEN(, ,); and at least some tetradic concepts like SELL(, , ,). I bet you have the triadic concept of giving, and if you can't distinguish giving from selling, let's meet. Birds, bees, and even ants can navigate by exploiting information that can be described in terms of how the sun (or a constellation of stars) is positioned in the sky; see Gallistel (200x). And it is, to understate, hard to see how animals could exploit such information without polyadic mental representations. A human child, presumably more sophisticated than an ant, has rich representational capacities.

Indeed, as I'll argue in chapter two, acquiring a Language requires second-order concepts. A child who acquires a word like 'everybody' or 'every' needs a concept that is second order in the sense of combining with one or two *predicative* concepts to form a thought. At least for illustration, let's say that while the first order monadic concept DOG() can combine with the singular concept FIDO to form the true thought DOG(FIDO), the concept EVERYTHING[] can combine with DOG() to form the false thought EVERYTHING[DOG()]; and while the dyadic concept CHASE(,) can combine with FIDO and GARFIELD to form CHASE(FIDO, GARFIELD), EVERY[,] can combine with DOG() and BROWN() to form EVERY[DOG(), BROWN()]. But again, I do not assume that unlexicalized quantificational concepts can combine systematically with *all* the first-order monadic concepts a child may have. Perhaps prelexical EVERY[,] is rooted in a faculty that makes it possible to represent inclusion relations between (extensions of) certain monadic concepts, but not others, with the result that EVERY[DOG(), BROWN()] is thinkable prelexically while EVERY[DOG-THAT-ATE-A-BONE(), CHASED-MOST-CATS()] is not.

Simplifying a bit, we can think of lexicalization as a process of associating concepts with words. From a child's perspective, the task is to connect each lexical item with one or more concepts, each of which has some adicity.¹³ But lexical items can only combine in ways

permitted by the available modes of combination. So each lexical item, of some permissible type $\langle\alpha\rangle$, must fit the associated concept(s) in a way that will make sense given how the lexical item can combine with others, *and* how the concepts can combine with others. This effectively requires that lexicalization be governed by a principle relating word types to conceptual types. Put another way, the space of word-to-concept projections that a child *can* acquire is constrained. Otherwise, there would be too many options for lexicalization, too many of which would not lead to the kind of systematicity in Language/thought relations that children regularly achieve.

It is, however, not obvious what the relevant constraints are. There are many grammatical distinctions to which a lexicalizer might attend. But suppose children hit on the following list of candidate lexical items: proper names, demonstratives, and indexicals; common nouns, adjectives, adverbs, and intransitive verbs; transitive verbs and prepositions; ditransitive verbs; negation, connectives, and quantifiers. It is not obvious how to relate these grammatical categories to the inventory of concepts—monadic and dyadic (first and second order), triadic, tetradic, and singular—for at least two reasons. First, some of the grammatical categories may not be *lexical*, in the sense of classifying *unstructured* expressions, as opposed to complex expressions with subtle structure. And second, since we don't yet know what the semantic composition operations are, we don't know what the semantic (as opposed to conceptual) types are.¹⁴ But we can consider some hypotheses about how meanings “fit” words to concepts.

Here is one proposal, rejected in chapter three, leaving adverbs and tetradic concepts aside until then. Second-order concepts are indicated with **SPECIAL** font.

<i>Grammatical Type</i>	<i>Semantic Type</i>	<i>Conceptual Type</i>
proper names, demonstratives, indexicals	$\langle e \rangle$	singular
common nouns, adjectives, intransitive verbs	$\langle e, t \rangle$	monadic
transitive verbs, prepositions	$\langle e, \langle e, t \rangle \rangle$	dyadic
ditransitive verbs	$\langle e, \langle e, \langle e, t \rangle \rangle \rangle$	triadic
negation	$\langle t, t \rangle$	monadic
connectives ('and', 'or', etc.)	$\langle t, \langle t, t \rangle \rangle$	dyadic
quantifiers (e.g., 'everyone')	$\langle \langle e, t \rangle, t \rangle$	MONADIC
determiners (e.g., 'every')	$\langle \langle e, t \rangle, \langle \langle e, t \rangle, t \rangle \rangle$	DYADIC

On this view, the semantic type of a lexical item matches the adicity of the concept lexicalized, at least as an idealization. The idea is that singular concepts and expressions of type $\langle e \rangle$ have adicity 0, and likewise for complete thoughts and expressions of type $\langle t \rangle$; monadic concepts and expressions of type $\langle e, t \rangle$ have adicity 1, and likewise for **monadic** concepts and expressions of type $\langle t, t \rangle$; etc. This “adicity matching” proposal is invited though not required by the FRET conception of permissible semantic types. *If* all these types are available to lexicalizers, then it would seem natural to pair words with concepts more or less as depicted above. Indeed, one might say that expressions of Languages have the various semantic types they do *because* they are associated with concepts that have the various adicities they do.

On the other hand, it is not clear how to determine the adicity of a concept indicated with a word. I take no stand on particular cases, apart from a few (relatively uncontroversial) logical concepts. My claim is that we have concepts of diverse adicities, including many singular and

polyadic concepts. But it doesn't matter, for my purposes, *how* polyadic the concept *SELL* is. It can be 5-place, with an event variable already there—or 8-place, with “slots” for the date, weather, and exchange rate in pounds sterling. By contrast, it is important for an adicity matching hypothesis that transitive verbs indicate concepts with adicity 2: not 3 or 4; 5 is right out. But how does one tell if the alleged confirming instances of such a view are really confirming? Does ‘triangle’ lexicalize *one* concept? And if so, is the relevant prelexical concept monadic or triadic? Competent speakers may have a concept with which they can think about trilateral figures that have vertices, *and* a concept with which they can think about three points that can be (and perhaps are) connected by lines. To evaluate an adicity-matching hypothesis, we need some principled way of saying which concept counts as the one lexicalized by the common noun.

I return to this point, and the related question of whether unconnected dots can (together) constitute triangles, squares, and circles that count as satisfiers of singular nouns like ‘triangle’, ‘square’, and ‘circle’. But here, let me offer another illustration of the general point that one cannot *intuit* the adicity of the “primary” concept indicated with a word. It can seem obvious that ‘eat’ indicates a dyadic concept. And we presumably do have a concept with which we can think about the relation an eater bears to the eaten. But we may also have a monadic concept, akin to *FUELED-UP*, with which we can think about those who eat. And even if the monadic concept is somehow derivative, that doesn't yet tell us which concept(s) are lexicalized with ‘eat’. Note that ‘mortal’—as in the classic ‘Socrates is mortal’—may be somehow derivative on a polyadic concept that relates entities to events of death; in which case, advocates of adicity-matching hypotheses may not want to insist that “derivative” concepts are irrelevant.

Moreover, the verb ‘eat’ can appear in nontransitive sentences like ‘John ate’; and contrary to initial expectations, the nontransitive sentence does *not* mean that there is something John ate. As Chomsky (196x) observed, if John ate a tack, there is something John ate—and so John ate something, on one construal of ‘ate something’. But if tacks are not edible for John, despite being things he can (unwisely) ingest, then it isn't true that John ate. In that case, John didn't dine, or even have a snack. Although if John is a robot for whom tacks are nutritious, then if John ate a tack, he did eat. This doesn't show that the basic concept in this cognitive vicinity *isn't* dyadic. But it is hardly obvious that ‘eat’ is of type $\langle e, \langle e, t \rangle \rangle$, and that the verb is an expression of this semantic type because the verb lexicalizes a dyadic concept.

Here is a different hypothesis, suggested by the PC conception of how words are related to semantic types, leaving prepositions aside for now. This proposal is defended in chapter three.

<i>Grammatical Type</i>	<i>Semantic Type</i>	<i>Conceptual Type</i>
proper names, demonstratives, indexicals, common nouns, adjectives, adverbs, verbs (intransitive, transitive, ditransitive), negation, connectives, quantifiers, and determiners	$\langle e, t \rangle$	singular, monadic, dyadic, triadic, and tetradic, monadic, dyadic, MONADIC and DYADIC

On this view, the semantic type of a lexical item is $\langle e, t \rangle$, *whatever* adicity the lexicalized concept has; where ‘e’ ranges over entities/events, truth values, and ordered pairs thereof. (The FRET conception ends up requiring a similar domain.) From this perspective, adicity mismatches

are rampant: singular concepts are lexicalized via expressions with adicity 1; dyadic and supradyadic concepts are lexicalized via expressions with adicity 1; even **DYADIC** concepts like **EVERY**[,] are lexicalized via expressions with adicity 1. And as we'll see, the requisite abstractions are not those one would expect in a theory of *truth*.

If $\langle e, t \rangle$ is the only type available to lexicalizers, then lexicalization must be a process in which nonmonadic concepts are somehow paired with monadic analogs, presumably with some nonmonadic residue. There is evidence in favor of this view; see chapter three. One can adopt a more traditional conception of semantic types, without the idea of adicity-matching, and still say that lexicalization is a process in which concepts are forced into certain molds. For example, perhaps any concept lexically associated with 'eat' or 'sell' gets formatted as a concept of type $\langle e, \langle e, t \rangle \rangle$, while any concept lexically associated with 'mortal' or 'triangle' gets formatted as a concept of type $\langle e, t \rangle$. But, I'll argue, this view predicts symptoms that are not there.

Still, even if the massively monadic proposal is motivated, one wants to see a specific formalism, if only to get a sense of how lexicalization and composition *could* work in the manner suggested. Section two included a suggestion about negation and existential closure. Section four is a sketch, filled out in chapter two, of how other conceptual types can be accommodated. As before, the technicalities can be skimmed by readers not yet concerned with about the details.

4. Proposal: Meanings Conjoin, with a Little Help from Some Friends

Begin with the familiar assumption that nouns and adjectives are of type $\langle e, t \rangle$. But suppose that a *name* is a complex expression whose constituents include a lexical noun (of a proper sort) and an index. For simplicity, let's say that each occurrence of the name 'Tyler' has the grammatical form $iD^{\wedge}Tyler_N$; where iD is an indexed covert determiner, and the noun is satisfied by individuals properly called (with the sound of) 'Tyler'. In truth-theoretic terms, one can hypothesize that relative to any assignment **A** of values to indices, an entity e satisfies iD iff $e = \mathbf{A}(i)$; where $\mathbf{A}(i)$ is the entity assigned to index i by **A**. If grammatical combination ($^{\wedge}$) signifies predicate-conjunction, it follows that e satisfies $iD^{\wedge}Tyler_N$ relative to **A** iff $e = \mathbf{A}(i)$ & e is properly called 'Tyler'. Given this idealization, $iD^{\wedge}Tyler_N$ is satisfied by $\mathbf{A}(i)$ or nothing, depending on whether or not $\mathbf{A}(i)$ is a Tyler.

One can think of 'D' as a device that converts an index, presumably of type $\langle e \rangle$, into a formally distinct predicate of type $\langle e, t \rangle$ with the same intuitive content. While an index cannot itself be conjoined with an expression of type $\langle e, t \rangle$, the abstract predicate iD can be.¹⁵ On this view, a singular concept used to think about a particular Tyler is not lexicalized with any particular word. Rather, the concept is expressed—in a context dependent way—with an index and a noun that can be used to talk about other Tylers. I think this cognitive "detour" is required because in Languages, there can be no lexical expressions of type $\langle e \rangle$. In one sense, the index is an atomic element of the name. But an index cannot be a free-standing expression. It must be plugged into a phrase via some "adapter" that allows for a smooth flow of conjoinability. And in any case, indices do not combine recursively with other atomic elements in the manner of open-class lexical items.

This model can be extended to deictic and indexical pronouns, by supposing that nouns can be common, proper, demonstrative, or, indexing. This is independently plausible if 'this' and 'that' are not strictly synonymous. Perhaps the constituent nouns in $iD^{\wedge}this_N$ and $iD^{\wedge}that_N$ impose different spatial proximity requirements, while the nouns in $iD^{\wedge}he_N$ and $iD^{\wedge}she_N$ are

like ‘male’ and ‘female’. Expressions like ‘me’ and ‘now’ may have grammatical forms like ‘sD’ and ‘tD’, where ‘s’ and ‘t’ are special indices associated with grammatically marked dimensions of communicative situations, as in Kaplan (198x); see chapter one.¹⁶

Intransitive verbs can be treated as lexical expressions of type $\langle e, t \rangle$. But instead of saying that a verb like ‘sang’ or ‘died’ is satisfied by entities who sang or died, I follow recent trends and say that such verbs are satisfied by events like singings and deaths. (Here, I speak in truth-theoretic terms for simplicity, instead of saying that verbs are instructions to construct monadic concepts.) The entity/event distinction makes no difference to semantic composition on my view; see chapter two. But we can mark this *conceptual* distinction by saying that ‘dog’ and ‘singer’ are of type $\langle e, t \rangle$, while ‘performance’ and ‘bark’ are of type $\langle \underline{e}, t \rangle$.¹⁷ The idea is that ‘brown’, ‘loudly’, ‘brown dog’, and ‘sang loudly’ are expressions of the same semantic type.

So far, so good. But a verb can be modified by a prepositional phrase like ‘in Boston’. And inferences from ‘sang in Boston’ to ‘sang’ are as good as inferences from ‘sang loudly’ to ‘sang’, which are as good as inferences from ‘red ipod’ to ‘ipod’. Similar remarks apply to ‘sang loudly in Boston yesterday’, strongly suggesting that ‘in Boston’ is understood as a conjoinable predicate of type $\langle \underline{e}, t \rangle$. (These are among the facts suggesting that grammatical combination at least often signifies predicate-conjunction, as opposed to function-application.) In the noun phrase ‘doctor in Boston’, the prepositional phrase seems to be of type $\langle e, t \rangle$; but all the more reason, I’ll argue, for not putting any theoretical weight on the e/\underline{e} distinction. Still, there is a potential snag. If we want ‘sang in Boston’ to be satisfied by events of singing that occurred *in* Boston, as opposed to things that are both events of singing and Boston, then we need to get a *relation* into the picture—a relation that can be exhibited by a city and an event. But this isn’t much of a snag, since introducing such relations is plausibly the work of prepositions.

Formally, it is easy to introduce a relation without violating the spirit of Conjunctivism: relative to any situation in which ‘in’ is used to indicate a relation ${}^{\text{in}}R$, ‘in Boston’ is satisfied by an entity or event e iff $\exists e' [{}^{\text{in}}R(e, e') \ \& \ \text{Boston}(e')]$. Though if you want to know what satisfies ‘in’, or what function from contexts to relations ‘in’ itself indicates, you’ll have to ask a Truth Conditional Semanticist. Given phrases like ‘in trouble’, ‘in haste’, ‘in a huff’, and ‘in a second’, I find it hard to pretend—even for simplicity of initial exposition—that ‘in’ is satisfied by pairs $\langle e, e' \rangle$ such that $\text{In}(e, e')$. What would ‘In(e, e’)’ mean in our theoretical metalanguage? Would ‘e’ range over things like cities, trouble, haste, huffs and seconds? My own view is that ‘in’ is an instruction to select an appropriate relational concept, on an occasion of use, and that the dramatic polysemy of prepositions reflects the imprecision of such instructions: when relations are introduced, speakers have more work to do, since Languages do less. But in any case, ‘in’ is not a monadic predicate that *conjoins* with ‘Boston’. In traditional terms, ‘in’ is more like a syncategorematic element, though of a special sort: it takes a predicate like ‘Boston’ and returns a predicate of type $\langle e/\underline{e}, t \rangle$ that can be coherently conjoined with a predicate like ‘doctor’ or ‘sang’. (Chapter two addresses plural predicates.)

The idea is that prepositions, which are also “closed class” words, are like the posited determiner ‘D’ in being adapters that preserve the flow of conjoinability.¹⁸ The proposal is *not* that thought is essentially monadic, or that meanings never reflect relational concepts. But one can defend the “neo-medieval” view that in Languages, relationality is introduced only through a small number of devices, despite the more pervasive relationality of thought. Prepositions are

paradigm examples of such devices. Though if verbs are semantically monadic predicates of events, and there is a genuine argument/adjunct distinction—in that arguments are not introduced by covert prepositions—then certain *grammatical relations* between verbs and their arguments are also devices for introducing *thematic relations* exhibited by events and their participants.

Grammatical subjects and objects are understood as if introduced by prepositions—i.e., as if the grammatical relations *being a subject of* and *being an object of* are vehicles for introducing thematic relations. The near synonymy of ‘Socrates sang’ and ‘There was a singing by Socrates’, with ‘by’ is used to signify agency, suggests that being a subject of a verb like ‘sang’ is like being introduced by an agent-indicating preposition. As discussed in chapter two, not all subjects are represented as agents; subjecthood does not itself indicate agency. But one can hypothesize that the grammatical relation between a verb and its subject introduces an abstract formal relation, $\text{External}(\underline{e}, e)$, such that: if \underline{e} satisfies a verb like ‘sang’, then $\text{External}(\underline{e}, e) \equiv \text{Agent}(\underline{e}, e)$; if \underline{e} satisfies a verb like ‘saw’, then $\text{External}(\underline{e}, e) \equiv \text{Experiencer}(\underline{e}, e)$; and so on, for the appropriate number of specific thematic relations indicated with subjecthood. Likewise, the near synonymy of ‘Socrates ate Boston’ and ‘There was an eating of Boston by Socrates’ suggests that being an object of a verb like ‘eat’ is a lot like being introduced by an patient-indicating preposition; although Boston is not intuitively represented as affected in ‘Socrates left/saw/recalled Boston’. So perhaps grammatical objecthood is semantically associated with an abstract formal relation, $\text{Internal}(\underline{e}, e)$, such that: if \underline{e} satisfies a verb like ‘ate’, then $\text{Internal}(\underline{e}, e) \equiv \text{Patient}(\underline{e}, e)$; but the internal participants of other events need not be Patients. In this sense, the relevant grammatical relations are polysemous, somewhat like prepositions.¹⁹

On this view, a supramonadic concept, used to think relational thoughts, is not lexicalized with a word of type $\langle e, \langle e, \mathbf{t} \rangle \rangle$ or higher. Instead, relational concepts are expressed via lexical monadic concepts and a few “all purpose” thematic relations, indicated with prepositions and certain grammatical relations. Expressions of thematic relations are situation dependent, like the expression of singular concepts via indices and monadic predicates. In chapter three, I’ll argue that this is a source of support for Conjunctivism: nonmonadic concepts are expressed the way they are in Languages, as opposed to other more straightforward ways, because (open class) lexical expressions must be of type $\langle e, \mathbf{t} \rangle$.

This also leads to an attractive semantic unification of arguments and adjuncts. The logical forms of ‘Socrates sang in Boston’ and ‘Socrates ate Boston’ are roughly as shown below.

$$\begin{aligned} & \exists[\exists e\{\text{Socrates}(e) \ \& \ \text{External}(\underline{e}, e)\} \cdot \text{PastSing}(\underline{e}) \cdot \exists e\{\text{Boston}(e) \ \& \ \text{inR}(\underline{e}, e)\}] \\ & \exists[\exists e\{\text{Socrates}(e) \ \& \ \text{External}(\underline{e}, e)\} \cdot \text{PastEat}(\underline{e}) \cdot \exists e\{\text{Boston}(e) \ \& \ \text{Internal}(\underline{e}, e)\}] \end{aligned}$$

As this notation suggests, ‘ $\exists e\{\text{Socrates}(e) \ \& \ \text{External}(\underline{e}, e)\}$ ’ is a complex monadic predicate that is conjoinable with ‘ $\text{PastSing}(\underline{e})$ ’ but formed via resources that go beyond predicate-conjunction. The symbol ‘ \cdot ’, which can only conjoin monadic predicates, is more restricted than ‘ $\&$ ’. So ‘ \exists ’, as opposed to ‘ $\exists e$ ’, always binds the *only* free variable in the relevant open sentence. Since the subexpression ‘ $\text{External}(\underline{e}, e)$ ’ lies within the scope of ‘ $\exists e$ ’, only the first variable position in ‘ $\text{External}(\underline{e}, e)$ ’ remains variable so far as ‘ \cdot ’ and ‘ \exists ’ are concerned.

In this sense, ‘ \exists ’ and ‘ \cdot ’ belong to a truly monadic/medieval system. Though a monadic predicate can be created “off line” by partially saturating a dyadic predicate like ‘ $\text{External}(\underline{e}, e)$ ’, using a singular term like ‘ δ ’ to create a “desaturated” predicate like ‘ $e = \delta$ ’.²⁰ So the logical forms above can be rewritten as below, with the “gap” indicating the variable bound by ‘ \exists ’.

$$\begin{aligned} & \exists[\exists e\{\text{Socrates}(e) \ \& \ \text{External}(_, e)\} \cdot \text{PastSing}(_) \cdot \exists e\{\text{Boston}(e) \ \& \ \text{inR}(_, e)\}] \\ & \exists[\exists e\{\text{Socrates}(e) \ \& \ \text{External}(_, e)\} \cdot \text{PastEat}(_) \cdot \exists e\{\text{Boston}(e) \ \& \ \text{Internal}(_, e)\}] \end{aligned}$$

This in turn suggests a very snug fit between these logical forms and the corresponding grammatical forms, represented below with *i* and *j* as indices.

$$\begin{aligned} & [\ iD \ \text{Socrates}_{N,D}]_D \quad \quad \quad \wedge \ [\text{sang}_V \quad \quad \quad \wedge \quad [\text{in}_P \ [jD \ \text{Boston}_{N,D}]_D]_P]_V \\ & [\ iD \ \text{Socrates}_{N,D}]_D \quad \quad \quad \wedge \ [\text{ate}_V \quad \quad \quad \wedge \quad [jD \ \text{Boston}_{N,D}]_D]_V \end{aligned}$$

One can view ‘ \exists ’ as binding the variable corresponding to the “spine” of the phrase-marker/tree that emerges as expressions are grammatically combined and semantically conjoined.

Given the near synonymy of ‘Socrates died’ and ‘There was a death of Socrates’, and given that deaths seem unlike actions done by agents, this account predicts that ‘Socrates’ is really the grammatical object (or internal argument) of ‘died’. But that is independently plausible; see Belletti (198x), Rizzi (198x), Williams (199x). In which case, the logical form of ‘Socrates died in Boston’ is roughly as shown below.

$$\exists[\exists e\{\text{Socrates}(e) \ \& \ \text{Internal}(_, e)\} \cdot \text{PastDie}(_) \cdot \exists e\{\text{Boston}(e) \ \& \ \text{inR}(_, e)\}]$$

5. Composition is Implemented, Truth is Complicated

Even if Conjunctivism is empirically plausible, questions remain. *Why* are words related to concepts in this way? Why erase adicity distinctions for purposes of semantic composition? One reason, hinted at above, is that a single word may be linked to concepts of varying adicity. Perhaps monadicizing massively is the most economical way of stably associating each word with a semantic type. But why not have one word per concept with adicity-matching as the norm?

Semantic composition is implemented. That is obvious but important. No operation that fails to be implemented by human children, in the understanding of Languages, is an operation of semantic composition. And that may leave few candidates. *If* children have a fully systematic language of thought *prior* to lexicalization, then perhaps several potential modes of semantic composition are already implemented and available for recruitment in Languages. That is one view, invited by the FRET conception of permissible semantic types and an “adicity matching” conception of lexicalization. But even setting aside facts which suggest that adicity mismatches and words of type $\langle e, \mathfrak{t} \rangle$ are ubiquitous, one might wonder if lexicalization is *driven* by powerful and general operations that make it possible to lexicalize any concept (of any permissible *conceptual* type) with a grammatically atomic expression of the corresponding adicity.

If children acquiring a Language *already* have a capacity to form complex concepts systematically—in that for any concepts *C* and *C'*, children can form a (nonpropositional) concept with *C* and *C'* as constituents—perhaps one can explain the source(s) of semantic compositionality in these terms; see Fodor (197x), though cf. chapter one. But then the posited capacity would have to be explained in terms of implementable modes of conceptual combination. And repeating an earlier question, if children enjoy a systematic language of thought *prior* to lexicalizing their concepts, then one wants to know how nature came to implement such an impressive (domain general) Mentalese, independent of Languages.

An alternative hypothesis is that in Languages, combining expressions—concatenating them in a way that creates certain grammatical relations between constituents—signifies predicate-conjunction because that is the only implementable mode of semantic composition, or at least the only one that human biology has yet come up with. Some evolutionarily recent change in primate biology led to a capacity to generate labeled phrase markers of the sort

characterized by transformational grammars.²³ But this change may not have introduced any novel semantic operations. One can imagine a hominid ancestor with a new talent for generating syntactic structures, which presented a cognitive challenge for an otherwise nonhuman primate mind: who ordered *those*, and how are they to be interpreted, now that they're here?

One can speculate that function-application was already available, independently implemented in such a mind. But then one wonders why such a mind needed Languages, whose *syntax* seems to differ importantly from the syntax of thought; see chapter one for review of Chomskyan illustrations. Why not cut straight to talking, by associating concepts with sounds, and letting the already systematic syntax of thought do double duty as the syntax of a language with overt signals? Perhaps this question can be answered. Maybe the need to *linearize* speech made it impossible to pronounce thoughts “directly.” Or maybe things started with everyone pronouncing thoughts directly, understanding one another just fine, but then there was pressure of some kind to segregate humans into distinct Linguistic communities.²⁴ And maybe not.

Another possible story is that our hominid ancestors, faced with the task of interpreting syntactic structures, did what they could: they interpreted concatenation as predicate-conjunction, which is *almost* like conjunction of thoughts; and they dealt with lexicalization on a case by case basis, using a small stock of devices for introducing relationality as needed. If it helps, imagine a spoken language Andish+ with a few expressions like ‘ext’, ‘int’ and ‘loc’, which can appear in expressions like [extC J], [intA F], and [A locC], with roughly the following meanings: a crocodile jumped; an armadillo fell; and an armadillo is in a crocodile.²⁵

Readers with impressive attention spans may recall that my initial remarks about compositionality—which led to discussions of Andish, how and what meanings fit, analogies between sentences and discourses, lexicalization, some details, abstraction, adjunction, thematic roles, hominid ancestors, and now Andish+—were supposed to be in the service of explaining the claim that meanings are concept construction instructions, and not truth-theoretic properties of expressions. On the negative front, I have suggested that meanings combine in Conjunctivist fashion, while hinting that the best candidates for truth-theoretic properties of Linguistic expressions do not. Putting this another way: if meanings are implemented in Conjunctivist fashion, but truth conditions are generally too complicated to be so implemented, that tells against Truth Conditional Semantics. It also suggests an alternative.

Suppose that Harold discovers yet another language, PolyAndish+, with lexical items that he translates into English as ‘line’ and ‘book’. Speakers sometimes use the first word to speak of spatiotemporally located queues, or things resembling queues, that fail to connect the shortest distance between two points. Other times, they use this word—and by all the usual tests, it is the same word—to talk about abstracta with no width. They also talk about lines of thought, lines of inquiry, blood lines, and even two different “kinds” of blood lines that forensic scientists can track. The second word, ‘W’, can be used to speak of spatiotemporally located tomes, or more abstract “things written” that have various instantiations, or even some blank pages glued together with a cover. The same lexical item can be also be used in complex expressions like [extC [W intA]], which can be used to say either that a crocodile charged an armadillo with a crime, or that a crocodile hired an armadillo to perform at a venue of entertainment.

Harold conjectures that speakers of polyAndish+ use ‘W’ to indicate various concepts, making this word nonhomophonous but polysemous. He has also been reading Marr (198x), and learning that implementation is sometimes achieved by performing computations in stages.

Maybe when an utterance of [extC [W intA]] is interpreted, interpretation *starts* with a recognition that the expression used has a logical form along the following lines:
 $\exists[\exists e\{\text{Crocodile}(e) \ \& \ \text{External}(\ , \ e)\} \ \bullet \ \text{Book}(\) \ \bullet \ \exists e\{\text{Armadillo}(e) \ \& \ \text{Internal}(\ , \ e)\}].$
But perhaps this is just a “primal sketch” of a thought, and not yet a truth-evaluable representation. After all, ‘External’ and ‘Internal’ signify mere formal relations. It also seems implausible, *prima facie*, that the middle conjunct is a predicate *satisfied* by tomes and propositional contents (expressions of which can fill a tome) and other “things” in which crocodiles and armadillos can participate. Maybe instead, speakers hear ‘W’ in [extC [W intA]] as an instruction to fetch an eventish concept that is suitably related to a concept with which one can think about tomes or propositions.

Harold pauses, dimly sensing that his career as an introductory trope is coming to an end,

Of course, one could *stipulate* a truth-theoretic interpretation for the logical form, in terms of crocodiles and armadillos (or prime numbers and toaster ovens). But that is a banality, which I will not confuse with empirical hypotheses in the neighborhood.

and then returns to his main thought.

If speakers “infuse” a logical form with specific and sensible content, the underlying process may be complex and situation sensitive, in a way that precludes execution by any simple compositional system—much less the one that generates infusable forms. There may be a gap between logical forms and truth conditions that is fillable, by the minds of those who speak polyAndish+, but not fillable by the cognitive subsystem that systematically generates primal sketches of thoughts that are formally Andish.

Harold wonders why a mind would be organized in such fashion, but reflects on his ignorance.

For all I know, initial sketches may be required to get a process of interpretation *started* in a constrained way that has any real chance of *ending* with a truth-evaluable representation that two speakers (or one speaker at two times) could focus on for purposes of evaluation, tweaking, reevaluation, retweaking, etc. Or maybe the sketches make polyAndish+ *acquirable* by the next generation. Perhaps a more “direct” system that mapped grammatical forms straight to representations of truth conditions would be unacquirable, or unimplementable, or both.

6. Proposal: Meanings are Conjunctive *Begriffsplans*

Let’s go back, now, to the discussion of ‘S•M’ in section two.

Instead of construing ‘M’ as a predicate *satisfied* by mortal beings, think of ‘M’ as an *instruction* that can be executed: *fetch* a concept that is both lexically associated with ‘mortal’ and of the right *form* for an operation of predicate-conjunction. Or more briefly: fetch a monadic ‘mortal’-concept. If the prelexical concept MORTAL is itself monadic, and conjoinable with other such concepts, then that concept can be fetched. If not, then lexicalization must somehow find or create a monadic alternative to the prelexical concept. And given that a mortal human can receive a mortal wound, someone who understands ‘mortal’ may have more than one monadic ‘mortal’-concept (each of which is somehow related to a concept of death), even if one these concepts is more basic than the others.

The prelexical concept *SOCRATES* may be singular. But it can be lexically linked to an analytically related monadic concept. So think of ‘S’, not as a predicate uniquely satisfied by a particular individual, but rather as an instruction to fetch a monadic ‘Socrates’-concept. (Remember that such a concept may have a deictic constituent, and that a thinker may have more than one such concept if she knows more than one Socrates.) Likewise, instead of construing ‘•’ as a sign for the operation of predicate-conjunction, think of ‘•’ as an instruction to conjoin concepts that have been fetched or constructed. The idea is that ‘S•M’ is an instruction to conjoin a monadic ‘Socrates’-concept with a monadic ‘mortal’-concept, and ‘ \exists [S•M]’ is the following instruction: fetch a monadic ‘Socrates’-concept and a monadic ‘mortal’-concept, conjoin them, and ‘ \exists ’-close the result. In which case, if ‘ \exists [S•M]’ is the logical form of ‘Socrates is mortal’, then to understand the sentence is to perceive it as a certain concept construction instruction.

Instructions can be satisfied. But this truism can confuse, given the history of semantics. Expressions of Frege’s (1879) *Begriffsschrift* can be viewed as unequivocal instructions, or algorithms, or blueprints for building concepts with truth-theoretic properties; see Horty (2007). And with regard to an invented formal language whose sentences are mirrors of truth-evaluable thoughts, saying that expressions are concept construction instructions—or *Begriffsplans*—may be equivalent to saying that expressions have satisfaction conditions in Tarski’s (193x) sense.²⁶ But the Languages that human children acquire may differ, in that expressions of these languages are “human *Begriffsplans*” that do *not* have satisfaction conditions in Tarski’s sense. A single word may be polysemously connected to a family of concepts, each of which is vague and context-sensitive in ways not indexed by the word itself. But a word, with its lexical meaning, can still be an executable (and in that sense satisfiable) instruction to fetch a monadic concept.

If this view is correct, then lexicalizing a nonmonadic concept *C* ensures at least one monadic analog of *C* that can be fetched and predicate-conjoined with other such concepts.²⁷ If the prelexical concept *EAT* is polyadic, it gets paired with a monadic concept of events. Not all concepts are eventish in this sense. But even for determiners like ‘every’, it is relatively easy to “monadicize” the lexicalized concepts. And while there is a sense in which lexicalizing increases the number of concepts a child has, there is another sense in which lexicalization adds nothing new. As discussed in chapter three, drawing on Horty, Frege’s notion of analysis is useful here. When lexicalizing a concept triggers creation of a monadic analog, the analog need not be logically independent of the initial concept.

On the contrary, we can think of the monadic analog as the result of applying a certain operation (of abstraction/analysis) to the initial concept; where this operation pairs concepts of a given type—say, singular concepts, or dyadic action concepts—with formally distinct concepts that have the same intuitive content. For example, given the concept *EAT*(*X*, *Y*), the concept *EAT*(*E*) might be introduced as follows: *EAT*(*X*, *Y*) \equiv *AGENT*(*E*, *X*) & *EAT*(*E*) & *PATIENT*(*E*, *Y*). Given a prelexical concept *C*, the corresponding monadic concept might also be a “layered” concept with *C* as the “core” surrounded by a Linguistic surface that fits the demands of predicate-conjunction. Arguing for this will require the whole book, and not just chapter three, where the idea is front and center. But for now, the aim is to make the proposal tolerably clear. So imagine a process that begins with a lexical sound *s* being paired with a prelexical concept *C*.

The sound paired with *C* *across* episodes of speech will vary dramatically, unless we employ a very abstract notion of sound. Not that there’s anything wrong with that. But for these

illustrative purposes, suppose that C gets paired with an *instruction to generate* sounds like s. Call this instruction ‘ $\pi:s$ ’, and assume that many different sounds can be *classified as* sounds that were produced in accord with $\pi:s$.²⁸ Now suppose that some process of abstraction relates C to a formally distinct concept MAC, a monadic analog of C. We can represent the emerging lexical item with the set $\{\pi:s, \{C, MAC\}\}$. If this proto-word is subsequently associated with another prelexical concept C' and monadic analog MAC', the result is a case of homophony or polysemy. I don't pretend to have a good account of the difference. (Maybe polysemy results when C' is in the same “cognitive field” as C; maybe not.) But lots of spoken words are polysemous. And such a word does not merely pair a sound-instruction with a particular concept and monadic analog.

Correlatively, lexical meanings are not concepts. The sound-instruction $\pi:s$ might be paired initially with C, but later with $\{C, MAC\}$, $\{\{C, MAC\}, \{C', MAC'\}\}$, and so on. None of these concepts is the lexical meaning. But suppose that $\pi:s$ is eventually paired with an instruction to fetch *some* monadic concept that has been lexically linked to C—either MAC, or a monadic analog of some concept subsequently tied to C. This existential instruction, call it ‘ $\lambda:C$ ’, is compatible with C being a default choice for fetching. But any concept lexically tied to C is thereby tied to $\pi:s$, and *vice versa*. So the semantic instruction, the meaning paired with $\pi:s$, can also be characterized as follows: fetch a monadic $\pi:s$ -concept. And the resulting word can be identified with $\{\pi:s, \lambda:C\}$; where the lexical meaning, $\lambda:C$, is an instruction to fetch a linked monadic concept. As Chomsky (1995) suggests, lexical items may just *be* pairs of instructions via which the human language faculty “interfaces with” other cognitive systems. Abstracting from connections to perceptual/articulatory systems, which seems appropriate for silent uses of Languages in thought, words may *be* concept construction instructions. In this sense, words may *be* pronounceable meanings. And this may explain why meanings are as composable as words.²⁹

In any case, we can at least imagine a process in which a mere sound-concept pair becomes a word like $\langle \pi:\mathbf{socrates}, \lambda:\text{SOCRATES} \rangle$ or $\langle \pi:\mathbf{mortal}, \lambda:\text{MORTAL} \rangle$, remaining agnostic about the adicity of the prelexical concepts. In the simplest case, the prelexical concept will be a monadic concept that remains the *only* lexicalized concept. And when a concept is lexicalized without polysemy, there may be exactly one way to carry out the semantic instruction. But in cases of polysemy, there will be more than one. Suppose that for competent speakers of English, ‘book’ is polysemously linked to at least two concepts: S-BOOK and P-BOOK, corresponding to the spatiotemporally located books and the more abstract (propositionish) books that get written and are harder to destroy. A child might pair the sound of ‘book’ with S-BOOK and then with P-BOOK, retaining the former as the default concept to be fetched. Suppose this lexicalizer also treats ‘chase’ as an indicator of single concept that is in fact dyadic. The result might be as depicted as below, with ‘()’ indicating the monadic analog concepts.

S-BOOK	P-BOOK	CHASE
S-BOOK()	P-BOOK()	CHASE()
$\langle \pi:\mathbf{book}, \lambda:\text{S-BOOK} \rangle$		$\langle \pi:\mathbf{chase}, \lambda:\text{CHASE} \rangle$

The concept CHASE() applies to events of chasing, while S-BOOK() and P-BOOK() apply respectively to concrete and abstract books, even if S-BOOK and P-BOOK are somehow relational.

If a lexicalizer starts with a monadic concept of eating, acquires a dyadic concept of ingesting, and overhears ‘Book him, Dano’, the results may be more interesting. The speaker might end up having a dyadic concept of eating that is the default concept for ‘eat’; and Linguistic understanding might trigger acquisition of an eventish concept E-BOOK() and a corresponding concept E-BOOK, which may be polyadic.

$\begin{array}{l} \text{EAT}_{\text{DINE}} \quad \text{EAT}_{\text{CONSUME}} \\ \text{EAT}_{\text{DINE}}() \quad \text{EAT}_{\text{CONSUME}} \\ \langle \pi: \mathbf{eat}, \lambda: \text{EAT}_{\text{CONSUME}} \rangle \end{array}$	$\begin{array}{l} \text{INGEST} \\ \text{INGEST}() \\ \langle \pi: \mathbf{ingest}, \lambda: \text{INGEST} \rangle \end{array}$	$\begin{array}{l} \text{S-BOOK} \quad \text{P-BOOK} \quad \text{E-BOOK} \\ \text{S-BOOK}() \quad \text{P-BOOK}() \quad \text{E-BOOK}() \\ \langle \pi: \mathbf{book}, \lambda: \text{S-BOOK} \rangle \end{array}$
---	--	--

But the general point—that words can have a common adicity, and yet serve as devices for indicating concepts with diverse adicities—is independent of any particular speculations about conceptual structure. This is good, given how little we know about which concepts with which adicities are initially lexicalized, and what happens in children after initial lexicalization.

Returning now to ‘Socrates is mortal’, instead of saying that this English sentence has a compositionally determined truth condition— \top iff something satisfies both ‘Socrates’ and ‘mortal’—we can hypothesize that the sentence is an instruction: fetch a monadic ‘Socrates’-concept and a monadic ‘mortal’-concept; predicate-conjoin them, and ‘ \exists ’-close the result. Likewise, the meaning of ‘Socrates ate Boston’ can be analyzed as a complex instruction to construct a complex monadic concept.

$\exists[\exists e\{\text{Socrates}(e) \ \& \ \text{External}(, e)\} \cdot \text{PastEat}() \cdot \exists e\{\text{Boston}(e) \ \& \ \text{Internal}(, e)\}]$

Spelling out this recipe: take a monadic ‘Socrates’-concept, proposition-conjoin it to the second variable position of a suitable ‘External’-concept like AGENT(,), and existentially close that variable position to prevent further modification; add a ‘PastEat’-concept—which can be prepared in advance by predicate-conjoining two concepts of events, given a good semantics of tense—and set the mixture aside; take a monadic ‘Boston’-concept, proposition-conjoin to the second variable position of a suitable ‘Internal’-concept like PATIENT(,), close that position as in step one, and add the result to the mixture; apply ‘ \exists ’ for 50 ms, or until firm.

In certain simple cases, such a recipe might be unequivocal in the following sense: there is exactly one way to carry it out, and the resulting thought has a truth condition. If this were true in general, then one could provide a Truth Conditional Semantics for the *thoughts constructed via* the meanings of symbols like ‘M’ and ‘S’. Indeed, I suspect that some neo-Davidsonians have this sort of view in mind, perhaps allowing for additional modes of semantic composition; see, e.g., Larson and Segal (1995). So my suggestion is not that the idea of “meanings as concept construction instructions” is itself hostile to Davidson’s Conjecture. But I do think that the requisite idealizations are implausible, in ways that tell against Truth Conditional Semantics; see chapters four and five. Here, I offer sentences (1-4) as an Austinian example of the general point.

- | | |
|--------------------------|---|
| (1) France is hexagonal | (3) France is hexagonal, and it is a republic |
| (2) France is a republic | (4) France is a hexagonal republic |

It is tempting to say that ‘France’ is satisfied by an entity, France, that can satisfy both ‘hexagonal’ and ‘republic’. This makes for a simple account of why (1) and (2) can each be used to say something true, and why (3) seems semantically equivalent to the conjunction of (1) and (2). But if the words in (1-3) have the requisite satisfaction conditions, and meanings are truth-

theoretic properties of expressions, it is hard to see why (4) is somehow *weird* in a way that (1-3) are not. For suppose that (3) is true, with ‘it’ as a bound pronoun, iff the satisfier of ‘France’ satisfies both ‘hexagonal’ and ‘republic’. Then (3) would seem to be semantically equivalent to (4), which would presumably be true iff the satisfier of ‘France’ satisfies ‘hexagonal republic’. But (4) differs *somehow* from (3). And I don’t think that (4) is ungrammatical.

My hunch is that (4) induces a “conceptual violation” less dramatic than the one induced by Chomsky’s famous (5), which is arguably just as grammatical as (6).

(5) Colorless green ideas sleep furiously

(6) Brilliant white stars shine beautifully

I return, in chapter three, to so-called nonconjunctive adjectives like ‘big’ and ‘small’. Big ants are smaller than small elephants, and there is indeed something amiss with (7).

(7) France is hexagonal for a republic

But I don’t think this is the *source* of anomaly in (4). Note that (8)

(8) Adam is big, and he is an ant

is not intuitively equivalent to (9), which is better glossed with (10).

(9) Adam is a big ant

(10) Adam is an ant, and a big one

And (11) is at least as weird as (4).

(11) France is a republic, and a hexagonal one

In any case, whatever the correct diagnosis, (4) and (11) feel weird to me in a way that (3) and (10) do not. If others share this intuition, it calls for explanation, since (12) and (13) are fine.

(12) Fido is a brown dog

(13) That is a brown house

My worry is not that (4) is necessarily false, or even false. I don’t think *any* sentence of a Language is itself true or false. But I grant that (4), like many sentences, can be used to say something true.³⁰ Nor is my primary concern that the world fails to include a republic that is hexagonal. As it happens, I am more ontologically conservative than some. I don’t believe there really is anything that is both hexagonal and a republic. But I am more ontologically liberal than some, in believing that there really are republics. So I am inclined to say: republics are not things of the right sort to *be* hexagonal. By itself, that is mere autobiography. But it does suggest that a competent speaker of English can understand (1-2) *without* understanding them as *entailing* (4). By contrast, I feel the pull of saying that (1) and (2) jointly entail (3). And I heartily endorse inferences like the following: France is a country, and France is in Europe, so France is a country in Europe. I just don’t hear the inference from (1) and (2) to (4) in the same way.

That inference sounds, to me, like a joke. But it shouldn’t sound like a joke given the following semantics: (1) and (2) are both true iff the satisfier of ‘France’ satisfies both ‘hexagonal’ and ‘republic’; and ‘hexagonal republic’ is satisfied by an entity *e*, like the satisfier of ‘France’ iff *e* satisfies both ‘hexagonal’ and ‘republic’. In chapters four and five, I consider some other things that a Truth Conditional Semanticist might say. But unsurprisingly, I think a better response is to view ‘hexagonal republic’ as an instruction to construct a monadic concept.

For simplicity, suppose the logical forms of (1) and (2) are as follows, ignoring the final \exists -closures: ‘F•H’ and ‘F•R’. Think of the first as an instruction to predicate-conjoin a ‘France’-concept with a ‘hexagonal’-concept, and the second as an instruction to predicate-conjoin a ‘France’-concept with a ‘republic’-concept. In my view, each of these two instructions can be

executed in a way that results in a complex concept that can apply to something, even if no concept applies to any thing that is both hexagonal and a republic. There's polysemy for you.

There may be *a* 'France'-concept with which we can think about certain hexagonal terrain, and *another* 'France'-concept with which we can think about a country that (now) has a republican form of government. And of course, these concepts may be related, in various ways, despite being distinct. Likewise, there may be *a* 'hexagonal'-concept with which we can think about spatially extended regions, and *another* (less vague) 'hexagonal'-concept with which we can think about certain abstract geometric figures. And of course, these concepts may be related, in various ways, despite being distinct. In which case, there are at least four possible ways of executing the instruction 'F•H'. But three of these ways may lead to the construction of monadic concepts that are somehow defective, despite being formally kosher on their Linguistic surfaces. And for simplicity, suppose there is exactly one 'republic'-concept, which we can use to classify things like the French polis but not the French terrain. Then there are two ways of executing the instruction 'H • R'. But both may lead to the construction of monadic concepts that are somehow defective, despite being formally kosher on their Linguistic surfaces.

7. Recapitulation and Final Advertisements

The previous paragraph may have gone by a little fast for some tastes, even for a first pass.

Recall the suggestion that a prelexical concept C might be linked to a monadicized analog concept, MAC, that is "layered" in the following sense: C itself is the "core" of MAC, whose Linguistic surface fits the demands of predicate-conjunction. That metaphor needs spelling out. But the idea will be that MACs are built to "interface" between a Linguistic system that wants to predicate-conjoin, all day long and without regard to intuitive content, and older cognitive systems that provide the diverse "core concepts" that children lexicalize; where the prelexical concepts are not, without help from their Linguistic relations, fully integrable. Spelke (200x) and Carruthers (200x) provide motivating data and arguments. But this is not to say that the prelexical concepts fail to be *composable* in every sense of 'able'. Leaving this delicate issue for later, it may be that without Languages, prelexical concepts would never get past certain barriers to integration of the sort required by Evans's (1982) Generality Constraint.

Suppose this idea can be made good. The four possible ways of executing 'F • H' correspond to four constructable monadic concepts that might be depicted as below.

FRANCOTERRAIN() • HEXREGION()
FRANCOTERRAIN() • HEXFIGURE()
FRANCOPOLIS() • HEXREGION()
FRANCOPOLIS() • HEXFIGURE()

If only for purposes of illustration, let's make the following assumptions: 'France' can be used (with an index) to indicate either of two *singular* concepts—FRANCOTERRAIN and FRANCOPOLIS—with which we might think about the French terrain or the French polis; these singular concepts have monadic analogs, FRANCOTERRAIN() and FRANCOPOLIS(); and 'hexagonal' can be used, in contexts, to indicate either of two *monadic* concepts, HEXREGION() or HEXFIGURE(), with which we can mentally classify spatial regions or geometric figures. But constructing a concept is one thing. Deploying it in thought is another.

The first of the four concepts indicated above is very closely related to the following thought: HEXREGION(FRENCHTERRAIN). This thought provides a specific, thinkable way for the

following existential thought to be true: $\exists[\text{FRANCOTERRAIN}() \bullet \text{HEXREGION}()]$. By contrast, each of following existential thoughts may be formally fine, and yet trigger an unsuccessful search for a thinkable value of the variable: $\exists[\text{FRANCOTERRAIN}() \bullet \text{HEXFIGURE}()]$; $\exists[\text{FRANCOPOLIS}() \bullet \text{HEXREGION}()]$; $\exists[\text{FRANCOPOLIS}() \bullet \text{HEXFIGURE}()]$. There is, of course, a trivial sense in which any thought you can construct is a thought you can think. Likewise, any monadic concept you can construct is a concept you can think with, and conjoin with other monadic concepts. On my view, that’s what MACs are for: conjoining and thinking with. But precisely because we humans have Languages with systematically composable lexical elements, we can—and I suspect we must—distinguish “mere formal” thinkability from a “more intuitive” notion of thinkability that can be cashed out in terms of whether a conjunction of “layered” monadic concepts is thinkable at the “core” level.

With that vague and hardly new idea in mind, let’s return to (4).

(4) France is a hexagonal republic

If ‘republic’ is lexically paired with just one monadic concept, there are still two ways of executing the instruction ‘hexagonal republic’, corresponding to two constructable concepts.

$\text{HEXREGION}() \bullet \text{REPUBLIC}()$

$\text{HEXFIGURE}() \bullet \text{REPUBLIC}()$

But with regarding to finding an “intuitively thinkable” value of the variable, these concepts may be no better than $\text{FRANCOPOLIS}() \bullet \text{HEXREGION}()$ and $\text{FRANCOPOLIS}() \bullet \text{HEXFIGURE}()$. In which case, there may be no way of executing instruction (4) that leads to a thought that is more than “formally thinkable”. By contrast, the following thought might be easily and intuitively thinkable: $\text{REPUBLIC}(\text{FRANCOPOLIS})$. And this thought provides a specific way for the following existential thought to be true: $\exists[\text{FRANCOPOLIS}() \bullet \text{REPUBLIC}()]$. If the pronoun in (3)

(3) France is a hexagonal, and it is a republic

is understood as an instruction to fetch a ‘France’-concept, not necessarily the one fetched via ‘France’, then instruction (3) can also lead to an intuitively thinkable thought—viz., the conjunction of the “easy” thoughts corresponding to (1) and (2). Nonetheless, (4) can be a weird instruction: one that cannot be carried out in a way that does not lead to a cognitive dead end.

Let me quickly enter three caveats. First, the point is not that some instructions cannot be carried out. (Square that circle!) The point is that understanding and executing an instruction, as well as possible, is no guarantee of a good result. (When I said “Wear a shirt and a tie,” I assumed you had a shirt and matching tie.) Second, I don’t want to overstress any example. Perhaps one can find or make an intuitive concept of a hexagonal republic—a concept with that very content. Perhaps we already have one, and I’m just confused. Maybe we also have a concept of a colorless green idea. But do we have, in the relevant sense, a *concept* of a colorless green idea who wore a hexagonal republic that ate Boston in Boston with sprinkles on top—as opposed to a *phrase* that serves as an *instruction to build* a concept from available lexicalized concepts? And can one tell the difference by introspecting? Third, I don’t mean to belittle “mere formal” concepts like $\text{HEXREGION}() \bullet \text{REPUBLIC}()$. I think MACs are extremely valuable things. I think we reason with them all the time. Our logical capacities may well depend on them. But I think we obscure their value by not distinguishing them from their more intuitive prelinguistic kin.

Unsurprisingly, I also have sympathy for the idea that a MAC may be “tweaked” or “enriched” in various ways prior to ‘ \exists ’-closure; cp. Bach (200x), Recanati (200x). If sentential

logical forms are of type $\langle e, t \rangle$, then quite apart from pragmatic considerations, there is generally a gap between logical forms and the thoughts that speakers can evaluate for truth or falsity in contexts. (In terms of the recipe metaphor, this leaves room for seasoning to taste, within constraints.) I am also open to the idea that binary evaluability of existentially closed thoughts does not always correspond to the classical true/false distinction; see, e.g., Dummett (197x) and Wright (198x). In some cognitive domains, we may need other notions of positive/negative evaluation. And we need not specify meanings for the “logical” vocabulary of a Language in truth-theoretic terms: ‘not’, ‘or’, and ‘every’ can be treated as instructions to construct monadic concepts. For these functional words, truth-theoretic models are especially illuminating. Still, these are hardly poster cases for capturing semantic compositionality in plausibly externalistic terms. It takes deep commitment to a theory to maintain that if every dog barked, ‘every’ is satisfied by $\langle \{x: x \text{ is a dog} \}, \{x: x \text{ barked} \} \rangle$. But these are matters for subsequent chapters.

By contrast, I can feel the pull of saying that ‘grommet’ is satisfied by grommets, and that one doesn’t really understand this word until one knows something substantive about what it applies to—e.g., that grommets are little metal rings, like the ones on tents that you can put a stake or piece of string through. But I think the pull should be resisted (cf. Higginbotham [199x]) in favor of the idea that ‘grommet’ is an instruction to fetch a concept that can be predicate-conjoined with others. Not even the meaning of ‘grommet’ is a grommet, or the grommets, or the logically possible grommets. If two speakers lexicalize substantively *different* concepts, then of course, there is a sense in which they end up having homophones as opposed to the same word. Such speakers associate the word-sound with different meanings. Instructions to fetch concepts are individuated in part by the concepts they are instructions to fetch. And perhaps some concepts are individuated in part by what they are concepts of. But that tells us little.

Suppose that a speaker S uses a word W to fetch a concept C, and thereby thinks about some things that differ from the things that S' thinks about by using a concept C' fetched with word W'. It *doesn't* follow that W and W' differ in meaning. In my view, it is unproductive for semanticists to worry about what it is for two speakers to link a word-sound to *the very same* concept, as if the study of meaning must conform to a certain picture of communication. And I don't think philosophers have established, or even made it plausible, that unnoticed differences in the environments of twins can make for a difference in concepts that precludes the twins from pairing a given word-sound with the same *meaning* (even if by all the usual linguistic tests, the twins are using the same word). I agree that Twin Earth is relevant *if* the meanings of expressions in a Language determine extensions for those expressions at every possible world. But why think meanings do that? Surely not because grommets tend not to be in the head.³¹

We can hypothesize that water is H₂O. And it is interesting that we humans seem to have “kind concepts” that let us classify things in terms of *their* natures, despite our ignorance. But we may also have more than one concept with which we can think about water. So why think the word ‘water’, regularly acquired by children, determines an extension that includes samples of H₂O *but not* superficially similar and equally drinkable substances? In my view, the meaning of ‘water’ is not tied to H₂O, and the meaning of ‘grommet’ is not tied to grommets.

Meanings are neither truth-theoretic properties of expressions nor concepts. Each of these views has virtues. But we can do a little better by taking meanings to be concept construction instructions that compose in Conjunctivist fashion. That’s the picture. The rest is development.

Notes

1. Some readers may not like talk of meanings or fitting. In the end, my claim is that such talk is fine if my proposals are correct. But I return to questions about how meanings are related to understanding. And I allow for the hypothesis that meanings are truth-theoretic properties of expressions, even this implies that intuitively nonsynonymous expressions can have the same meaning—e.g., being true iff a lawyer/attorney xeroxed/photocopied a book by Twain/Clemens. *If* this hypothesis is correct, it may be better to eschew talk of meanings, instead of allowing for expressions that have the same meaning despite being understood differently. But that’s a big if.

2. A language with a trillion lexical items and a billion modes of grammatical combination, each with its own significance, could be semantically compositional in the sense just described. And this principle says nothing about composition being *local* in any sense that requires (for example) the meaning of ‘The dog saw me’ to be determined by the meanings of ‘The dog’ and ‘saw me’, given how those *phrases* are combined to form the sentence; see Pelletier (19xx), Szabo (200x), and further references there. For now, I set aside the delicate question of whether substitution of synonyms always preserves meaning, modulo quotation. But at least for these purposes, let’s say that each expression has exactly one meaning, treating ambiguity (in spoken languages) as homophony: cases of distinct expressions that sound the same. I also defer discussion of the radical suggestion that for languages like English, all semantic compositionality theses are false because expressions of such languages don’t have meanings. COMPLICATION...

In light of Mates (195x) and discussions of “logophors,” one might think that certain lexical items and/or modes of combination—say, those associated with sentential complements—can make the meaning of a phrase depend on *nonsemantic* properties of its constituents; see, e.g., Forbes (19xx, 19xx) and Pietroski (199x). If meanings are truth-theoretic properties, then most accounts of opacity imply that the meaning of ‘thinks that a lawyer bought a book by Twain and xeroxed it’ need not be preserved by replacing a constituent with a synonymous expression.

3. See Szabo (200x), who usefully discusses semantic compositionality as a logically contingent supervenience thesis along the following lines: the meaning of a complex expression *depends* on the relevant lexical meanings and their arrangement, at least in the sense that two expressions cannot *differ* in meaning without differing in some lexical or compositional way. I assume, at least for the illustration, that ‘or’ signifies inclusive disjunction. But we can imagine a third language in which ‘[[AB] Y]’ applies to x iff (x is an armadillo *xor* x is blue) *xor* x is yellow; where ‘*xor*’ signifies exclusive disjunction.

4. A certain kind of philosopher will refuse, insisting that no real question was left behind. But given our theorist’s interrogative soliloquy, I will carry on and address such philosophers later.

5. See, e.g., Chomsky (...), Jackendoff (...) Baker (...), Lasnik (...), Hornstein (...), Uriagereka, other CITATIONS. One needn’t endorse Chomsky’s “minimalist” program to see a trend towards simpler grammars, over the last fifty years. Of course, theoretical simplicity often comes at the cost of positing complexity in the relation between what is readily observable (by a child or a theorist) and the underlying reality. But it is equally true that an intuitively simpler syntax (see

Culicover and Jackendoff [200x]) may posit complexity in the relation between grammatical form and meaning. So a simple conception of semantic composition, along lines defended here, can be part of a more broadly minimalist conception of the human language faculty.

6. See, for example, Higginbotham (1985), Larson and Segal (1995), Heim and Kratzer (1998)... My argument will be that predicate-conjunction is at least one of the operations, and that given this, others are unneeded and unwanted.

7. CITATIONS...Parsons, Hobbs, reference to Carlson on thematic roles, etc.

8. One can think of grammatical forms as expressions of Languages, abstracting away from any associated signals, and logical forms as expressions of Mentalese that “interface with” grammatical forms; where the relation of grammatical form (often called ‘LF’) to word order is abstract enough to make a straightforward mapping to logical form plausible. See Pietroski (200x) for review of this idea—with roots in both Frege (189x) and the study of transformational grammar—in the framework of a Conjunctivist semantics.

9. Readers of a Tarskian bent may have long since learned to translate “type $\langle t \rangle$ ” talk into talk of sentences as predicates of truth values—and likewise for the Fregean idea that sentences are (conditional) truth value *designators*, like the invented name ‘Flip’ whose *bearer* is the number 1 if Socrates is mortal and 0 otherwise— But for the purposes at hand, this is to grant that sentences are of type $\langle t, t \rangle$, given the standard conventions governing types. I assume, of course, that predicate-conjunction permits conjunction of two predicates of this type.

10. Given a Tarski-semantics, ‘S(e)’ is equally of type $\langle t \rangle$; it is true or false, relative to each assignment of values to variables, depending on whether or not Socrates is the value assigned to ‘e’. Given a Church/Montague-semantics, ‘S’ is of type $\langle e, t \rangle$; but at least recently, this kind of metalanguage has been used in combination with the hypothesis that in Languages, grammatical combination typically signifies function-application. And for reasons that will become clear below, I want to highlight certain differences between predicate-conjunction (‘•’) and proposition-conjunction (‘&’), as well as some intuitive similarities. In any case, we can stipulate that ‘ $\exists t$ ’ is of type $\langle \langle t, t \rangle, t \rangle$, and that expressions of type $\langle t \rangle$ include ‘ $\exists t \{ \exists e [M(e)] = t \}$ ’ and ‘ $\exists t \exists t' \{ \exists e [M(e)] = t \ \& \ \sim \exists e [M(e)] = t' \ \& \ \sim (t = t') \}$ ’. If ‘ δ ’ is of type $\langle e \rangle$, then ‘M(δ)’ is also of type $\langle t \rangle$. And we can either stipulate that ‘M(x)’ maps an entity x to \top iff x satisfies M, understanding satisfaction independently of truth, or stipulate that an entity x satisfies M iff ‘M(x)’ maps x to \top .

11. As its name suggests, function-application is an operation that maps a function F and a element x in the domain of F to F(x)—the value of the function given the entity as argument. For reasons that will soon become clear, function-application is equivalent (for these purposes) to an operation of concept-saturation.

12. Though we can also say that conceptualization is something that happens to words—or perhaps word-sounds, leaving room for the hypothesis (not mine) that meanings *are* concepts,

and that words of a spoken Language are sound-concept pairs.

13. Here, I say ‘one or more’ because *polysemy* is pervasive in Languages, and I don’t want to prejudge how many concepts are lexically associated with (for example) ‘book’ or ‘line’. The concept with which one thinks about a spatiotemporally located line may differ from the concept(s) with which one thinks about the abstract lines of geometries. But the sound of ‘line’ is not thereby ambiguous like the sound of ‘pen’. I assume that children can and do treat ambiguity as homophony: the animal-enclosure concept presumably differs from the writing-implement concept in a way that leads to the acquisition of distinct though homophonous words; while ‘line’ can serve as a locus for formally distinct concepts, which are similar along a significant cognitive dimension, which corresponds to the intuition that a line in the supermarket is more like a line in a right triangle than a ballpoint pen is like a pig pen. For better or worse, Languages exhibit both polysemy and homophony; and ‘ambiguous’ is polysemous, since it can be used to talk about . T polysemy, homophony, and various phenomena under the heading of “pragmatic” ambiguities.

14. That is, it’s not obvious to us theorists. But it is, evidently, child’s play.

15. But ‘*iD*’ will still be as rigid, in Kripke’s (198x) sense, as the index itself. And note that if *i* combined with a predicate of type $\langle e, t \rangle$ in a language where concatenation signifies predicate *saturation* (or function application), the result would be an expression of type $\langle t \rangle$.

16. We can say that such dimensions are **A**-tracked; see Larson and Segal (1995). An assignment of values to *variables of a Language* is not a mere sequence of *d* potential values of an indexed demonstrative. There are also “assignment slots” corresponding to indexical terms. So we can speak of **A**(*s*), **A**(*t*), etc. It is an empirical question how many dimensions of communicative situations are **A**-tracked. Speaker, time, place, and perhaps a few others are plausible candidates. But in my view, many dimensions that are not **A**-tracked are still relevant to the truth or falsity of what is said in such situations; cf. Stanley (200x). Indexicals are very special cases. Moreover, even the meanings of ‘now’ and ‘here’ do not always determine spatiotemporal coordinates finely enough for purposes of truth. But here, in this context, it is important to distinguish assignments (Kaplan-contexts) from communicative situations; see chapter four.

17. This does not require the further hypothesis that ‘performance’ is somehow derived from ‘perform’...see Remarks on Nominalization...

18. In chapter three, I’ll suggest that this is *why* Languages have prepositions, and why they differ from tensed verbs in ways that make little sense given FRET. But I don’t deny that prepositions can be homophonous. The ‘by’ of agency and the ‘by’ of spatial proximity are presumably distinct words. And there is probably more than one word pronounced ‘in’. But not *many*.

19. Citations, including Gruber, Fillmore, Chomsky, Jackendoff, Dowty, Baker...
More details in chapter one

20. I do not deny that predicate-conjunction is derived from a more basic notion of proposition-conjunction; on the contrary, see chapter three. But if ‘Txyz & Nzxy’ is well-formed, then ‘&’

hardly signifies the simplest kind of proposition-conjunction. One can imagine a system with AND-gates, but without polyadic predicates, variables, and an operator that permits conjunction of open sentences with different numbers of variables; where the n th variable in one open sentence can be the i th variable in the other ($n \neq i$). I find it hard enough to imagine how a *single* variable might be biologically implemented. So I want to see how far we can get by assuming that long ago, animal somehow got *one* step beyond AND-gates in having monadic predicate-conjunction (\bullet). I have no objections to rewriting ‘S•M’ as ‘Sx & Mx’, so long as one does not stipulate that a capacity to generate representations like ‘S•M’/‘Sx & Mx’ already brings with it the conjunction operation required for representations like ‘Txyz & Nzxy’. I have nothing against the stipulation; indeed, I adopt it. But given this standard use of ‘&’, replacing ‘S•M’ with ‘Sx & Mx’ is not mere *rewriting*.

23. Or more precisely, to generate Linguistic structures, theoretical description of which apparently requires appeal to labeled phrase markers of the sort regularly employed in transformational grammars....

cite Hornstein, and others, and paper on “Basic Operations”...cite Hurford

24. References, including Kayne and Baker

25. $\exists[\exists e\{\text{Crocodile}(e) \ \& \ \text{External}(\ , e)\} \bullet \text{PastJump}(\)]; \exists[\exists e\{\text{Armadillo}(e) \ \& \ \text{Internal}(\ , e)\} \bullet \text{PastFall}(\)];$ and $\exists[\text{Armadillo}(\) \bullet \exists e\{\text{Crocodile}(e) \ \& \ \text{R}(\ , e)\}]$

26. While ‘*Begriffsplans*’ is compact, and resonant of ‘*Bauplans*’, a compound like ‘*Begriffskonstructionsanleitungen*’ better displays the idea (and reminds us that words can be complex expressions as opposed to lexical elements).

27. Call them MacConcepts if you like. They’re not offended.

28. Relevant references here, including Halle...Chomsky and Halle...minimalist Chomsky
But I’m not making any commitments about phonology.

29. If $\pi:s$ can be identified with the lexical pronunciation, then meanings are like pronunciations.

30. Click: Did you just say that political organization is determined by geography?

Clack: I did. Democracy never flourishes in triangular regions, and people who inhabit a hexagonal country always have a constitutional monarchy.

Click: France is a hexagonal republic, you dolt.

Clack: By which you mean that something is France and hexagonal and a republic?

Click: No, I mean that your generalization is false—counterexamples by France and its inhabitants—and that you’re a dolt.

31. There is H₂O in the head, on a regular basis. But I have been assured, by experts, that this is philosophically irrelevant. I have also been assured that *if* there a language faculty in the head, that too is philosophically irrelevant—and that given arguments for externalism, a science of linguistics cannot teach us that meanings are internalistic properties of linguistic expressions.

