Chapter Five: Massively Monadic, Potentially Plural

In this chapter, I return to the idea that lexicalization is a process in which available concepts are used to *introduce* concepts that can be conjoined in certain ways.

1. Introduction and Foreshadowing

In chapter one, I posited two combinatorial operations, M-junction and Θ-junction. Recall that only monadic concepts can be M-joined: combining Φ(_) with Ψ(_) in this way yields another monadic concept, Φ(_)^Ψ(_), which applies to whatever both Φ(_) and Ψ(_) apply to. But a dyadic concept can be Θ-joined with a monadic concept: combining Θ(_, _) with Φ(_) in this way yields the monadic concept ∃[Θ(_, _)^Φ(_)]; where the internal/right slot of Θ(_, _) is linked to *the* unsaturated position of the adjacent monadic concept, which is also the position that ∃ closes. So ∃[Θ(_, _)^Φ(_)] is a concept of things that bear the dyadic relation—perhaps a thematic relation that events bear to certain participants—to at least one thing that falls under the concept Φ(_). In my view, a phrase formed by combining two lexical items is an instruction to fetch a corresponding pair of concepts and join them in one of these two ways. If this is correct, then concepts fetched via lexical items must be monadic or dyadic, even if this requires that fetchable concepts be *introduced* as opposed to being merely *labeled* in the process of lexicalization. I think there is considerable evidence for this view. More generally, I think lexicalization is like Fregean definition in being a cognitively creative process, though without the hierarchy of conceptual types; see §§1.4 and 1.5 of chapter two.

1.1 Junctions without Variables

The operation of M-junction permits construction of concepts like RED(_)^COAT(_). As discussed in chapter one, such concepts seem to be governed by a natural logic that typically licenses conjunct-*reduction*, but licenses conjunct-*expansion* in certain special contexts like the scope of negation. As the medieval logicians knew—and as Frege's work highlighted—a mere logic of predicates is inadequate even for natural thought, which involves at least some relational notions. But given just a few *atomic* dyadic concepts, Θ-junction makes it possible to build boundlessly many concepts that have dyadic constituents, as in STAB(_)^∃[DONE-WITH(_, _)^KNIFE(_)] and ∃[DONE-TO(_, _)^COLONEL(_)]^∃[DONE-BY(_, _)^[SPY(_)^∃[HAS(_, _)^[RED(_)^COAT(_)]]]]. This is so, even though neither form of junction yields non-monadic concepts. Correspondingly, both operations can apply to unsaturated concepts whose "slots" are not indexed with *variables*.[1]

For many purposes, one can ignore the differences between ∃[DONE-WITH(_, _)^KNIFE(_)] and its Tarskian analog ∃X[DONE-WITH(X′, X) & KNIFE(X)]. But the Tarskian apparatus, which also permits ∃X′[Θ(X, X′) & R(X′, X″, X‴)], is far more powerful than the restricted operation of Θ-conjunction. Likewise, RED(_)^COAT(_) is more modest than its Tarskian analog. My hypothesis is that M-junction is a natural operation, somehow implemented by biological circuitry, and that this mode of combination identifies the unsaturated position of one monadic concept with that of another. Likewise, Θ-junction is posited as an operation that links "the first" unsaturated position of a dyadic concept with the unsaturated position of a monadic concept.

If human linguistic expressions are concept construction instructions (Begriffsplans) that repeatedly call for these operations, then "expression reduction"—e.g., from 'red coat' to 'coat', or from 'stabbed a colonel with a knife' to 'stabbed a colonel'—would often correspond to compelling inferences. But lexicalization would have to be a process that links pronunciations to

---

[1] Jacobson (1999) also urges a "variable free" semantics, though with different initial assumptions human languages, linguistic meanings, and the importance of constraints on homophony.

concepts that can be M-joined or Θ-joined, even if such concepts have to be introduced. So evidence that lexicalization *is* such a process can support the proposed conception of composition, at least if many prelinguistic concepts are like GAVE(_, _, _) and BESSIE in *not* being joinable via the restrictive operations. In this chapter, I provide such evidence. Chapter six is devoted to compositional details and the more general claim that all phrasal meanings are instructions for how to assemble conjunctive monadic concepts.

On this view, lexical items can be used to fetch dyadic concepts, but composition preserves only *one* unsaturated position. In terms of the Fregean hierarchy of concepts: phrases correspond to concepts of types <e, **t**>; and the lexicon of a human language is only slightly less restricted, permitting access to some concepts of type <e, e**t**>. But if meanings are not used to access or build concepts of types <e> or <**t**>, with talk of entities and truth values being otiose, we may as well call the two basic semantic types <M> and <D> for 'monadic' and 'dyadic'. And if all phrases correspond to concepts of type <M>, then not only is there no hierarchy of non-basic types of the form <α, β>, there are *no* non-basic types. That said, my proposal allows for (and requires) lexical items and phrases corresponding to concepts that are *plural*—and quantification that is *second-order*—in Boolos' (1998) sense. Indeed, the idea developed in this chapter and the next is that human I-languages let us introduce and systematically combine concepts that can be viewed as composable symbols of a *second-order extension* of a massively *monadic fragment* of the (first-order polyadic) language, PL, that was described in chapter two.

1.2 Introducing vs. Labeling

In chapter one, I readily granted that humans have relational concepts like CHASE(_, _) and GIVE(_, _, _); where the dyadic concept applies to ordered pairs <α, β> such that (at some time t) α chases β, and the triadic concept applies to triples <α, β, γ> such that α gives β to γ. But it doesn't follow that CHASE(_, _) is the meaning of 'chase', or that this word meaning is dyadic.

For reasons discussed in chapter three, I assume that verbs are semantically associated with "event positions" of some kind, even if their relation to mind-independent events remains unclear. But this assumption is compatible with many specific claims about how words of a human I-language combine, and how verb meanings are related to concepts and to the meanings of sentences like (1), in which 'chase' combines with more than one grammatical argument.

(1) A girl saw a dog chase a cow.

One can hypothesize that 'chase' is true of ordered triples <ε, α, β> such that ε is a chase by α of β—or in overtly mentalistic terms, that 'chase' accesses a triadic concept with slot for events. And perhaps 'give' accesses a tetradic concept that applies to <ε, α, β, γ> iff ε is a giving by α of β to γ. This assumes an operation of saturation, or its equivalent, as the correlate of combining a verb (or verb phrase) with a grammatical argument. But another possibility is that both 'chase' and 'give' access monadic concepts of events, with grammatical arguments corresponding to concepts like ∃[AGENT(_, _)^DOG(_)] and ∃[PATIENT(_, _)^COW(_)]. Given Θ-junction, it is also possible that verbs access dyadic concepts like PASSIVE-CHASE(_, _) and PASSIVE-GIVE(_, _); where such concepts apply to event-individual pairs (see Kratzer 199x). This permits concepts like ∃[AGENT(_, _)^DOG(_)]^∃[PASSIVE-CHASE(_, _)^COW(_)]. But if combining a verb with an argument always calls for Θ-junction, this constrains hypotheses about the meaning of 'chase'; CHASE-BY-OF(_, _, _) cannot be Θ-joined with any concept.

In this sense, it isn't obvious which Fregean type 'chase' corresponds to, even in a specific construction like (1): CHASE(_) is of type <e, **t**>; PASSIVE-CHASE(_, _) is of type <e, e**t**>; CHASE-BY-OF(_, _, _) is of type <e, <e, e**t**>>. But it seems unlikely that in (1), 'chase' is used to access a concept of chaser-chasee pairs. Perhaps children initially connect the pronunciation of

‘chase’ with CHASE(_, _). But if so, then connecting the pronunciation with the meaning of ‘chase’ requires further work. Similarly, if combining ‘chase’ with a grammatical argument calls for Θ-junction, then acquiring the lexical item requires more than connecting the pronunciation with CHASE-BY-OF(_, _, _). If children simply acquire and label PASSIVE-CHASE(_, _) or CHASE(_), they don’t need to *introduce* some such concept that has no slot for chasers. That is compatible with my proposal. But then we need to explain the intuition that every chase has an agent.

If combining a verb with an argument always calls for Θ-junction, this also constrains hypotheses about the meaning of a name like ‘Bessie’; BESSIE cannot be Θ-joined with any concept. But in any case, the available combinatorial operations constrain the options regarding which concepts can be accessed with words. And if only for this reason, it seems unlikely that lexicalization is always a process in which an available concept is simply labeled with an item that inherits its type from the labeled concept. It seems that any plausible view of composition will require a dovetailing conception of lexicalization according to which some concepts are used to introduce formally distinct but analytically related concepts.

One might have hoped for the following transparent mapping from conceptual types to linguistic types: lexicalizing a denoting concept yields a proper name of type <e>; lexicalizing a monadic concept yields a noun or intransitive verb (or adjectival/adverbial expressions) of type <e, **t**>; lexicalizing a dyadic concept yields a transitive verb, of type <e, e**t**>; etc. Ignoring tense for simplicity, one might have expected ‘chase’ and ‘Bessie’ to be labels for CHASE(_, _) and BESSIE, with ‘chase Bessie’ indicating CHASE(_, BESSIE); where this partly saturated concept applies to an individual α iff (at some time t) α chases Bessie. Analogizing to 19th century chemistry, one could hypothesize that ‘chase’ inherits a valence of –2 from the concept lexicalized, ‘Bessie’ inherits a valence of +1, and so ‘chase Bessie’ has a valence of –1; cp. the compound Hydroxide, $(O^{-2}H^{+1})^{-1}$. This would capture the idea that ‘chase Bessie’ is like the intransitive verb ‘bark’, which can combine with a single argument to form a complete clause, much a Hydroxide ion can combine with a Hydrogen atom to form a stable molecule of water.

As we’ll see, there is more than one way of understanding the valence analogy. But for now, think of a verb’s valence as a (hypothesized) *structural* property of that verb—akin to the adicity of an unsaturated concept. The familiar picture is that a verb—or more generally, a lexical predicate—provides a “sentence frame” in much the way that an unsaturated concept provides a “thought frame;” see chapter one. Correlatively, one might say that names—or more generally, referential devices—can saturate lexical predicates, each of which “takes” a certain number of *grammatical* arguments, perhaps subject to certain typological or “selectional” restrictions. On this view, Frege was wrong to think that human linguistic expressions exhibit subject-predicate structure *as opposed to* predicate-argument structure; expressions are said to reflect both the saturated/saturating contrast and adicity contrasts among unsaturated concepts.

As noted in chapters one and two, this simple idea does not yet accommodate quantificational constructions like ‘chased every cow’, given the limits of saturation (as opposed to Church-style function-application) as a composition operation. But for these purposes, one might want to bracket questions about quantification and ignore the formal differences between CHASE(_, _) and λX.λX′.⊤ ≡ CHASE(X′, X). So let’s concessively extend the notion of labeling to include the additional apparatus required to construct λX.λX′.⊤ ≡ CHASE(X′, X) from CHASE(_, _). This relatively minimal re-presentation of a conceptual content preserves the concept’s original “argument structure;” likewise for recoding GIVE(_, _, _) as λX.λX′.λX″.⊤ ≡ GIVE(X″, X′, X). But it isn’t mere labeling if a mind initially connects the pronunciation of ‘chase’ with CHASE(_, _) but then uses this dyadic concept to introduce λX.λX′.λX″.⊤ ≡ CHASE-BY-OF(X″, X′, X) and

connect this triadic concept with the pronunciation of 'chase'.

Creating a concept of *higher* adicity, in ways that allow for conjunctive adverbial modification, reflects a more significant departure from the initial concept. It also reflects a capacity to go *up* the Frege-hierarchy, using a concept of type <e, e**t**> to introduce one of type <e, <e, e**t**>>. If a mind initially connects the pronunciation of 'chase' with CHASE(_, _) but then uses this concept to introduce concept to introduce λX.λX′.⊤ ≡ PASSIVE-CHASE (X, X′) or λX.⊤ ≡CHASE(X), that is also interesting. Of course, evaluating any such claim about lexicalization is hard, if only because it is unclear which concepts are natural for children independent of acquiring words. To understate, we don't have good theories of which concepts are available to children as natural candidates for initially labeling with pronunciations. But some familiar generalizations tell against the idea that lexicalization is merely a process of labeling concepts, and in favor of the idea that at least often, lexicalization is a process in which available concepts are used to introduce formally new concepts that can be M-joined or Θ-joined.

These same considerations may also tell against appeals to "grammatical argument structure," depending on how that vexed notion is understood; see Borer (2005), cp. Ramchand (2008); Lohndal (2014). But in the next section, I just state the generalizations somewhat starkly. Section three offers clarifications and qualifications, along with discussion of some implications. In section four, I extend the notion of concept introduction beyond cases of new adicities, and suggest that lexicalization is a process in which available concepts are used to introduce concepts that are "number neutral"—neither singular nor plural, neither count nor mass. This will connect back to reasons for analyzing verbs in terms of monadic predicates of events, following Schein (1993, 2002, forthcoming), who drew on Boolos' (1998) account of second-order quantification as plural quantification. This will also prepare the way for chapter six. But the next order of business is to argue against the idea that lexicalization is merely a process of labeling concepts.

2. Four Generalizations

In this section, I note four generalizations, each of which is familiar but often set aside as a complication as opposed to a telling symptom of lexicalization and composition in human languages. First, while humans have denoting concepts, proper nouns seem not to be lexical items of type <e>, textbook treatments to the contrary notwithstanding. On the contrary, the lexical noun 'Bessie' seems to fetch a monadic concept, perhaps one that applies to things called 'Bessie'; though the atomic noun can combine with a functional item to form a name that corresponds to a complex concept that applies (in contexts) to only one individual. Second, unsaturated concepts seem to be lexicalized with words that are often flexible about how many arguments they combine with in sentences. And the pattern of flexibility suggests that these words are used to fetch monadic or dyadic concepts. Third, "supradyadic" concepts like GIVE(_, _, _) are not lexicalized with verbs that take more than two grammatical arguments. Conceptually essential participants—e.g., the recipients in events of giving (or donating)—are often represented with prepositional phrases that are grammatically optional. Fourth, there is a constraint on how the subject/object asymmetry is related to the agent/patient distinction, as if concepts that have an unsaturated position corresponding to agents are lexicalized with words that subsequently fetch related but distinct concepts that do not have such a slot.

2.1 Limited Relations

This last generalization is often discussed in other contexts. Transitive constructions respect a substantive constraint on how "thematic order" can be expressed.[2] Presumably, 'chase'

[2] See Gruber (1965), Fillmore (1968), Dowty (1991), Baker (1997), and further references there.

lexicalizes a relational concept. But even if this concept already has a slot for events, that doesn't explain why in (1), 'a dog' corresponds to the chaser of the chasing.

(1) A girl saw a dog chase a cow.

We can assume that 'a dog chase a cow' has the following constituency structure, with 'a cow' as the internal argument of 'chase': [[a dog] [chase [a cow]]]. But this still doesn't explain why 'a cow' corresponds to the thing chased. Likewise, (2) implies that was Juliet was a kisser,

(2) Everyone saw Juliet kiss Romeo.

but not that Romeo was.

This thematic contrast might be described by saying that 'chase' corresponds to the concept indicated with (1a) as opposed to (1b), and likewise for 'kiss';

(1a) $\lambda X''.\lambda X'.\lambda X.\top \equiv \text{CHASE-BY-OF}(X, X', X'')$

(1b) $\lambda X'.\lambda X''.\lambda X.\top \equiv \text{CHASE-BY-OF}(X, X', X'')$

where by stipulation, CHASE-BY-OF(_, _, _) applies to $<\varepsilon, \alpha, \beta>$ iff ε is a chase *by* α *of* β, and order of function-application tracks order of grammatical arguments. On this view 'Fido chase Bessie' corresponds to $\lambda X.\top \equiv \text{CHASE-BY-OF}(X, \text{FIDO}, \text{BESSIE})$; where this function maps ε to $\top$ iff ε is a chase *by* Fido *of* Bessie. But this just encodes the contrast. We could have stipulated that CHASE-OF-BY(_, _, _) applies to $<\varepsilon, \alpha, \beta>$ iff ε is a chase *of* α *by* β. In which case, (1b′)

(1b′) $\lambda X''.\lambda X'.\lambda X.\top \equiv \text{CHASE-OF-BY}(X, X', X'')$

is equivalent to (1b), not (1a). And we have no reason for thinking that CHASE-OF-BY(_, _, _) is less natural, as a concept, than CHASE-BY-OF(_, _, _).

Put another way, we can imagine a verb 'asechay' that corresponds to (1b). So we can imagine a language such that 'Bessie' is the direct object of the verb in (3),

(3) Fido asechayed Bessie.

but this imagined sentence means that there was a chase by Bessie of Fido. Similarly, we can imagine a language Glisheng, in which the pronunciation of 'chase' is associated with the concept CHASE-OF-BY(_, _, _). In the Glisheng sentence (4),

(4) Fido chased Bessie.

'Bessie' is the direct object of the verb and so corresponds to the chaser-role, while 'Fido' corresponds to the chasee-role. So we face the question of why human children regularly acquire verbs that are thematically like 'chase' as opposed to 'asechay'. Or put another way, we face the question of why human children regularly acquire languages like English as opposed to Glisheng.

One can describe the generalization by elaborating (1a) and (1b) as (1aa) and (1bb).

(1aa) $\lambda X''.\lambda X'.\lambda X.\top \equiv \text{AGENT}(X, X') \;\&\; \text{CHASE}(X) \;\&\; \text{PATIENT}(X, X'')$

(1bb) $\lambda X'.\lambda X''.\lambda X.\top \equiv \text{AGENT}(X, X') \;\&\; \text{CHASE}(X) \;\&\; \text{PATIENT}(X, X'')$

Then one can say that the external argument of 'kiss' is associated with AGENT(_, _) while the internal argument is associated with PATIENT(_, _); cp. Parsons (1990) on "subatomic" semantics. Alternatively, one can replace CHASE(X) & PATIENT(X, X″) with PASSIVE-CHASE(X, X″). But either way, this still leaves the question of why the pronunciation of 'kiss' is connected with the first triadic meaning as opposed to the second. One possible answer is that while verb meanings can be triadic, meanings like (2bb) are prohibited. Perhaps human *composition* operations can generate the meaning of (3), but while (2aa) is a possible *word* meaning, (2bb) is not. I don't think there is a plausible source for any such thematic constraint on lexical items. So my suspicion is that neither (2aa) nor (2bb) can be a human word meaning: 'chase' fetches the monadic concept CHASE(_), or perhaps the dyadic concept PASSIVE-CHASE(_, _); 'asechay' would fetch the same concept, not any triadic concept with a slot for chasers. And if that's right, I think it tells against the idea that verbs simply label the concepts they lexicalize.

On my view, verbs do not have thematically elaboratable meanings. So the order of the chaser/chasee slots in CHASE(_, _, _) or CHASE(_, _) is irrelevant. Indeed, we need not assume that all polyadic concepts impose an *order* on their arguments. At least some relational concepts may be more like hubs with spokes; a conceptual hub may associate each of its spokes with a particular role, though not in any canonical order. In any case, whatever the nature of the concept lexicalized with 'chase', a child may associate this lexical item with a (perhaps introduced) concept that has a slot for chases, no slot for chasers, and perhaps no slot for chasees. If 'chase' corresponds to CHASE(_), the verb phrase [chase [a cow]] corresponds to the monadic concept of events CHASE(_)^∃[PATIENT(_, _)^COW(_)]. If 'chase' corresponds to PASSIVE-CHASE(_), then [chase [a cow]] corresponds to ∃[PASSIVE-CHASE(_)^COW(_)]. Either way, the phrase corresponds to a concept of cow-chasings; and while such events have agents, neither Θ-junction reflects this fact. But one can say that the verb phrase [a dog [chase [a cow]]] corresponds to ∃[AGENT(_, _)^DOG(_)]^CHASING-OF-A-COW(_); where the second conjunct is spelled out in terms of CHASE(_) or PASSIVE-CHASE(_, _).[3]

The familiar idea is that one way or another, the external argument of 'chase' is associated with the concept AGENT(_, _). That still leaves the question of why children don't acquire languages that associate the external argument of 'chase' with PATIENT(_, _). Saying that 'chase' corresponds to PASSIVE-CHASE(_, _) begs this question; let UNERGATIVE-CHASE(_, _) apply to <ε, α> iff ε is a chase *by* α. I suspect that it is not arbitrary that internal arguments are not associated with AGENT(_, _). But even if it is, we can at least capture the relevant generalizations by saying that certain grammatical relations are associated with certain thematic relations. By contrast, it's hard to see how to even report the generalizations if verbs can have meanings that are triadic, as opposed to at most dyadic.

As briefly noted in chapter two, there seems to be an analogous constraint on sentences formed by combining a quantificational determiner like 'every' with two predicates as in (5).

(5) Every cow is brown.

Presumably, 'every' lexicalizes some relational concept—perhaps INCLUDE[Ψ(_), Φ(_)], with Φ(_) corresponding to 'cow'—that yields a complete thought when saturated by two monadic concepts. One can imagine a word 'ryev' that is otherwise like 'every', except for expressing the concept SUBSET[Ψ(_), Φ(_)]. Then (6) would have the meaning of (7).

(6) Ryev cow is brown.

(7) Every brown thing is a cow.

But as reviewed in chapter six, there is no such determiner. I think this is because *pace* standard claims, determiners do not fetch concepts of second-order relations; see Pietroski (2005a, 2006). But in any case, I will be highlighting the *absence* of words that might be expected if lexicalization were simply a matter of labeling available concepts.

To take another example, why do we construct phrases like 'is bigger than Fido'? Why don't we label BIGGER(_, _) with a monomorphemic transitive verb, and use sentences like (8)?

(8) *Sadie bigs Fido.

Or put another way, why don't we use *more* morphemes to say what we say with (9)?

---

[3] On the former view, verbs are treated like common nouns in terms of valence. On the latter, verbs are treated more like dyadic prepositions. For now, I don't want to take a hard line on this issue; see Svenonius (2010) for relevant crosslinguistic discussion. I return to the issue in chapter six. Appeal to PASSIVE-CHASE(_, _) is compatible with a saturationist conception of composition according to which 'chase Bessie' corresponds to PASSIVE-CHASE(_, BESSIE). But the idea that 'Bessie' is of type <e> may not be sustainable, given facts about names; see §2.4 below. So I don't think that this is a reason for favoring PASSIVE-CHASE(_, _).

(9) Sadie is big.

One can imagine a language in which (9) and (10) are ungrammatical, but (8) and (11) are fine.

(10) Sadie is a big mare.

(11) *Sadie bigs most mares.

So assuming that the concept lexicalized with 'big' is somehow relational, one wants to know why (9) and (10) are even possible, if relational concepts are typically lexicalized with predicates that take two or more arguments in sentences. One can posit covert morphemes. But this is to grant that the relational concepts are not simply labeled with open class items.

More generally, it seems that many "relational thoughts" are expressed with words that fetch monadic concepts. If only because (12) implies (13),

(12) Trees dropped leaves onto the ground, as a student dropped her books.

(13) Leaves dropped onto the ground.

there is pressure to say that the lexical item 'drop' *fetches* a concept that has no unsaturated slot for droppers, and that the transitive verb in (12) is formed by adding a covert morpheme to the intransitive verb in (13). But it seems less plausible that 'drop' *lexicalized* a concept with no slot for droppers, and that our initial concept of one thing dropping another has a monadic component. Thinking about dropping seems to differ from thinking about falling, in that the former seems to involve thinking about actions that have agents.

Relatedly, the concept of one thing *felling* another may have constituents that include a concept of falling. Perhaps in this case, conceptual structure reflects morphological structure, with the verb in (14) having the verb in (15) as a constituent; cp. Fodor and Lepore (200x).

(14) Lumberjacks felled trees in the forest.

(15) Trees fell in the forest.

But if DROP(_, _) is lexicalized with a "root" expression that does not itself take two arguments, so that using the lexical item to express a relational thought requires an extra morpheme, one wants to know why. My proposal is that 'drop' lexicalizes a relational concept that gets used to introduce DROP(_), which can be M-joined with concepts like ∃[AGENT(_, _)^STUDENT(_)].

2.2 Never Three or More

Constraints on the expression of *supra*dyadic concepts seem to be even more severe.

While (16) and (17) are possible, (16a) and (17a) are also good sentences.

(16) I bet you ten dollars that Fido is a mutt.

(16a) I bet ten dollars.

(17) I sold her a horse for ten dollars.

(17a) I sold a horse.

So prima facie, nether 'bet' nor 'sold' requires four arguments: 'that Fido is a mutt' may specify the content of a bet without being a grammatical argument of 'bet', even if the propositional content of a bet is conceptually essential to it, just as 'for ten dollars' may specify the amount for which a horse was sold without being a grammatical argument.[4] One can posit covert arguments in (16a) and (17a). But 'her' seems to be an optional constituent in (17). And with present tense, (16) implies an *offer*—not yet uptake—raising further questions about the role of 'you'.

I can't prove that no verbs take four arguments, two of which are often covert. But if such verbs are possible, in human I-languages, one wants to know why they are not more common.

---

[4] See Munro (1982) on the adjunctive character of complementizer phrases. There is also evidence that verbs like speech act (and propositional attitude) verbs fetch monadic concepts of event/states that have contents, as dyadic concepts of individual-content pairs; see Pietroski (2000c), Motomura (2003), Moltmann (201x).

For plausibly, we have tetradic concepts, even if we often lexicalize them with words like 'square' and 'trade'. that do not take four grammatical arguments. Moreover, given a saturationist conception of composition, one might expect verb phrases that exhibit the following structure: [[[sold a horse] her] [ten dollars]].[5] So perhaps human I-languages *don't* allow for verbs that take four arguments. And perhaps that is because phrasal combination does not signify saturation.

Similarly, we seem to have triadic concepts that get lexicalized with words like 'triangle' or 'between', which we then use in odd constructions like (18), as opposed to (18a).

(18) Sadie is between a dog and a cow.

(18a) *Sadie betweens a dog a cow.

But if we *could* label BETWEEN(_, _, _) with a verb that takes three arguments, why do we circumlocute and use (18)? Studies of ditransitive constructions like (19)

(19) Hubbard gave a dog a bone.

also suggest that the constituent verbs do not combine with three arguments corresponding to saturaters of GIVE(_, _, _); see Larson (1988), Baker (199x), and references there. Rather, it seems that (19) should be analyzed as a superficially different variant of (20),

(20) Hubbard gave a bone to a dog.

whose prepositional phrase corresponds to a monadic conjunct in a concept of the following form: $\Phi(_)\wedge\exists[\text{RECIPIENT}(_, _)\wedge\text{DOG}(_)]$.

This is the obvious analysis for (21) and (22).

(21) Hubbard kicked a bone to her dog.

(22) Hubbard kicked a dog a bone.

But if 'give' lexicalizes a concept that has a slot corresponding to 'a dog' in (19) and (20), one wants to know why 'give' does *not* express/fetch the concept $\lambda X.\lambda X'.\lambda X''.\top \equiv \text{GIVE}(X'', X', X)$—or (using 'E' mnemonically) $\lambda X.\lambda X'.\lambda X''.\lambda E.\top \equiv \text{AGENT}(E, X'') \,\&\, \text{CHASE}(E) \,\&\, \text{PATIENT}(E, X') \,\&\, \text{INSTRUMENT}(E, X)$—and *always* require three saturating arguments, as suggested by the "surface form" of (19). My proposed answer is that words cannot fetch supradyadic concepts. A verb cannot fetch $\lambda X.\lambda X'.\lambda X''.\top \equiv \text{GIVE}(X'', X', X)$ any more than it can fetch $\lambda X.\lambda X'.\lambda E.\top \equiv \text{CHASE-OF-BY}(E, X', X)$. A human I-language allows for some words that fetch atomic dyadic concepts, subject to constraints, but not words that introduce further polyadicity. And if 'give' fetches a concept that has no variable corresponding to indirect objects, this already tells against the idea that lexical items have grammatical valences that reflect the adicities of lexicalized concepts.

2.3 Flexibility is Normal

If a lexical item L labels a concept that has some adicity *n*, then other things equal, one would expect L to combine with *n* arguments in any complete sentence where L does not appear as an argument of a higher-order predicate. But lexical predicates seem not to exhibit fixed selectional features concerning the number of mandatory/permitted grammatical arguments. On the contrary, it seems that a concept is often lexicalized with a word that can combine with one or two arguments—or more, if indirect objects and/or 'that'-clauses count as arguments—to form a complete, active voice sentence. Moreover, lexical items often appear in passive forms, or nominal forms with *no* accompanying arguments.

[5] I assume that 'a horse' is the most internal argument in (17); but this isn't essential to the point. One wonders why 'for' is required in (17) if 'ten dollars' is an argument of the verb. Perhaps for whatever reasons, 'sold' can only assign three cases; though cp. 'traded a horse for a cow' and 'replaced a horse with a cow'.

For example, 'kick' is not limited to transitive constructions like (23).

(23) Hubbard kicked a bone.

Along with (22), (24-26) are also possible.

(22) Hubbard kicked her dog a bone.

(24) A bone was kicked.

(25) The baby kicked.

(26) Hubbard gave the bone a swift kick.

One can say that 'kick' is passive in (24), unaccusative in (25), and a noun in (26). But the point remains: 'kick' need not appear with two arguments. This flexibility bolsters independent arguments for "separating" thematic roles from lexical items, with grammatical arguments corresponding to independent conjuncts in complex concepts, and verbs corresponding to monadic concepts of events; see Schein (1993, 2002). But whatever the best diagnosis of lexical flexibility, if 'kick' lexicalizes a concept that has a slot for kickers, this raises the question of why (24) and (26) are even possible constructions; see Borer (2005), Ramchand (2008).[6]

We face further questions, given the need for "event positions," if combining verbs with arguments indicates function-application. Even if we set (25) and (26) aside for special treatment, we should ask if the meaning of 'kick' is triadic, with a position corresponding to the *subject* of (23). If the verb meaning is like $\lambda X.\lambda X'.\lambda E.\top \equiv$ KICKED(E, X′, X) in this respect, then not only does the passive (24) call for special treatment, active voice phrases like 'kicked the ball west' have dyadic meanings like $\lambda X'.\lambda E.\top \equiv$ KICKED(E, X′, THEBALL) & WEST(E). In which case, adding an adverbial modifier prior to the subject requires more than mere predicate conjunction: 'west' has to target the correct variable in a dyadic predicate; see Chung and Ladusaw (200x).

This invites the thought that verb meanings are more like $\lambda X.\lambda E.\top \equiv$ PASSIVE-KICK(E, X); where the internal variable 'X' corresponds to internal arguments of the verb (see Kratzer 199x, Harley 200x). But then verb *phrases* like 'kicked a ball' express monadic concepts, even if verbs themselves express concepts of certain event-participant pairs. This is already a large step in the direction urged here. Moreover, if combining a verb with an internal argument indicates saturation of a relational concept, one wants to know why verbs don't regularly express concepts of polyadic relations that are exhibited by independent entities.

We can instead view (22-26) as evidence that verbs do not express such concepts. And however we encode event positions, we need to distinguish (23) from this sentence *minus* whatever aspect of it corresponds to existential closure of the 'E'-variable. Note that (27)

(27) That Hubbard kicked it.

cannot be understood as a sentence which means that the demonstrated event was a kicking by Hubbard of a bone: while KICKED(THAT, HUBBARD, IT) seems to be a fine thought, we cannot express it with (27). One can say that event positions of verbs are special, perhaps in ways related to tense. But overt adverbial quantification is possible, as in (28);

(28) Usually, Hubbard kicked her dog a bone after dinner.

see Lewis (1975b), Kratzer (1986, 2012). So the question is *why* the remaining event position in (27) cannot be overtly saturated; see §2.1 of chapter six. My answer is that combining expressions does not signify saturation. In any given phrase, the grammatical arguments of a verb V are phrasal constituents to which V bears certain structural relations, which correspond to thematic concepts that can be conjoined with others. In (26), 'kicked' bears the relevant

---

[6] Perhaps several words share the pronunciation of 'kick'. But it is no explanation to say that each such word labels a different concept. If the noun corresponds to KICK(_), one needs reasons for saying that the verbs do not.

grammatical relation to 'That Hubbard'. But the verb cannot bear this relation to 'Hubbard' and *another* one to 'That'; there is no *third* structural relation for 'kicked' to bear to 'That'.

On this view, (27) presents no special difficulty. The hard questions are more about what sentences are, and how they differ from clauses like 'Hubbard kick a bone', if phrasal composition always yields an instruction for how to build a *monadic* concept. I return to this chapter six, where I show how a "complete sentence" can be viewed as an instruction for how to build a special kind of monadic concept, rather than a thought of type <**t**>.

2.4 What's in a Name

Saturationist conceptions of semantic composition are initially motivated by sentences like (28),

(28) Fido barked.

in which a name combines with a predicate. This is unsurprising, since Frege's hierarchy of concepts is rooted in appeal to thoughts and denoting concepts. So if lexical items label concepts, one expects proper nouns to be parade cases, with the subject of (28) standing for a denoting concept that can saturate $\lambda X.\top \equiv \text{BARKED}(X)$ or perhaps $\lambda X.\lambda E.\top \equiv \text{BARKED}(E, X)$. But empirical and theoretical considerations suggest that lexical proper nouns do *not* fetch concepts of type <e>. Rather, such nouns fetch monadic concepts of a special sort, as suggested by (29) and (30),

(29) Each of the Fidos barked.

(30) That Fido is the biggest one.

according to which the demonstrated Fido (as opposed to some other) is the biggest Fido.

Similar remarks apply to deictic pronouns, as in (31).

(31) He barked.

As we'll see, the subject of (31) is not plausibly viewed as a context sensitive lexical item of type <e>. The meaning of 'he'—as opposed to 'she', 'they', 'it', etc.—along with the need for indices as in (32), suggests that the subject of (31) is more like the noun phrase in (33).

(32) Fido_1 barked when he_1 arrived.

(33) A certain male animal barked.

And examples like (34) invite an independently plausible idea about proper nouns.

(34) That individual called 'Fido' is the biggest one.

These nouns, which can combine with (perhaps covert) determiners/demonstratives to form complex subjects, fetch metalinguistic monadic concepts; cp. Burge (1973), Katz (1994). That raises the question of why lexical proper nouns do *not* simply label concepts of type <e>. My suggestion is that human I-languages do not tolerate expressions of type <e>. In the course of lexicalization, denoting concepts are used to introduce monadic concepts that can be fetched and *conjoined* with others. If that's right, human I-languages may not tolerate expressions of type <**t**> either. Saturation and truth may be features of thought, not the linguistic faculty that lets us generate instructions for how to build monadic concepts. But to make that suggestion plausible, I need to flesh out the argument sketched in this first section.

3. Lexical Flexiblity

Suppose children lexicalize concepts that exhibit a range of Fregean types. If lexical items of a human I-language can exhibit the same range of types, and combine in Frege-Church fashion, then one would expect lexical valences to reflect the adicities of lexicalized concepts. So if this expectation turns out to be wrong, we need to consider other conceptions of human I-languages. Perhaps these biologically implemented procedures let us recursively generate (instructions for how to build) concepts of a special sort, and lexicalization lets us employ available concepts to introduce formally new *and less diverse* concepts that can be systematically combined in useful ways. In my view, the generalizations mentioned above point in this direction, and there is little

if any evidence for saying that lexical valence reflects conceptual adicity. To be sure, (35-38)

(35) Hubbard's dog arrived.
(36) *Hubbard arrived a dog.
(37) Hubbard put a bone on the floor.
(38) *Hubbard put a bone.

illustrate contrasts that must be accommodated. But examples of lexical constraint do not show that verbs have valences, at least not in any structural sense that *explains* why (36) and (38) are anomalous. We also need to accommodate the kind of lexical *flexibility* that 'kick' illustrates. Though before discussing words in more detail, let me enter a caveat about concepts.

3.1 Adicity Uncertainty

My view is not that speakers can intuit the adicities of lexicalized concepts. Pretheoretic suspicions may reflect, among other things, statistics regarding how often a given verb is used with how many arguments. Indeed, when a child acquires a word, there may not be any *one* concept that gets lexicalized; talk of "the" concept lexicalized may be a simplification. Some may doubt that children have concepts—mental representations that exhibit adicities—prior to having words. But I don't see how to explain human lexical acquisition, or even animal navigation, without ascribing concepts to creatures that do not (yet) have words. Still, one can have good reasons for positing lexicalized concepts of diverse adicities without knowing which adicity any particular lexicalized concept has.[7] And it is worth being clear about this, since part of the argument for my account of semantic composition does rely on a general assumption: children lexicalize concepts that are not already massively monadic, occasionally dyadic, and systematically conjoinable in ways that support the cognitive productivity described in chapter one. Again, I think lexicalization *introduces* concepts that can be M/Θ-joined, whatever the adicities of the concepts lexicalized.

Recall from chapter two that Frege showed how to reformat ordinary arithmetic thoughts, formulated in terms of concepts like NUMBER(_), in terms of logically interesting concepts like NUMBER-OF[_, Φ(_)]. For purposes of logic, one can bracket the question of whether ordinary humans somehow "have" the higher-order polyadic concepts, albeit inchoately. The important point here is that Frege's Begriffsschrift does not merely let us express thoughts in a way that supports communication; it lets us re-present conceptual contents in ways that permit fruitful definitions and new uses of certain formal operations. In my view, humans use I-languages to re-present conceptual contents by introducing logically *un*interesting (but M/Θ-joinable) concepts like CHASE(_) and GIVE(_). So I do assume that in many cases, such concepts were not already available for labeling via lexical items that simply pair the concepts with pronunciations.[8]

This is not a strong claim about what was available in principle, given the expressive power of pre-human languages of thought. An animal might have the concepts and capacities needed to introduce CHASE(_), yet never exercise the capacities as a human child does, frequently and spontaneously. Once my proposal is in place, there will also be room for a modification, according to which the reformatting job that I assign to lexicalization was partly done by an older biological innovation. Perhaps we had ancestors who used (and cousins who use) diverse

---

[7] Though see Lidz. et.al. (201x) for interesting experiments inspired by Gordon (199x). One might have hoped to reduce talk of KICKED(_, _) to talk of 'kicked' and "pre-conceptual" mental representations. But especially in light of examples like (22-26), I see no way of avoiding appeal to at least one concept corresponding (and prior) to the word.

[8] Given our natural concepts, we can represent flexibility with invented concepts that have no fixed adicity; see, e.g., Fine (2007). Perhaps natural concepts can likewise be "multi-variate." But I find it hard enough to think about concepts that have fixed adicities. And in any case, lexical items still seem to be monadic or dyadic.

concepts to introduce atomic concepts that can be joined, in ways that the initial concepts cannot, but without a capacity for pronunciation. But in any case, the cognitive productivity of human thought is mirrored by the systematic combinability of our words. And human children have a special capacity for lexical acquisition. So this capacity is presumably related to whatever is special about human thought.

I don't deny that many lexicalized concepts are monadic. But in any particular case, it is hard to know. And this point is not limited to verbs. Perhaps 'number' lexicalizes a monadic concept, and Frege's relational analog was invented. But 'triangle' may lexicalize a naturally triadic concept; cp. 'between'. We have sentences like (39) rather than (40), suggesting that 'triangle' is used to fetch a monadic concept.

(39) That is a triangle.

(40) *This side triangles that side the hypotenuse.

Though it doesn't follow that children simply label TRIANGLE(_) with 'triangle'. Lexicalizers may use a polyadic concept to introduce TRIANGLE(_); where the underlining indicates an introduced concept, though perhaps one that is also available to other animals. Similar points apply to 'square', 'trio', 'quartet', etc. This raises familiar questions. Do words for artifacts lexicalize concepts that relate constructed things to purposes? Do some nouns lexicalize concepts that relate things to parts/histories/times/places? But for just these reasons, one might be skeptical of any generalization like "common nouns lexicalize monadic concepts."

Recall 'mortal'. This classical example of a monadic predicate arguably lexicalizes a concept that relates individuals to events of death. We can speak of mortals, who fall under the concept of being (a) mortal. But we can also speak of mortal wounds, which are fatal. This is not to deny that we have an atomic concept MORTAL(_). But as discussed in chapter two, atomic concepts need not be primitive; cp. MARE(_). nd reflection on cases suggests that monadic concepts often mask underlying relationaltity. From this perspective, there was something right about the classical/medieval focus on monadic predicates, at least with regard to the distinctively human thoughts that are also natural for humans; although Frege was also right to stress the importance of a logic that accommodates relational notions. I'll return to this point.

Likewise, I assume that many verbs are results of lexicalizing polyadic concepts, even if the details are unclear. Children presumably have a concept—EAT(_, _) or EAT(_, _, _)—with which they can think about some things as (the agents of events of) eating other things. But we humans can also think about some things as (the agents of events of) grazing, dining, and "fueling up." We also know that cows who graze "graze on" something, even though we don't say 'Cows graze grass'. So it isn't crazy to think that both 'eat' and 'graze' lexicalize relational concepts yet fetch monadic concepts. Moreover, (41-44)

(41) John ate a tack. (42) John ate something.

(43) John had a snack. (44) John ate.

suggest that 'eat' lexicalizes a concept that is somehow normative; cp. 'dine'.

Note that (41) implies (42) only on a purely existential reading of (42) that does not imply (43); and so read, (42) does not imply (44). In this sense, (42) and (44) differ in meaning. Likewise, (41) does not imply (44), unless it is assumed that a tack is nutritive for John; see Chomsky (1986). So even if 'ate' has a grammatical valence of -2, and (44) has a covert direct object, it doesn't follow that the verb's valence matches the adicity of the lexicalized concept, which may have a *further* slot for the relevant norm. Of course, a concept like NUTRIFY(_, _) or REFUEL(_, _) can be satisfied by event-individual pairs, and be normative in the relevant sense. But just as a "normative consumption concept" need not have a slot for a norm, it need not have

a slot for a thing consumed. So whatever one says about the grammatical valence of 'eat', it is a substantive hypothesis that this valence matches that of any particular lexicalizable concept. More generally, it hard to see how the number of required arguments for verbs of consumption—'eat', 'dine', 'snack', 'nosh', 'nibble', 'graze', 'refuel', 'drink', 'gulp', 'sip', 'chug', etc.—could generally track the adicities of the lexicalized concepts.

This is compatible with various views about how the meanings of these verbs are related. But given that ATE-A-TACK(_) is not the result of imposing a restriction on ATE(_), but more like INGESTED-A-TACK(_), one possibility is that the verb in (44) has a covert internal argument that corresponds to a conjunct like ∃[PATIENT(_,_)^APPROPRIATE(_)]; where the normative monadic concept applies to things that are "fit" for the agent in question. Alternatively—or perhaps expressing the same idea another way—the address for 'eat' may be shared by a concept CONSUMED(_) that applies to events in which the thing consumed is nutritious for the relevant agent, and a more permissive concept INGESTED(_). Then the instruction ***fetch***@'eat' can be executed in more than one way. But perhaps the covert argument in (44) is an instruction such that executing the verb phrase requires that 'ate' be executed by fetching CONSUMED(_); cp. 'France is hexagonal'. I return to analogous issues in section four, in the context of thinking about "count to mass" implications like 'The dragons ate a cow, so each dragon ate some beef'.

3.2 Denoting Concepts and Proper Nouns

With the caveats in mind, let's go back to the point about names noted in section two. This can help bring out how empirically implausible the "lexicalization as labeling" idea is, at least if we assume that humans and other animals have denoting concepts. And if one wants to argue for a saturationist account of semantic composition, it does not help to deny that children have denoting concepts.

Consider (45-48), which suggest that 'Caesar' can be used to fetch a monadic concept.

(45) Every Caesar I saw was a tyrant.

(46) Every tyrant I saw was a Caesar.

(47) There were three Caesars at the party.

(48) That Caesar stayed late, and so did this one; but the other Caesar left early.

More precisely, there seems to be a *lexical proper noun* (LPN) 'Caesar' that is like the common noun 'tyrant' in this respect. Of course, names differ from common nouns. And evidently, given the contrast between (49) and (50), the subject of (49) is not a mere LPN.[9]

(49) Caesar arrived.

(50) *Tyrant arrived.

The subject of (49) is a name. But names may be complex expressions, consisting of an LPN and a determiner akin to 'That' in (48); where this determiner, covert in English, combines with LPNs but not common nouns. The sound of 'Caesar' might be the sound of both an LPN and a determiner phrase whose head is silent. This hypothesis is not *ad hoc*, given overt analogs of the posited determiner in other languages. For example, Spanish allows for both 'Juan' and 'El Juan' as devices for referring to a certain Juan. And even English allows for 'our John' ('the John I know', etc.) as a way of referring to a certain John who is suitably related to the speaker.[10]

---

[9] See Burge (1973) and many others, e.g., Katz (1994), Longobardi (1994), Elbourne (2005), Matushansky (2006).

[10] Similar remarks apply to Basque, German, Scandanavian languages, many dialects of Italian, and Greek (discussed below). One can call LPNs names. But I assume that for purposes of taxonomizing human linguistic expressions, theorists need to speak of *nouns* and subcategories thereof; cp. Chomsky (1965). Segal (2001) offers considerations that seem to favor <e>-type analyses of proper nouns. But in my view, these considerations can be accommodated by saying that LPNs *lexicalize* singular concepts and *fetch* monadic concepts.

I won't defend a specific proposal about the posited covert functional element. But to illustrate, and to stress that complex names can be used to designate rigidly (cp. Kripke 1982), imagine an indexable determiner *D* such that relative to any assignment **A** of values to indices: for each entity e and index *i*, (i) e is a value of 'D_i' iff **A** assigns e to *i*, and (ii) e is a value of 'D_i Juan' iff e is a value of both 'D_i' and 'Juan'. As discussed in chapter six, any such analysis can be recast in terms of instructions for how to build monadic concepts, and extended to pronouns: e is a value of 'D_i she' relative to **A** iff **A** assigns e to *i*, and e is female. One can further hypothesize that the indices index individuals via denoting concepts stored in memory, thus fostering the illusion that names are semantically atomic expressions of type <e>.

This leaves room for various views about the meanings of LPNs. In chapter one, I pretended that an LPN fetches a partly saturated concept of identity. But given multiple Caesars, it won't do to say that 'Caesar' fetches IDENTICAL(_, CAESAR). The LPN might, however, fetch a concept that applies to individuals called—or called with the pronunciation of—'Caesar'; cp. Katz (1994). That would accommodate the meanings of (45-48).

Like any other lexical item, the LPN 'Caesar' pairs a pronunciation π with a meaning μ. Suppose that π initially connected with a denoting concept JULIUS. If a child has access to a relational concept CALLED(_, _) that applies to individual-pronunciation pairs, and to some *concept* Π with which she can *think about* the pronunciation π, she might form a thought like CALLED(JULIUS, Π) and connect π with a lexical meaning that is used to fetch the complex monadic concept CALLED(_, Π). Initially, the child might apply this concept to exactly one individual, Julius. This is compatible with Julius being called with many things besides π. Indeed, friends of Julius may have never used any English pronunciation to call him; cp. Kripke (198x) on whether Socrates *was* so-called. But more importantly, a child might *think* that Julius is the only bearer of the name-sound π, yet have the capacity to think otherwise by forming a thought like CALLED(ROMERO, Π). If the child is led to form and endorse such a thought, she might eventually conclude that more than one thing is called with π.[11]

In short, an LPN can be used to *lexicalize* an atomic concept of type <e>, but used to *fetch* a complex (and in some sense metalinguistic) monadic concept that applies to more than one thing. On this view, it is not surprise that (51) is a possible sentence,

(51) The Smiths are coming to dinner.

and that it is easily heard as a claim about some people who share a surname. Moreover, as surnames remind us, even surface considerations suggest that many "proper names" are *not* grammatically atomic. The direct object of (52) seems to have two words as parts.

(52) At noon, I saw Tyler Burge.

Prima facie, 'Tyler Burge' is semantically related to 'Tyler' and 'Burge': a Tyler Burge is both a Tyler and a Burge. Though in a context where the only Tyler is also the only Burge, one can use (53) or (54) to talk about that one person. Examples like (55) extend this point.

(53) I saw Tyler at noon.

(54) I saw Burge at noon.

(55) Professor Burge saw a certain Doctor Smith, but not *the* Doctor Smith.

If 'Professor' and 'Burge' are expressions of types <e, **t**> and <e>, then 'Professor Burge' denotes a truth value. One can reply that the honorific differs from the common noun 'professor'. But that is a cost. It raises questions about how to accommodate 'a certain Doctor Smith' and

[11] Or perhaps endorse a thought like IDENTICAL(CAESAR, ROMERO). Tiedke (201x) and Erickson (2013) develop different but compatible applications of this Burge/Katz-inspired suggestion; cp. Graff-Fara (201x).

'*the* Doctor Smith', with italics indicating the familiar kind of emphasis.

I return to ambiguity hypotheses. But if LPNs fetch monadic concepts, this presents a puzzle if nouns of a human I-language *can* be used as labels for denoting concepts like CAESAR. If lexicalizers could label denoting concepts with distinctive nouns, and thereby acquire names that can combine with a predicate of valence *n* to form a predicate of valence *n-1*, then one would expect lexicalizers to do just this and become adults for whom (45-48) and (51-52) are defective. One can imagine children who grow up to be speakers for whom 'every Caesar' is like the nonsensical '∀x:c', where 'c' is a logical constant. But we were not such children. We acquired a language in which (45-48) and (51-52) are fine.

(45) Every Caesar I saw was a tyrant.

(46) Every tyrant I saw was a Caesar.

(47) There were three Caesars at the party.

(48) That Caesar stayed late, and so did this one; but the other Caesar left early.

So assuming that we had denoting concepts, and often used them to think about named individuals, lexicalizing these concepts was not a matter of simply labeling them with LPNs of type <e>. Such lexicalization led to the acquisition of nouns like 'Caesar', which have the distributions they have. These proper nouns, like their common counterparts, show all the signs of being lexical items that fetch monadic concepts. They can be pluralized as in (47), or be constituents of complex demonstratives as in (48). And note that 'one', as it appears in (48), is ordinarily a pro-form for nouns that are *not* used to denote particular individuals.

These uses of proper nouns tell against the idea that LPNs fetch denoting concepts.[12] But one might posit ambiguities. Perhaps speakers who use 'Caesar' to talk about two people—Julius and Romero—have *three* homophonous LPNs that are used to fetch JULIUS, ROMERO, and a complex concept like CALLED(_, Π). This posits both "saturating-LPNs" of type <e> and "monadic-LPNs." Ambiguity hypotheses are notoriously hard to refute; see Kripke (1979b). But examples like (45-48) suggest that for each name-sound, there is a monadic LPN. And positing additional LPNs, with meanings of another type, is unattractive in several respects.

Since many human languages permit complex names, in which monadic LPNs combine with a determiner, children acquiring English can presumably formulate such an analysis of (49).

(49) Caesar arrived.

But if complex-name analyses are available to children, positing an analysis with 'Caesar' as a saturating LPN makes (49) strangely ambiguous: the first word might be either an instruction to fetch a monadic concept like CALLED(_, Π) or an instruction to fetch any of several denoting concepts; yet the posited denoting meanings for (49) could be plausibly redescribed in terms of the monadic meaning. In this sense, appeals to "saturating-LPNs" are not just *theoretically* redundant, they require children who posit lexical ambiguities beyond necessity. And one might think that children would, if possible, adopt uniform analyses of LPNs.

Relatedly, appeal to saturating LPNs makes "noun" a disjunctive grammatical category, and not just because of the two semantic types. Saturating LPNs would not head phrases, much less phrases of their own type. Since common nouns can combine with other expressions to form noun phrases like 'angry tyrant who died', much as verbs can combine with other expressions to form verb phrases like 'quickly ate his lunch', one needs special reasons for saying that 'Caesar' is ambiguous as between (i) a monadic-LPN that can head the noun phrase 'angry Caesar who arrived yesterday', which can combine with a determiner like 'every', and (ii) a saturating-LPN

[12] If 'Tyler Burge' is as semantically distinct from 'Tyler' as 'Twain' is from 'Clemens', then we also want to know why 'Tyler Burge is a philosopher.' seems to follow from (52) and 'Every Tyler I saw is a philosopher.'

that cannot head such a phrase, but must always appear as a complement of some predicative phrase whose valence is then different from that of *its* head. It is worth remembering that given a valence/saturationist conception of composition, predicate-argument phrases differ crucially from their heads. But even accepting this departure from a minimalist ("bare phrase structure") ideal, appeal to (ii) seems to complicate grammar needlessly, given appeal to (i).

In English, (45-48) are in some sense "marked" constructions, compared with (49). This makes it tempting to think that the question is whether or not English happens to be a language in which name-sounds are ambiguous, as if both options are available to children. But it is an empirical hypothesis that saturating-LPNs are even *possible* human linguistic expressions. And if lexicalization can introduce monadic concepts, any appeal to saturating-LPNs stands in need of support—especially since there are human languages in which complex names, with monadic LPNs as constituents, appear regularly. Such languages highlight the question of whether or not the human language faculty even *permits* saturating-LPNs.

For example, in Greek, analogs of (49) are as anomalous as analogs of (50).

(50) *Tyrant arrived.

A bare LPN like 'Petros' is no better, at least discourse initially, than a bare common noun as the subject of a sentence. To talk about the contextually relevant person called 'Petros', one uses the masculine determiner 'o' to form the phrase 'o Petros'.[13] Any child can acquire such a language. And *if* English has saturating LPNs, along with monadic LPNs, any child can acquire such a language. So *if* the ambiguity hypothesis for English is correct: experience with English leads every normal acquirer to a lexicon with *enough* LPN entries, despite homophony and the grammatical possibility of monadic LPN analyses that would shorten the lexicon; and experience with Greek leads every normal acquirer to a lexicon without *too many* entries, despite the possibility of ambiguity and saturating LPN analyses that would lengthen the lexicon.

Usually, children treat lexical sounds as ambiguous only given reason to do so. So we have to ask what would lead children to conclude that English name sounds are ambiguous? One might speculate that *not* hearing the determiner lets children know that English has lexical names. But even if children are shielded from phrases like 'the John we met yesterday', they may not use "negative" evidence in acquisition; see, e.g., Crain and Pietroski (2001). Worse, a special lexical type must be posited to let children use negative evidence to acquire a grammar that admits theoretically superfluous ambiguities. By contrast, if LPNs fetch monadic concepts, languages that regularly combine LPNs with determiners are expected. English can be viewed as a slightly quieter language in which the determiners are often though not always covert.

Thus, several considerations converge to suggest that there are no saturating LPNs, even though children plausibly lexicalize many denoting concepts with LPNs. This is unexpected *if* human I-languages permit proper nouns of type <e>, with verb-name combinations signifying saturation of predicative concepts by denoting concepts. By contrast, if semantic composition is fundamentally conjunctive, then given ways of using denoting concepts to introduce monadic concepts, it is no surprise that lexicalizing denoting concepts involves such introduction. With this in mind, let's return to cases in which the lexicalized concepts are plausibly *polyadic*.

3.3 Supradyadic Concepts and Verbs

If there are no 17-place concepts to lexicalize, then the absence of verbs with valence –17 tells us

[13] See Giannakidou and Stavrou (1999) for discussion in a broader context of Greek nominals. There is a vocative case form, used in address, that precludes determiners (and the final 's' of masculine nouns). But vocatives are plausibly used to call people, rather than fetch concepts of those called with certain LPNs. And even English has names like 'the Thames', 'the Bronx', 'the Met', 'the Mets', etc.

little. But as Gallistel (1990, 200x) discusses, it seems that animal navigation already requires polyadic representations: a forager bee can somehow relate itself to a target via the position of the sun and the time of day; an ant can somehow keep track how many steps are needed, along which vector, to relative to the hive. Humans, who can distinguish selling from giving, certainly seem to have concepts that are at least tetradic. As noted in section one, our concepts of selling and giving seem to differ, in the that former has a greater adicity: a seller gets something back in exchange for the thing transferred to the recipient. So if 'give' is alleged to have valence –3 because it lexicalized a triadic concept, one wants to know why 'sell' does not have valence -4?

We could invent a language in which (56) is a sentence with the meaning of (56a).

(56) *The colonel sold the knife the professor ten dollars.

(56a) The colonel sold the knife to the professor for ten dollars.

But in English, (56) is anomalous. And prima facie, 'sold' takes two arguments in (57).

(57) The colonel sold the knife.

Positing two further covert arguments seems *ad hoc*, especially as part of a proposal that eschews a covert constituent of names in English. Similar remarks apply to 'bought'; though (58)

(58) Professor Plum bought Miss Scarlet the knife.

is roughly synonymous with (58a), which has a "benefactive" implication, unlike (59).

(58a) Professor Plum bought the knife for Miss Scarlet.

(59) Professor Plum bought the knife for ten dollars.

Presumably, the valence of 'bought the knife' is not reduced by 'for Miss Scarlet' in (58a), else 'bought' would have valence -5 given (60).

(60) Plum bought the knife from Mustard for Scarlet for ten dollars.

But then the valence of 'bought' is not reduced by 'Miss Scarlet' in (58). This bears on (61),

(61) Plum gave Scarlet the knife.

since (58) has a closely related meaning which adds the implication that Plum paid for the gift And famously, (62) is odd while (62a) is not.

(62) *Scarlet donated the museum the art.

(62a) Scarlet donated the art to the museum.

All of which suggests that 'gave' and 'donate' do *not* fetch concepts that are saturated by correlates of three grammatical arguments. It seems to be an idiosyncratic fact that 'donate', unlike 'give', resists ditransitive constructions. But however the oddity of (62) is explained, the acceptability of (63-65) is puzzling if 'give' has valence -3.

(63) Scarlet gave the money away.

(64) Plum gave at the office.

(65) The money was given to charity.

Repeating a point from section one, the mere acceptability of (61) cannot show that 'give' has valence -3, else (22) would show that 'kick' has the same valence.

(22) Hubbard kicked her dog a bone.

And (22) is not an isolated case of "ditransitivization." We usually think of 'cooked' and 'sang' as combining with one or perhaps two arguments, as in (66) or (67).

(66) Mrs. White cooked, while Colonel Mustard sang.

(67) White cooked an egg, while Mustard sang a lullaby.

But consider (68), synonymous with (68a), and (69).

(68) White cooked Mustard an egg, while he sang the baby a lullaby.

(68a) White cooked an egg for Mustard, while he sang a lullaby to the baby.

(69) White bounced/fed/stole/left Mustard a tennis ball.

It seems that many (intuitively transitive) verbs can appear in ditransitive *constructions*, in which a "recipient" is specified, without lexicalizing concepts that have a slot for recipients.

We also seem to have triadic concepts like BETWEEN(_, _, _) that don't get expressed with verbs that take three arguments. The claim made with (70)

(70) Plum was between Scarlet and White.

is correct only if Plum, Scarlet, and White exhibit a certain triadic relation. But (70) is a subject-predicate sentence whose predicate employs three functional words to accommodate two names. This raises the question of why we don't have a verb 'twixt'—of type <e, e, e**t**>—such that (71)

(71) Plum twixted Scarlet White.

has the meaning of (70). But there is a puzzle here only if verbs *can* fetch triadic concepts.[14]

In this context, it is relevant that 'jimmied' takes *two* arguments.

(72) Mister Green jimmied the lock (with a screwdriver).

(72a) *Mister Green jimmied the lock a screwdriver.

Reference to an implement appears within a grammatically optional adjunct, as in (72), not as a mandatory third argument.[15] The concept lexicalized presumably has a slot for that with which the jimmier jimmies the jimmied. So we need some explanation why (72a) is *not* a perfectly good ditransitive construction with the meaning of (72). Similar points apply to (73/73a).

(73) Green nailed Plum (with a devastating objection/right hook/fastball/lawsuit)

(73a) *Green nailed Plum a lawsuit

One can hypothesize that "supratransitive" verbs are disallowed for reasons stemming from the underlying *syntax* of human languages; see Hale and Keyser (1993), etc. But this idea, fully compatible with a "conjunctivist" conception of semantic composition, fits ill with the idea that combining expressions often signifies saturation. If a mode of composition can be employed twice in a sentence, why not thrice or more? And if syntax precludes saturating a predicate more than twice—so that at least in this sense, saturation is not available as a *recursive* mode of composition—we have to ask whether this mode of conceptual composition is invoked by any mode of linguistic combination.

As a final illustration of argument flexibility, consider the concept of marriage. Whatever its adicity, this concept can be indicated with a noun. Yet each of (74-78) might be used to describe the same wedding.

(74) Scarlet married Plum, but their marriage was doomed.

(75) Scarlet got married to Plum, with the Reverend Green officiating.

(76) With reservations, Green married Plum and Scarlet.

(77) Plum and Scarlet married, and they got married in a hurry.

(78) It was Scarlet's first marriage, though Plum married for third time.

This suggests that given three acting participants—Scarlet, Plum, and Green—we can describe various events of marrying that fall under a monadic concept (of marriage) that a competent speaker has given the word 'marry', which lexicalizes a relational concept.

3.4 Dyadic Concepts

One might concede that when a denoting or supradyadic concept is lexicalized, the result is a

---

[14] Note that we can impose a meaning, Jabberwocky style, on (71): Plum did something to a certain Scarlet White; or perhaps Plum did something that involved twirling Scarlet to White.

[15] I am indebted to Alexander Williams for this example, and more importantly, for a series of conversations that deeply influenced the development of this section and the next. He has discussed closely related matters in Williams (2005, 2007). Note that if 'He jimmied me the lock' has a meaning, it is that he jimmied the lock *for* me, as opposed to he jimmied the lock with me, and likewise for 'He jimmied the screwdriver the lock'. Cp. note 7.

word that has a valence of –1 or –2. But this doesn't yet show that all lexical items fetch monadic or dyadic concepts. In particular, one might think that given the "event positions," lexicalizing a relational concept often results in a word that fetches a triadic concept.

For these purposes, let's distinguish "robustly" dyadic concepts like KICK(_, _) and CHASE(_, _)—concepts of genuine relations exhibited by pairs of independent entities—from concepts like AGENT(_, _) that are formally dyadic, but apply to pairs consisting of an entity and something in which that entity participates. We can think of KICK-BY-OF(_, _, _) as formally triadic but robustly dyadic, since it applies to <ε, α, β> iff ε is an event of α kicking β. If the concept lexicalized with 'kick' robustly dyadic, one might expect 'kick' to fetch a robustly dyadic concept, at least when it appears as an active voice verb. Similarly, if the valence of 'give' is capped at –2 because syntax allows only two grammatical arguments, one might expect 'give' to fetch a robustly dyadic concept. So if verbs that lexicalize robustly relational concepts do not fetch such concepts, one wants to know why.

I stress this point because many considerations press in the direction of saying that transitive constructions do not include verbs that have robustly relational meanings. To be sure, many examples are compatible with the hypothesis that some verbs have such meanings. But in my view, little if any evidence favors this hypothesis, once we take its negation seriously.

Given (23), it might seem that the lexical item 'kick' has a robustly dyadic meaning,

(23) Hubbard kicked a bone.

specified in terms of KICK(_, _) or the formally triadic concept KICK-BY-OF(_, _, _). But then (22)

(22) Hubbard kicked her dog a bone.

calls for a theoretical choice. Perhaps the verb in a ditransitive *construction* can have a robustly dyadic *meaning*, and the "indirect object" of (22) reflects a covert preposition. Or maybe the lexical item 'kick' is used, perhaps along with a covert functional item, to *introduce* a related word that has a robustly triadic meaning. In my view, each of these options is coherent.[16] But I assume that no process can change the adicity of a concept or the valence of any particular linguistic expression.

Given the requisite conceptual tools, KICK-BY-OF(_, _, _) and RECIPIENT(_, _) can be used to *introduce* a formally tetradic though robustly triadic concept KICK-BY-OF-TO(_, _, _, _) that applies to <ε, α, β, γ> iff ε is an event of α kicking β to γ. But KICK-BY-OF(_, _, _) remains formally triadic. Likewise, one can posit a process—perhaps an ordinary process of grammatical composition—in which a verb that (by itself) takes exactly two grammatical arguments is used to create an expression that takes exactly three grammatical arguments. If 'kick' has a valence of –2, perhaps it can be used to create an analytically related and homophonous expression that has a valence of –3.[17] But in any case, composition is not alchemy. If a hypothesized notion of grammatical valence is to explain why a verb can and must take a certain number of arguments (in a complete sentence), then valence cannot be a *changeable* property of the verb. If a verb V has a valence of –2, and a verb V′ appears (in a complete sentence) with one or three arguments, then V′ is not identically with V. It may be that one verb is a constructed variant of the other. But another possibility is that few if any lexical items have valences, in any explanatory sense, even if verbs often impose requirements on the sentences in which they appear.

---

[16] Though it can be hard to distinguish [[kicked (~~to~~ her dog)] (a bone)] from [[kicked+~~to~~ (her dog)] (a bone)].

[17] In terms of the 19th century chemical-compounding analogy, one might think of iron, which permits both FeO and Fe_2O_3; though if one thinks of O has having a valence of –2, one will say that Fe can have a valence of +2 or +3. So perhaps the analog of *kicked+~~to~~* is a pretended ion that has a valence –3 but remains a variant of Oxygen (as opposed to, say, Nitrogen). Or perhaps the analogy should be dropped, given modern conceptions of electron shells.

I'll return to this last point. But note that if the meaning of 'kick' is robustly dyadic, then (22) is hardly the only construction that calls for hard theoretical choices. Recall (24-26).

(24) A bone was kicked.

(25) The baby kicked.

(26) Hubbard gave the bone a swift kick.

One can say that 'kick' permits passivization as in (24), object dropping as in (25), and nominalization as in (26). Indeed, this kind of descriptive vocabulary can be useful, at least initially: by characterizing some constructions as transformation or variants of others, one can highlight commonalities across constructions.[18] But prima facie, (24-26) are grammatically complete expressions in which 'kick' does not take two arguments. I can't prove that (25) lacks a covert direct object that does not induce the implication that the baby kicked something; nor can I can prove that (25) lacks a covert indirect object that does not imply a recipient. But even if (25) has a covert direct object, as Hale and Keyser (1993) suggest, this is compatible with 'kick' fetching the monadic concept KICK(_). And since (25) does not imply that the baby kicked anything, theoretical considerations that motivate appeals to a covert *grammatical argument* do not motivate appeals to *verb meanings* that are robustly dyadic yet compatible with semantically anemic direct objects.

A competent speaker will know that if (24) is used to make a true claim, then there was an agent of the relevant bone-kicking. But such a speaker will also know that the bone-kicking occurred in a particular place at a particular time, that the agent had legs (or suitable analogs), etc. Competent speakers may well know that kickings have agents, especially if KICK(_) is introduced in the course of lexicalizing a robustly dyadic concept like KICK(_, _). But it doesn't follow that (24) and (26), generable expressions of a human I-language, have constituents that somehow represent kickers. Endlessly many properties of events, essential and contingent, can go unrepresented by (instructions for how to build) concepts that can be used to think about events.

With enough ingenuity, one can preserve the idea that 'kick' is fundamentally a label for a robustly dyadic concept, even given constructions like (24-26). But once one appeals to processes that use certain representations to *introduce* other representations that have *different* valence properties, we face the question of whether lexicalization is itself such a process. Instead of saying that children acquire a robustly dyadic verb that can be used to introduce a passive variant that does not require a grammatical subject, one can say that when children lexicalize a robustly dyadic concept, they introduce a monadic (passive-friendly) concept and acquire a word that does not have a valence of –2. From this perspective, the main theoretical questions concern human linguistic capacities, not how we choose to describe the constructions generated via those capacities. So it is worth noting the absence of logically possible words that would be expected given a labeling conception of lexicalization.

In discussing ditransitive constructions, I noted that the absence of certain verbs—and the need for certain kinds of circumlocution—bolsters the considerations in favor of denying that supradyadic concepts are lexicalized with grammatical predicates of matching valence. Similar points apply with regard to transitive constructions and dyadic verbs. It seems that prepositions often lexicalize robustly dyadic concepts like FROM(_, _) or perhaps formally triadic analogs that permit temporal modification. So it is interesting that in (79),

(79) Plum is from Devon.

the prepositional phrase combines with a copula to form a tensed monadic predicate. One can

[18] See Chomsky (1965, 1970), Larson (1988), Baker, Johnson, and Roberts (1989).

imagine a language with a corresponding transitive verb, as in (79a); cp. (80).

(79a) *Plum froms Devon.

(80) Plum hails/comes/arrived from Devon.

But it seems that 'from' cannot be used to label a concept like λY.λX.T ≡ ISFROM(X, Y) or λY.λX.λE.T ≡ COMESFROM(E, X, Y). Circumlocution is required.[19]

Similarly, we use (81), not (81a),

(81) Plum is taller than Green

(81a) *Plum talls Green

as if relational concepts cannot be lexicalized with open-class monomorphemic expressions. So perhaps robustly dyadic concepts like FROM(_, _) and TALLER(_, _) cannot be labeled/fetched with verbs, not even verbs that take two grammatical arguments.[20] Once we take this hypothesis seriously, standard accounts of causatives can be seen as more instances of the general trend.

It is a familiar idea that (82) and (83) have the indicated grammatical structures;

(82) Green broke the glass. (83) The glass broke.

(82a) [Green [[*v* broke$_i$] [_$_i$ [the glass]]]] (83a) [[the glass]$_i$ … [broke _$_i$]]

where the lexical item 'broke' appears as a constituent of both sentences. On this view, (82) is formed by combining the overt verb with a functional item *v*, which is covert in English but overt in many other languages.[21] It is independently plausible that the internal argument of 'broke' in (83) displaces to a matrix subject position without becoming the external argument of 'broke'; see Burzio (19xx, etc). So it seems that lexical items like 'break', which can appear in transitive *constructions*, do not fetch concepts like λY.λX.λE.T ≡ BREAK-BY-OF(E, X, Y). One can say that the concept lexicalized with 'break' is—perhaps modulo an event variable—a monadic concept of broken things, rather than a robustly dyadic concept. But this hypothesis, about how (all?) children acquire the word 'break', seems rather implausible for the full range of examples.

Recall (12) and (13). Does 'drop' lexicalize a concept that is *not* robustly relational?

(12) Trees dropped leaves onto the ground, as a student dropped her books.

(13) Leaves dropped onto the ground.

Perhaps we understand 'kill' (Old English, 'cwell') as the transitive form of 'die' (Old English, 'cwell'); see Traugott (1972). And perhaps we understand the verb in 'They felled trees.' as a transitive form of 'fall'. It doesn't follow that we have *no* concept of killing/felling that does not include a concept of dying/falling. But even if many causative verbs lexicalize complex concepts, *pace* Fodor and Lepore (2002), I see no reason to deny that many verbs lexicalize atomic concepts that are robustly relational.[22] Such verbs may, however, *fetch* monadic concepts.

---

[19] My own view is not that prepositions let us introduce and fetch monadic concepts of frommings. Perhaps 'from' can fetch a formally but not robustly dyadic concept FROM(_, _) that applies to <ε, β> iff ε is an event/state of β coming/being from somewhere; see Schein (2002), Svenonius (2010). But given polysemy, it may well be that for any given preposition, more than one concept can be fetched from its lexical address.

[20] One might reply that 'tall' is of this type, with relation holding between entities like Plum and abstracta like heights; cp. Kennedy (1999). But while I agree that 'tall' is used to signify a relation, in a way that invites talk of heights (or degrees thereof), circumlocution is still required, as Kennedy's work highlights.

[21] See, e.g., Chomsky (1995) and Kratzer (1996), drawing on Baker's (1988) revival of an old idea. For further references and discussion in a neo-Davidsonian context, see Pietroski (2003b).

[22] Given how little we know about lexicalized concepts, 'break' may be an example. A child might use BREAK(X, Y) or BREAK(E, X, Y) to introduce both BREAK(E) and another monadic concept BREAK(F) such that: ∃F[BREAK(F) & TERMINATER(E, F)] ≡ BREAK(E); where BREAK(F) applies to events of something breaking, regardless of the cause, and TERMINATER(E, F) is a concept of a whole-to-part relation that processes bear to their final parts. Cp. Levin and

In any case, once causative constructions like (77) are set aside, we need to ask how many of transitive constructions plausibly include a lexical verb that has a valence of –2. Since (25) does not imply that the baby kicked, one have hoped that (23) would be a paradigm case.

(25) The baby kicked.

(23) Hubbard kicked a bone.

I think the following quote—from exceptionally good psycholinguists, who have thought long and hard about lexicalization—reflects a common but mistaken view.

> Clearly, the number of noun phrases required for the grammaticality of a verb in a sentence is a function of the number of participants logically implied by the verb meaning. It takes only one to sneeze, and therefore *sneeze* is intransitive, but it takes two for a kicking act (kicker and kickee), and hence *kick* is transitive (Lidz, Gleitman, and Gleitman 200x).

The authors immediately qualify this remark, saying that "there are quirks and provisos to these systematic form-to-meaning-correspondences." But even setting (22) and (26) aside as quirks,

(22) Hubbard kicked her dog a bone.

(26) Hubbard gave the bone a swift kick.

one might wonder if (25) and (24) are plausibly accommodated with provisos.

(24) A bone was kicked.

Prima facie, (25) challenges the idea that an *act* of kicking requires *two* participants, even if one grants that paradigmatic acts/processes of kicking do involve both a kicker and a kickee. The important truth behind the quote is that the paradigmatic cases, when described with transitive constructions like (23), may well be the ones that help children figure out which concept to lexicalize with 'kick'. And this may feed our sense, as adults, that the transitive uses of 'kick' are somehow primary. But in my view, passive constructions like (24) are neither grammatical quirks nor by-products of a special process that can somehow make a transitive verb intransitive. We should, I think, be puzzled by talk of passivization (and nominalization). If the meaning of 'kick' somehow involves a variable position for kickers—and this explains why 'kicked a bone' is a mere verb phrase as opposed to a "semantically complete" sentence—then how can a word have the meaning of 'kick' but no argument position for kickers? I think that 'kick' can appear as a passive verb (or even as a noun), despite lexicalizing a *concept* that has an unsaturated slot for kickers, because lexical items need not fetch the concepts they lexicalize. Perhaps there are other explanations for why (24) and (25) can be understood as complete sentences. But I don't think the facts are *explained* by saying that verbs have valence properties whose canonical manifestations are limited to active voice uses.

If only because passive uses seem to manifest core grammatical competence, one might follow Kratzer (1996) and "sever" variables for agents from verb meanings. One can say that the lexical item 'kick' fetches $\lambda Y.\lambda E.T \equiv$ PASSIVE-KICK(E, Y); where this concept is perhaps introduced in terms of the robustly dyadic $\lambda Y.\lambda E.\lambda X.T \equiv$ KICK(E, X, Y). On this view, transitive uses of 'kick' require a second "small" verb that supports addition of an external argument, as indicated in (23a); cp. Chomsky (1995).

(23a) Hubbard [*v* [kicked a bone]]

For these purposes, we needn't worry about whether the verbs "incorporate;" cp. Baker (1988). The idea is that one way or another, [kicked a bone] and [Hubbard [*v* [kicked a bone]]

Rappaport (1995). See Pietroski (1998, 2003b, 2005a) for discussion of TERMINATER(E, X), its relation to PATIENT(E, X), and the extensive literature on these topics.

correspond respectively to the distinct *monadic* concepts λE.⊤ ≡ ∃Y[PASSIVE-KICK & BONE(Y)] and λE.⊤ ≡ AGENT(E, HUBBARD) & ∃Y[PASSIVE-KICK & BONE(Y)].

One can preserve the idea that arguments of verbs *saturate* verb meanings, by assigning to *v* a meaning of type <<e, **t**>, <e, **t**>> such that [*v* [kicked a bone] corresponds to the formally (but not robustly) dyadic concept λX.λE.T ≡ AGENT(E, X) & ∃Y[PASSIVE-KICK & BONE(Y)]. But as noted above, my view permits a variant according to which 'kick' fetches the introduced concept PASSIVE-KICK(_, _); where this concept applies to certain event-participant pairs. Then one can say that [kicked a bone] and [*v* [kicked a bone] both correspond to ∃[KICK(_, _)^BONE(_)]. One needn't assign a special meaning to *v* itself in order to say that being an external argument of a verb has thematic significance. And if *v* mediates a thematically significant relation that 'Hubbard' (but not 'a bone') bears to 'kicked' in (23a), then we can treat (23a) as an instruction for how to build a concept like ∃[AGENT(_, _)^THAT-HUBBARD(_)]^∃[KICK(_, _)^BONE(_)]. But compositional details aside, I think Kratzer's proposal is an important step in the right direction. Once variables for agents are severed from verb meanings, we need to rethink standard conceptions of how lexical meanings are related to concepts lexicalized.

In this context, recall that (84) lacks the reading indicated with (84c); see chapter zero.

(84) a spy saw a man with a telescope
(84a) A spy saw a man by using a telescope.
(84b) A spy saw a man who had a telescope.
(84c) #A spy saw a man and had a telescope.

The prepositional phrase can be understood as modifying 'saw a man' or 'man'. But in the former case, the implication is that an event of seeing was done with a telescope, and not merely that a spy was in possession of a telescope. Similar remarks apply to (85).[23]

(85) A doctor rode a horse from Texas.

If 'see' and 'ride' have robustly relational meanings like λY.λE.λX.⊤ ≡ SEE-BY-OF(E, X, Y) and λY.λE.λX.⊤ ≡ RIDE-BY-OF(E, X, Y)—with variables corresponding to the grammatical subject—the question is why we *can and evidently must* skip over this variable when interpreting the prepositional phrase. Like Kratzer, I think the verb meaning provides no such variable to skip over. One can instead posit a special rule of composition that encodes how we can understand (79): given two expressions that have meanings of the form λE′.λE.⊤ ≡ R(E, E′) and λ_.T ≡ M(_), combining the expressions creates a phrase with a meaning of the form λE′.λE.⊤ ≡ R(E, E′) & M(E); see Chung and Ladusaw (2003). But this way of describing the facts does not explain them, absent an account of why the phrasal meaning is not of the form λE′.λE.⊤ ≡ R(E, E′) & M(E′).

If there were no independent arguments for "thematic separation," and lots of evidence that supports a saturationist conception of semantic composition for verb phrases, it might be reasonable to set this issue aside for future work. But at some point, we have to ask what data actually *confirms* (as opposed to being consistent with) the claim that verbs can and often do have robustly relational meanings. If ditransitives and causatives and passives are special cases, one wonders which "core cases" motivate treating the others as somehow quirky. Reports like (86-88) are often set aside as special cases. But they may be instructive.

(86) It is snowing in London, where it often rains.
(87) Snow is falling, but at least it isn't raining frogs.

[23] If it helps, suppose that in the imagined language, this reading can be expressed with 'From Texas, a doctor rode a horse'. One might speculate that in English 'from Texas' cannot combine with 'ride the horse', but must rather attach "higher." But this seems unmotivated, at least as a strategy, given the parallel point about 'with binoculars'.

(88) A spy saw it snow in London as we watched rain fall in Venice.

Prima facie, the verbs 'snow' and 'rain' fetch event-concepts that lack a variable for that which falls when it snows/rains. But one can use a direct object to say what is falling; and if 'rained frogs' and 'rained blood' are grammatical, along with 'rained dihydrogen oxide', it is hard to see how 'rained water' could be ungrammatical as opposed to redundantly weird.

From a saturationist perspective, (86-88) are puzzling. Given the need for event variables, the verbs in 'It rained/snowed/poured/drizzled' cannot be treated as devices for expressing thoughts with no unsaturated elements, even if there are such thoughts; cp. Montague (1974). But if 'rain' expresses RAIN(E), and the argumentless verb corresponds to an argumentless concept modulo the event variable, we need some explanation for why (89) is acceptable and why it implies (90).

(89) Rocks rained down on the village.

(90) Rocks fell on the village.

An obvious initial suggestion is that RAIN(E) is an essentially plural variant of FALL(E), which is introduced via FALL(E, X), a concept that relates falls to fallen; cp. Boolos (1998). If some events satisfy RAIN(E), they are falls; if their patients were rocks that ended up on the village, they were falls of rocks that ended up on the village. And if we typically use 'rain' to think/talk about waterdrops, we might add a nominal use as in (88).[24] My aim is not, however, to provide a theory of weather reports. It is rather to highlight two points. First, if verbs express monadic concepts of things that can have participants, then "argument optionality" is not surprising. If 'rain' expresses RAIN(E), then absent lexical restrictions of the sort imposed by 'put', (86-90) do not present puzzles.

3.5 LANs, SCANs, and POSSEs

In thinking about the difficulties for any view that eschews appeal to verbal valences, it is useful to distinguish cases in which certain arguments seem to be *required* from cases in which certain arguments see to be *prohibited*. As illustrations of the first type, consider (91) and (92).

(91) *John devoured. (92) *John put.

(91a) #John devoured something. (92a) #John put the book somewhere.

As indicated, not only are the shorter sentences anomalous, they are not understood as slightly defective ways of pronouncing the meanings of (91a) and (92a); cp. 'seems sleeping/sleepy'. As illustrations of the second type, consider (93) and (94), which are quite defective,

(93) *Brutus sneezed Caesar. (94) *Brutus arrived Caesar.

but can be heard as having the meanings of (93a) and (94a); cp. 'burped the baby'.

(93a) Brutus made Caesar sneeze. (94a) Brutus caused Caesar's arrival.

And of course, these are not the only facts.

Recall that (42) does not follow from (44), which does follow from (95),

(42) John ate something.

(44) John ate.

(95) John dined.

[24] Note that 'Cats and dogs rained down on Rome' does not have the idiomatic meaning of 'It rained cats and dogs in Rome', which is roughly that it rained heavily in Rome. One might argue that 'snow' expresse SNOW(E, L), with a variable for locations. But even if this is correct, it is little comfort to saturationists. For unlike the variable for the fallen in FALL(E, X), the location variable is not saturated by the concept expressed with any argument of the verb. We can say 'Snow fell' and 'Rome fell', but not 'Rome snowed'. And if one insists that 'It snowed' has a covert *saturating location argument*, as opposed to a covert *conjoining location adjunct*, one needs appropriate analyses of (89) and (90).

even though the preposition in (96) is obligatory,

(96) John dined on shrimp.

as if the concept expressed with 'dine' does *not* have a variable for the food eaten—even though 'dine' imposes some condition on the manner of eating and/or food eaten. Any theory will have to encode, one way or another, the differences among 'ate', 'dine', and 'devour'. One option is to say that 'devour' always has a valence –2, while 'ate' can appear in *either* a (primary) transitive form that has valence –1 *or* an (object-dropping) intransitive form that has valence –1. But in my view, this sacrifices the potential explanatory force of saying that verbal valences are determined by lexicalized concepts. And as a description of the facts, I don't see why the two alleged forms of 'ate' should be viewed as "variants" rather than homophones, perhaps linked by a meaning postulate that captures the (one-way) implication from (44) to (42).

We do need to account for the anomaly of (91), since there seems to be a perfectly fine corresponding thought: John was the agent of a devouring. If (44) is an instruction for how to build a thought like ∃[AGENT(_, _)^JOHN]^ATE(_), then other things equal, (91) should be an instruction for how to build a thought like ∃[AGENT(_, _)^JOHN]^DEVOURED(_). So if 'ate' does not require overt specification of the eaten, why does 'devour' require overt specification of the devoured? This is a fine question. But the contrast between (91) and (44) *does not* support the idea that verbs have valences. If we diagnose the anomaly of (91) in terms of a "missing" internal argument, then the puzzle is why (44) is acceptable. One can say that 'ate' has a covert direct object in (44); but that seems less plausible for 'dined' ('snacked', 'noshed', etc.). Moreover, if (44) includes a covert object that satisfies the demands of 'ate', why can't (91) be understood as including a covert object that satisfies the demands of 'devoured'?

If appealing to valences was required to account for the anomaly of (91), we might have to live with the "two forms" account of why both (42) and (44) are acceptable. But here, 'put' is instructive, and in some ways a helpful example for me. Consider (97), which is also anomalous.

(97) *Chris put the book.

Appeal to valences accounts for the unacceptability here only if 'put' has a valence of -3. This is the prediction if verbs inherit their valences from concepts lexicalized, and 'put' lexicalized a robustly triadic concept (with a slot for where the agent puts the thing in question). But in (92a), 'somewhere' seems to be an adjunct rather than a grammatical argument of the verb. So to maintain a valence explanation for why (97) is unacceptable, the notion of valence must be understood so that a verb with a valence of *–n* can combine with fewer than *n* arguments to form a grammatically complete active voice construction—at least if *n* exceeds 2—so long as the "valence remainder" corresponds to a (non-optional) adjunct. One can make this hypothesis seem less *ad hoc* by claiming that human I-languages permit at most two arguments per verb. I think this plausible claim about grammar fits ill with the idea that grammatical arguments correspond to saturatable slots of concepts fetched by verbs. But in any case, if a verb can require an "associate" that is not an argument, why not say that a verb can require an associate that is not a valence-reducer?

We can say that 'put' demands a location-specifier, while neither 'eat' nor 'devour' do, and that 'devour' demands a theme-specifier in a way that 'eat' does not. In the latter respect, devour is like 'fell' and unlike 'rain'. One way or another, such lexical idiosyncracies must be encoded. Some verbs are more flexible than others with regard to how many arguments they can/must appear with in active voice constructions. But that is not yet a reason for encoding the idiosyncracies in terms of valences and the idea that a flexible verb has more than one form.

Let's say that each lexical item L has a "Semantic Composition Adicity Number"

(SCAN) corresponding to the adicity of the concept(s) fetched with L. A lexical item that always fetches a monadic concept has a SCAN of –1; a lexical item that always fetches a dyadic concept as a SCAN of –2.[25] Let the "Lexicalized Adicity Number" (LAN) of L reflect the adicity of the concept lexicalized with L: if L lexicalized a dyadic concept, its LAN is –2; if L lexicalized a triadic concept, its LAN is –3; etc. Finally, let's say that each verb also has a "Property of Smallest Sentential Entourage" (POSSE), corresponding to the number of grammatical arguments and/or mandatory adjuncts that must appear with the verb in a grammatical active voice sentence. Given this terminology, one can hypothesize that while all verbs have a SCAN of –1, POSSEs vary in part because LANs vary in ways that can affect the data available to children.

As noted above, passive constructions already suggest that SCANs do not reflect external arguments of verbs. And like weather reports, "nominalizations" suggest SCANs of –1; cp. Chomsky (1970), Marantz (1984), Borer (2005), Ramchand (2008). One can cut to the chase, expect an onsides kick, or give someone a break. So prima facie, 'chase', 'kick', and 'break' fetch monadic concepts. One can instead posit linguistic "SCAN-reducing" operations that create nouns from verbs, passives from actives, etc. But this is hardly costless. And to repeat an earlier point, if a representation can be used to a representation that has a different valence, we have to consider the possibility that lexicalization is itself a reformatting process of this sort.

Still, even if lexicalized concepts are not fetched for purposes of semantic composition, they may leave "marks" on the resulting verbs. In particular, if 'put' lexicalizes a triadic concept, the verb may record this fact. And if for whatever reason, the introduced concept PUT(_) is especially *un*restrictive—one can put on a happy face, put up with a colleague, put someone on, etc.—this may have lead speakers to regularly use 'put' with locative adjuncts. Over a few generations, this may have led children to conclude that while 'put' fetches a monadic concept of events, it *also* imposes a constraint on complex expressions in which the verb appears: such expressions should represent something partly in terms of how its location changed. (Recall that phrases like 'in an hour' can only modify "telic" predicates that represent endpoints of events.) The actual facts are surely more complicated. But one can certainly imagine that details concerning lexicalized concepts lead to patterns in usage that lead children to encode certain lexical restrictions on sentential expressions.

Relatedly, let's distinguish at least two ways in which a string of words might fail to be (the string determined by) a grammatical expression. The contrast between (98) and (99)

(98) *Chris been have might there

(99) Chris might have been there

is apparently due to the limited ways in which expressions of a human I-language can be generated; see Chomsky (1957). The string 'been might have there' is not the "linearization" of any expression of (an I-language that is a version of) English. No procedure for assembling such an expression—forming 'have might there' and combining with 'been', forming 'been have' and 'might there' and combining, etc.—is a human way of generating expressions.[26] By contrast, while (100) is deviant for adult speakers of English, it is as generable as (100a).

(100) *Chris goed to the store

(100a) Chris went to the store.

---

[25] A lexical item that always fetches a singular concept has a SCAN of +1. A lexical item that sometimes fetches a monadic concept and otherwise fetches a dyadic concept has a SCAN of {–1, –2}.

[26] Recall that expressions are, or at least determine, articulation-meaning pairs. Obviously, a grammar can generate (an expression that determines) a word-string that sounds like 'might been have there' but transliterates the English expression 'is easy to please'. But that is irrelevant.

The deviance of (90) lies with "irregularity" of a particular verb. Competent speakers know that the past tense of 'go' is 'went'. In this sense, 'goed to the store' is generable but "filtered out."

Recall that (101) has the meaning of (101a) and not (101b), the deviance of (101) differs from that of (98).

(101) *The child seems sleeping
(101a) The child seems to be sleeping
(101b) #The child seems sleepy

A competent speaker of English can somehow pair the articulation of (101) with the meaning of (101a). The deviance of (102) can make it hard to hear the meaning indicated with (102a).

(102) *Which linguist did the philosopher think that went home?
(102a) Which linguist did the philosopher think went home?

But one might think that (102) is a generable violation of a filter that precludes (pronunication of) the complementizer 'that' when it immediately precedes a trace.[27] I don't want to take a hard stand on any particular case. My point is just that we can retain the traditional idea that a full characterization of (un)grammaticality will require appeal to some idiosyncratic constraints, at least for cases like (100), in addition to an intrinsically constrained generative procedure. Ungenerability may be a "deeper" sort of ungrammaticality. But just as we've grown used to the idea that there are many sources of unacceptability, and often more than one way for a grammar to generate (an expression that determines) a certain string, we can and should get used to the idea that a grammar may fail to license a string that is determined by a generable expression.

In thinking about these issues, it is crucial that one not forget the reasons for adopting an I-language perspective on human languages. It is easy to forget that an I-language, consisting of a lexicon and a grammar, generates *structured expressions* that determine strings (relative to certain assumptions about how these expressions are linearized by relevant articulatory systems). It is easy to slide into thinking of languages as directly generating strings that either do or don't meet "well-formedness" conditions of the sort specified when we specify familiar invented languages in the usual way. From this perspective (97) and (103) and are not expressions,

(97) *Chris put the book.
(103) *Chris sent.

and hence not meaningful expressions. But another option is to say (97) and (103) are generable but defective expressions expressions: 'send' and 'put' are instructions to access concepts like SEND(_) and PUT(_): but 'send' also calls for specification of a patient; and 'put' calls for specification of a location. Likewise, one can say that (91) and (92) generable but defective.

(91) *John devoured.
(92) *John put.

As Levin and Rappaport (1995, 2005) show, *merely* adopting a version of this second option would miss generalizations that we want to capture. But these may be generalizations about human concepts and how children tend to lexicalize them; and these facts can be described in terms of the variable POSSEs that children associate with lexical items, instead of assuming that the operations of semantic composition permit significant variation in LANs.

4. Number Neutrality

I have focused on adicity as an important dimension in which lexicalized concepts can differ from—and exhibit more typological variation than—introduced concepts that can be joined in simple ways. But I suspect that lexicalized concepts exhibit other typological distinctions that lexically introduced concepts do not, making the latter more systematically combinable than the

[27] Note to Lasnik and ellipsis repair of island violations.

former. In particular, the concepts that children lexicalize may include both "mass concepts" and "count concepts;" where many of the latter are "singular," in the sense of being applicable to a single thing, while others are "essentially plural." But in my view, introduced analogs of these concepts are "number neutral," and this neutrality—between mass and count, singular and plural—is related to the limitations on how human languages employ second-order quantification. Here, I will be applying and extending Boolos' (1998) treatment of second-order quantification in an invented *monadic* calculus as *plural* quantification. The details will carry over into chapter six, which shows how phrases like 'most of the cows that the dogs chased' can be analyzed as instructions for how to build (second-order/plural) monadic concepts. I want to end this chapter by laying the ground for such analyses, while also showing that one need not focus on adicity to suspect that lexicalization is a process of using typologically diverse concepts to introduce *less* diverse concepts that can be systematically M-joined or Θ-joined. The idea will be that lexically introduced concepts are the atomic expressions of a mental language, SMPL, that is a second-order extension of a fragment of PL, the Tarskian language discussed in chapter two. And since SMPL only generates complex *monadic* concepts (a.k.a. mental predicates), it is far less powerful than CPL, Church's lambda-extension of PL. But it is powerful enough to make it plausible that human linguistic meanings are instructions for how to build SMPL concepts.

4.1 Typology and Flexibility Again

I assume that many lexicalizable concepts are "mass concepts" in a familiar and much discussed sense. To think about some wood *as* wood is to think about it as some *stuff* of a certain sort, and thereby think about it in a way that does *not* involve thinking about it as (or even as constituted by things that are) *countable*. In thinking about anything as wood, we think about it in a way that lets us wonder *how much* of it there is in a given situation. If only for simplicity, I will also assume that in thinking about wood as such, we use a monadic/predicative concept as opposed to a relational concept. One can think that all wood is good, and recognize the implication that all brown wood is good; cp. §2.2 of chapter one. So I will speak of the mass-concept WOOD(_).[28]

Other lexicalizable concepts are presumably "count concepts." To think about a tree or some trees as such is to think about one or more things in a certain way—e.g., as a tree or as some trees—that lets us wonder *how many* of them there are in a given situation. At least many of these concepts are singular, or individualizing, in the following sense: if such a concept applies to anything, it applies to *each one* of the one or more things it applies to. The concept TREE(_) applies to many things, but only in the distributive sense of applying to each of many trees. Each tree may comprise many things. But if TREE(_) applies to any such things, taken together, that is because they constitute a tree; and while it may turn out that some trees are parts of other trees, TREE(_) applies to some tree-parts only if each of them is itself a tree. By contrast, the complex concept THREE-TREES(_) seems to be "essentially plural" in the sense of being potentially applicable to some things (taken together) but not to any single thing. Some atomic concepts may also be essentially plural; indeed, THREE(_) may be a such a concept.

A certain kind of theorist will say that THREE(_) is really a singular concept that applies distributively to certain "plural entities"—e.g., *sets* that have three elements. Likewise, one can hypothesize that TREES(_) and THREE-TREES(_) and even WOOD(_) are singular concepts that

[28] Our concept of good(ness) may underlying relational; cp. §4.2 of chapter three. But our concept of wood seems to be classificatory. Though talk of a single concept, WOOD(_), is idealization; see §1.2 of chapter one. In particular, a person who knows some science may have more than one concept of wood (water, fire, etc.). But I assume that we have a concept of wood that makes it possible to entertain the (false but coherent) thought that wood is basic stuff, not composed of anything else, and even that there is no smallest portion of wood. ADD REFERENCES.

apply to plural entities. I don't deny that we can form concepts like SET-OF-TREES(_) and SET-WITH-THREE-ELEMENTS(_) and SET-OF-WOOD-PARTICLES(_). For reasons discussed below, I think we also form essentially plural concepts that *are not* singular concepts that apply to plural entities; and I think that many expressions, including 'three trees', are instructions for how to build such concepts. But for the moment, it doesn't matter whether we distinguish singular concepts from plural concepts (and distinguish all of these from mass concepts), as opposed drawing distinctions among singular concepts. The more important point is that just as lexicalizable concepts seem to differ typologically in ways that we can describe with talk of adicity, they seem to differ along other dimensions, including mass/count and singular/plural; where these distinctions intended to describe kinds of *concepts*, or *ways of thinking about* things, not kinds of things that might be thought about in a uniform way; see §1.1 of chapter one.

I am not assuming that the world contains stuff that is not constituted by countable things. I am also not assuming that the world doesn't contain any such stuff. (I am focused on meaning, not ontology.) But I do assume that humans and other animals often think about stuff-or-things *as* countable things—or at least as things whose numerosity can be estimated—and that we often think about stuff-or-things *without* thinking about it-or-them as countable. When we think about stuff-or-things *as* countable, we sometimes do so via concepts that are essentially plural, and we often use concepts that are singular/individualizing/distributive.

However this familiar typology is described, the conceptual distinctions seem pretty robust. So as many authors have discussed, it is striking that linguistic expressions seem to be rather flexible in this respect. Examples like (104-106) abound.

(104) The red rock that is scattered in the yard was inexpensive.
(105) It came from a single red rock.
(106) There are more red rocks over there.

Similarly, a small portion of fish may have come from a fish that was among three fish caught this morning. One can say that 'rock' appears as a "mass noun" in (104) and as a pluralizable "count noun" in (105) and (106). But talk of mass/count/plural nouns may layers of grammatical structure, as opposed to typologically distinct kinds of lexical items; see, e.g., Gillon (1987), Borer (2005). I'll return to the idea that just as 'rocks' is complex expression, not a lexical item, the count noun 'rock' is a kind of noun *phrase* that combines a lexical *root*, with a functional element. Unsurprisingly, I like this idea; the functional element can be viewed as an instruction to conjoin a concept like ONE-OR-MORE(_) with a neutral concept, √ROCK(_), which can be introduced and then fetched via the lexical item 'rock'. But however we encode the distinction between the uses of 'rock' in (104) and (105), it seems clear that 'rocks' in (96) has the lexical item as a constituent. So let's focus on rocks before turning to rock.

If only for simplicity of exposition, let's assume that in acquiring 'rock', a child lexicalizes a concept ONE-ROCK(_) with which the child can think about a single rock as such. This leaves room for various hypotheses about (i) how the lexicalized concept is related to the concept(s) that can be fetched via the lexical item, and (ii) the role of the plurality marker. But given a Frege-Church conception of composition, an obvious first thought is that 'rock' fetches $\lambda X.T \equiv$ ONE-ROCK(X)—a concept of type <e, **t**>—and the plurality marker fetches a concept of type <<e, **t**>, <e, **t**>> like $\lambda X.\lambda P.T \equiv$ PLURALITY(P) & $\forall X$:ELEMENTOF(X, P)[$X_X = T$]; where PLURALITY(_) is a concept of plural entities, and the lambda concepts may have to be introduced, rather than simply labeled by infants who already function-denoters independent of pressure from lexicalization. Given the operation of function application, as discussed in chapter two, 'rocks' then be viewed as an instruction for a how to build a complex concept of plural entities—

perhaps collections or lattices—whose basic elements are countable rocks. This is an overtly mentalistic version of a familiar kind of account that is often presented in truth-theoretic terms; see Link (1983), Schwartzschild (1996), Chierchia (1998).

I don't deny that infants have concepts like PLURALITY(_), which can be used to form concepts of pluralities of rocks. But it doesn't follow that infants acquire a corresponding concept of <<e, **t**>, <e, **t**>> and use the plurality marker to fetch it. Another possibility is that 'rocks' is an instruction for how to build a concept like #ROCK(_)^PLURAL(_); where PLURAL(_) is an essentially plural concept that applies (nondistributively) to some things iff there are two or more of them, and #ROCK(_) is a number-neutral concept that can apply to a single rock or to some rocks (however numerous). Put another way, #ROCK(_) can be introduced so that it applies to *one or more things* iff ONE-ROCK(_) applies to each of them; cp. the number-neutral concept EACH-A-ROCK(_). I spell out this idea in section one of chapter six, drawing on Boolos (1998). But if it helps, think of #ROCK(_) as akin to the concept ROCK-OR-ROCK-COLLECTION(_), which applies to entity e iff e is either a rock or a collection of rocks. Similarly, #ROCK(_)^PLURAL(_) is akin to ROCK-OR-ROCK-COLLECTION(_)^PLURALITY(_), which applies to e iff e is a collection of rocks. However, kin differ. The concept #ROCK(_)^PLURAL(_) applies to rocks, not to collections.

I think we use 'rocks' to form concepts of rocks, not concepts of collections. Like Boolos, I don't think (107) *implies* that there is a collection of the rocks in question,

(107) Theo saw the rocks.

much less that Theo saw any collection. There may *be* a collection of the relevant rocks, which may have been widely scattered, with each seen only by chance. There may also be a collection of those rocks and Theo, a collection of them and the event(s) of Theo seeing the rocks, etc. But I don't think that (107)—or any thought built in accord with this instruction—implies that there is a collection of anything; though (107) does imply that there is at least one rock that Theo saw. Similarly, (108) can be used to make an assertion that implies the existence of various things,

(108) The red rocks that are scattered in the yard fell from the sky.
Yesterday, they rained down and formed a circle around a cactus.

but not a collection of rocks that rained down, formed a circle, and is scattered. And (109)

(109) Some logician iteratively specified the Zermelo-Frankl sets.

does not imply there was an iterative specification of a collection whose elements are the ZF sets.

One can imagine minds that use concepts like PLURALITY(_) to introduce higher-order concepts like $\lambda X.\lambda P.T \equiv \text{PLURALITY}(P) \,\&\, \forall x{:}\text{ELEMENTOF}(x, P)[Xx = T]$, or perhaps use concepts like ONE-ROCK(_) to introduce concepts that apply to collections. But ascending the Fregean hierarchy types, or introducing concepts of abstractions, may be more of a logician's move than a natural strategy for human children. And there is an alternative: given concepts that cannot be coherently M-joined or Θ-joined, try to introduce concepts that can be so combined. The concept ONE-ROCK(_) applies to each rock, but PLURAL(_) does not. Simply M-joining these concepts, as if PLURAL(_) were like RED(_), would be incoherent. (I return to the mass-reading of 'red rock'.) But if children have singular and plural concepts, independent of the I-languages they acquire, perhaps lexicalization is a process in which they acquire neutral concepts like #ROCK(_). This concept can be coherently M-joined with PLURAL(_) to form a concept of rocks, not collections.

By hypothesis, a nonmonadic concept cannot be M-joined with anything. But the monadic concepts available to children, prior to lexicalization, may also exhibit typological distinctions that preclude systematic M-junction. There may be various ways of introducing more neutral concepts. Though one obvious (and perhaps biologically workable) strategy is to introduce concepts that are effectively disjunctive, taking advantage of the fact that predicate

reduction is a natural form of inference. If a child can use ONE-ROCK(_) to introduce #ROCK(_), given some ancillary resources, then it might be easy to form concepts like #ROCK(_)^PLURAL(_) and #ROCK(_)^ONE(_); where is ONE(_) effectively equivalent to ~PLURAL(_), and so #ROCK(_)^ONE(_) is effectively equivalent to ONE-ROCK(_). Correlatively, the resources required to introduce #ROCK(_) are relatively modest.

Suppose that a child acquiring an I-language has a concept INCLUDE(_, _) that can be used to think that one or more independently specified things include an independently specified thing. Thinking that Bessie is a cow is not very different than thinking that the one or more cows include Bessie, or that Bessie is among the one or more cows. Indeed, INCLUDE(_, _) can be viewed as the inverse of ONE-OF(_, _); each cow is one of the cows, each of which is a cow. If it helps, think of ONE-OF(_, _) as akin to ELEMENT-OF(_, _). But my suggestion is that when children are introducing fetchable concepts, in the course lexicalization, they can use concepts like ∀:INCLUDE(_, _)[ONE-ROCK(_)]; where this complex monadic concept applies to one or more things iff each of them is a rock. For some purposes, this might count as a variant of the idea that children can introduce the concept λX.λ*P*.⊤ ≡ PLURALITY(*P*) & ∀X:ELEMENTOF(X, *P*)[XX = ⊤], fetch it via the plural marker, and saturate it with the concept λX.⊤ ≡ ONE-ROCK(X). But here, the issues concern the range of human linguistic types and the nature of composition operations.

Recall that Θ-junction permits construction of ∃[UNDER(_, _)^ONE-ROCK(_)], which applies to things that are under a rock. This monadic concept is akin to the Tarskian open sentence ∃X[UNDER(X′, X) & ONE-ROCK(X)]. But if the "slots" in ∃[UNDER(_, _)^ONE-ROCK(_)] are number neutral, and ∃[FORM(_, _)^ONE-CIRCLE(_)] can apply to some things that together form a circle, then Tarskian sentences are both too sophisticated and too limited as models for these Θ-junctions: too sophisticated because they impose no limit on the number of open variables in a conjunct; too limited, because each variable has *exactly one* value relative to each assignment of values to variables. And I don't want to say that ∃X[FORM(X′, X) & ONE-CIRCLE(X)] is a concept of sets. Likewise, I am positing concepts like ∀:INCLUDE(_, _)[ONE-ROCK(_)] instead of positing concepts like ∀X:INCLUDE(X′, X)[ONE-ROCK(X)] and saying that they true of certain sets that are extensions of functions of type <e, **t**>.

4.2 Two Ways of Interpreting Capitalized Variables

Luckily, there is a better model. The trick is to allow for quantification into predicate positions without thereby quantifying over anything new.

One can invent a Frege-Church language that generates (110-112),

(110) Pa

(111) ∃x[Px & (x = a)]

(112) ∃X[Xa & ∀x(Xx ≡ Px)]

with each of these sentences being equivalent to the others. If 'a' denotes an entity, and 'P' denotes a function from entities to truth values, then (110-112) are mutual paraphrases. One can also invent a variant of the Tarskian language PL that generates the same strings, but still has no denoters, letting (112) be shorthand for (113). Then (114) can be a theorem while (115) is not.

(113) ∃s:Set(s)[a ∈ s & ∀x(x ∈ s ≡ Px)]

(114) Pa ≡ ∃x[Px & (x = a)]

(115) Pa ≡ ∃X[Xa & ∀x(Xx ≡ Px)]

Even if (115) is true, its right side has implications like (115a) and (115b);

(115a) ∃s:Set(s)[a ∈ s]

(115b) ∃s:Set(s)[∀x(x ∈ s ≡ Px)]

and one can stipulate that these are not implications of (110) or (111) in the variant of PL.

In the second kind of language, the "second-order" quantification in (102) and (105) is disguised first-order quantification over sets. Quine (196x) and others have asserted that eschewing second-order quantification in favor of explicit first-order quantification over sets—or perhaps merelogical sums, collections, lattices, or other "potentially plural" entities that can somehow "include" more basic entities—is virtuous in some way that matters for science. I am not persuaded. But whatever one thinks about the languages that mathematicians and physicists (along with linguists and philosophers) *should* use, for scientific purposes, there is no reason for thinking that human children have been acquiring such languages for millennia.

Humans, perhaps along with some other animals, may enjoy one or more languages of thought that permit quantification into "predicate positions." This is not to say that children can formulate mental analogs of (112), which already includes an ampersand along with a first-order quantifier that can appear in sentences like (116).

(116) $\forall x(Rxx' \ \& \ Sx'x''x''')$

But children might be able to formulate mental predicates that are more like (117),

(117) $\exists$[AGENT(_, _)^DOGS(_)]^CHASE(_)^PAST(_)^$\exists$[PATIENT(_, _)^CATS(_)]

which is a concept of past chases whose agents are some dogs, and whose patients are some cats. If the unsaturated slots are like Tarski-style first-order variables, then any assignment of values to the slots would assign *exactly one entity* to each slot, perhaps allowing sets as special entities. If the unsaturated slots are like Church-style second-order variables, then any assignment of values to the slots would assign *exactly one function* to each slot. For many purposes, including mine, that is an unimportant contrast. But a third interpretation for (112)

(112) $\exists X[Xa \ \& \ \forall x(Xx \equiv Px)]$

provides a helpful model for the slots in (117).

Boolos offers an interpretation according to which second-order variables are *formally distinctive* while having values of the *same sort* as first-order variables. Instead of treating a capitalized variable like 'X' as a disguised first-order variable, ranging over sets or other abstracta that bear a special relation to values of 'x', Boolos treats each capitalized variable as a number-neutral variable that can have *one or more* values relative to each assignment of value*s* to variables. Each lower case variable is still construed as a singular variable that always has *exactly one* value relative to each assignment of value*s* to variables.

On this interpretation, 'Xx' indicates that *the* value of 'x' is *one of* the *one or more* values of 'X'. Or put the other way, the one or more values of 'X' *include* the one value of 'x'. For any two or more things (the Xs), each thing (x) bears the "one of" relation to them iff it is (identical with) one of them. Each thing also bears this relation to itself, but to no other single thing: if the one or more Xs are just one, then trivially, that lone X is still *one of* the one or more Xs; it is the X. On this view, to say that Aristotle is *one of* the one or more philosophers is not to say Aristotle is an element of a certain set, or that he is mapped to a truth value by a certain function. The claim is simply that there are one or more philosophers, and Aristotle is one of them. Likewise, '$\forall x(Xx \equiv Px)$' indicates that each domain entity is one of the Xs iff it is one of the Ps. If these values of 'P' are the philosophers, and 'a' indicates Aristotle, then (102) says that one or more things are such that Aristotle is one of them, and each thing is one of them iff it is a philosopher. Which is just to say, elaborately, that Aristotle is a philosopher.

In English, one says 'one or more things *are*', not 'one or more things *is*'. But nonsingular *agreement* does not imply more than one thing. Consider (118) and (119).

(118) There are one or two philosophers here.

(119) There are *zero/between zero and two/1.0/no* philosophers here.

And while (120) implies that the room contains more than one philosopher, (121) does not.

(120) There are philosophers in the room, and he is one.

(121) There are one or more philosophers in the room, and he is one.

Grammatical agreement cannot keep us from using 'one or more' number neutrally. And a tradition of set-theoretic construal cannot force us to interpret 'X' as having exactly one value relative to each assignment of values to variables. We can interpret 'X' as a variable that can have many values, none of which need to be sets.

This is to deny that an assignment must be characterizable as a function that maps each variable to exactly one entity. And this is to make the notion of an assignment more general than Tarski's original notion. But there is nothing mysterious about the more general notion. On the contrary, the apparatus usually invoked to make the set-theoretic construal explicit can be used to make the Boolos-construal equally explicit. For illustration, assume a domain of exactly 4 things. This yields 15 potential assignments, which can be depicted in either of the ways shown below.

```
              abcd                          0001  0010  0011
      abc   abd   acd   bcd          0100   0101  0110  0111
  ab    ac   ad    bd    bc    cd    1000   1001  1010  1011
a           b           c           d       1100  1101  1110  1111
```

The representation on the left can be used to illustrate the set-theoretic construal: the bottom nodes of the "lattice" indicate "basic" elements of an extended domain; every other node indicates a non-primitive entity that can be *the* value of a second-order variable. Then 'abd' signifies the set {a, b, d}, or the "mereological sum" a⊕b⊕d. But instead of introducing an extended domain, one might take 'abd' to indicate a *plural* assignment of three *values* to a second-order variable. Recoding in base two—a = 1, b = 10, c = 100, d = 1000—makes this explicit. For while '1011' can be viewed as a label for a thing that has three elements, the binary code can also be read as providing answers to yes/no questions. For each entity—d, c, b, and a—we can ask whether it is one of the potentially many values assigned to the variable, with '1011' providing the following answers: yes, no, yes, and yes.

Given this apparent equivalence, one might wonder if the plural interpretation of 'X' differs at all from the set-theoretic interpretation. But consider the infamous formula (122),

(122) $\exists X \forall x[Xx \equiv (x \notin x)]$

with '$x \notin x$' meaning that the value of 'x' is not selfelemental. And let's stipulate a domain that is *exhausted* by two things—say, Aristotle and Frege—neither of which is an element of itself.[29] Given such a domain, (122) is true on the Boolos interpretation: there are some things such that for each thing, it is one of them iff it is not selfelemental. But since nothing in the stipulated

[29] Chomsky can be the only linguist, and so one of the one or more linguists, without being selfelemental. Insisting that each thing is a "part of" itself is, I think, to interpret 'part' (misleadingly) as 'one'. A related point: competent speakers can understand and assent to the sentence 'There are many sets, each of which is a set, but there is no set whose elements are all and only the sets'. This suggests that speakers need *not* assume that each noun corresponds to the set of things satisfied by that noun. Likewise, speakers can understand and assent to 'No set contains all and only the sets that do not contain themselves', just as they can understand and assent to 'No barber shaves all and only the barbers who do not shave themselves'.

domain has each thing as an element, (122) is false on the more standard interpretation. Given a *different* domain that includes {Aristotle, Frege} along with its two elements, (122) can be true on both interpretations. But the interpretations still differ. For there are many domains relative to which (122) is true on the Boolos construal but false on the set-theoretic construal. Consider the Zermelo-Frankl (ZF) sets. Relative to this familiar domain: there are some things—namely, the ZF sets—such that each thing (i.e., each ZF set) is one of them iff it is not selfelemental; yet there is no set whose elements include all and only the nonselfelemental things. So even if *other* domains include a "collection" of all and only the ZF sets, at least one perfectly standard domain makes (122) true on one interpretation but false on the other. So the interpretations differ.

Correlatively, (123) is trivial on the Boolos interpretation.

(123) $\exists x(Fx) \equiv \exists X \forall x(Xx \equiv Fx)$

Given an F, there are one or more Fs—i.e., one or more things are such that each thing is one of them iff it is an F—and given one or more Fs, there is an F. This does not imply that the Fs form a set, although they may well. Likewise, (124) is trivial.

(124) $\forall X{\sim}\exists x(Xx \,\&\, {\sim}Xx)$

For any (choice of) one or more things, nothing is both one of them and not one of them. Or put another way, replacing '∀X' with '~∃X~' and eliminating the double negation:

(124a) ${\sim}\exists X \exists x(Xx \,\&\, {\sim}Xx)$

it's false that one or more things are such that something is both one of them and not one of them. These formulations of monadic second-order noncontradiction are not claims about sets. If it helps, imagine replacing the singular variable 'x' with '☝', and the number-neutral 'X' with '✋', as in (124b).

(124b) ~∃✋∃☝[OneOf(☝, ✋) ≡ ~OneOf(☝,✋)]

By contrast, (125/125a) is true, letting '$\underline{s}$' be a special variable ranging over sets.

(125) $\forall \underline{s}{\sim}\exists x[(x \in \underline{s}) \,\&\, {\sim}(x \in \underline{s})]$

(125a) ${\sim}\exists \underline{s} \exists x[(x \in \underline{s}) \,\&\, {\sim}(x \in \underline{s})]$

But this seems more like a special case of noncontradiction: for any set, no thing both is and isn't an element of it. In this sense, (125) seems like (126).

(126) $\forall y{\sim}\exists x[In(x, y) \,\&\, {\sim}In(x, y)]$

As Boolos notes, his interpretation does not extend to second-order *relations* that resist characterization as predicates of n-tuples. But instances of (127) can be glossed as in (127a).

(127) ∃R...Rxy…

(127a) one or more ordered pairs, the Rs, are such that…they…

And even with just a (very) finite number of dyadic relations, including some thematic relations, this yields a lot of expressive power—perhaps enough, or even more than enough, to adequately model the concepts that can be constructed by executing the assembly instructions that human linguistic expressions connect with pronunciations.

One can hypothesize that each human language is a fragment of a language like PL or CPL. Though if one suspects that sentences of a human language do not themselves have truth conditions, one might say that meanings are instructions for how to build expressions of a mental language that is a fragment of like PL or CPL. My suggestion is that meanings are instructions for how to build expressions of a mental language—call it SMPL—that is more-second order than PL, but massively monadic. While SMPL allows for finitely many atomic dyadic concepts, every complex concept of SMPL is monadic; although many atomic concepts of SMPL may be *introduced* by using available polyadic concepts. And while SMPL is like CPL in allowing for quantification into predicate positions, as in the Θ-junction ∃[AGENT(_, _)^DOGS(_)],

the concepts generated by SMPL are predicative, not denoters; although some atomic concepts of SMPL may be *introduced* by using available denoting concepts.

Moreover, the monadic concepts generated by SMPL are not instances of the Fregean type <e, t>. Concepts of the SMPL type <M> can be described from a Fregean perspective; but they do not themselves relate entities to truth values. Likewise, ∃[AGENT(_, _)^DOGS(_)] is not an instance of the Fregean type <e, **t**>; AGENT(_, _)^DOGS(_) is not an instance of the Fregean type <e, <e, **t**>>; and ∃ is not an instance of the Fregean type <<e, <e, **t**>>, <e, **t**>>. There is an operation of Θ-junction that permits combination of atomic dyadic concept with a monadic concept; and since the combined concepts may be plural, the kind of existential closure involved is better viewed as number neutral, *a la* Boolos, than as a Church-style second-order quantifier (or a Tarski-style first-order quantifier) whose variables range over sets of some kind.

If the slots in thematic concepts are characterized number-neutrally, we can say that AGENT(_ ,_) applies to some entities and some events iff the latter are (all and only) *the agents* of the former; cp. the discussion of thematic "exhaustion" in §3.2 of chapter three. Following Schein (1993, 2002, forthcoming), I assume that 'branded fifty cows with twenty irons'—which can combine with 'Five ranchers'—can be understood distributively, implying that each brander branded fifty cows, or collectively. The collective reading, which implies only fifty brandings, implies nothing about how the cows and irons were distributed across those who got the work done. And this neutrality is captured if each thematic role has its own number-neutral conjunct, like ∃[AGENT(_, _)^[FIVE(_)^RANCHERS(_)]], according to which five ranchers did the work.

Likewise (128) exhibits a "semi-distributive" reading:

(128) Three trainers taught five dogs four tricks yesterday

which three trainers together taught four tricks to each of five dogs, for a total of twenty processes of teaching a dog a trick—perhaps with more than one trainer involved in a given trick-teaching, but with no implication (or prohibition) of cooperation. As Schein (1993, 2002) argues, such facts tell against analyses in terms of polyadic concepts like TEACH(E, X″, X′, X′) that are satisfied by *n*-tuples of *collections* of events/trainers/dogs/tricks. Though even setting such examples aside, relative clauses like 'who branded fifty cows with twenty irons yesterday' apparently correspond to concepts that can be satisfied by one individual who did a lot of work, or by several individuals who together did the requisite branding (with or without cooperation).

Even if 'dogs chase cats' can be analyzed as instruction for how build (117),

(117) ∃[AGENT(_, _)^DOGS(_)]^CHASE(_)^PAST(_)^∃[PATIENT(_, _)^CATS(_)]

that is still a long way from making it plausible that meanings are more generally instructions for how to build SMPL concepts. But idea is to figure out what must be *added* to M-junction and Θ-junction, given a suitable stock of *potentially plural* atomic monadic and dyadic concepts, to generate the meanings of human I-langauge expressions—instead of starting with a far more powerful model like CPL, and asking what must be *subtracted* to not overgenerate meanings.

4.3 Mass without Church

If children can form concepts like #ROCK(_), which can apply to one thing or to many things, that may help with (105) and (106), but it doesn't yet accommodate the use of 'rock' in (104).

(104) The red rock that is scattered in the yard was inexpensive.

(105) It came from a single red rock.

(106) There are more red rocks over there.

Perhaps the lexical item 'rock' is used to fetch *either* #ROCK(_) *or* a concept of certain sets whose elements are parts of rocks. But if so, then a certain simplicity of SMPL is lost. If ONE-ROCK(_) can be used to introduce SET-OF-ROCK-PARTS(_), which can then be fetched and

conjoined with a concept that can be fetched via 'red', then it would seem that SET-OF-ROCKS(_) can be introduced, fetched, and conjoined with a concept that can be fetched via 's'. In which case, SMPL starts looking like a fragment of CPL. Of course, the concepts we build by executing meanings are what they are. But perhaps Boolos's strategy can be extended to allow for predicates and quantifiers that are mass/count neutral.

Let's return the idea that the lexical item 'rock' is a constituent of the count noun 'rock', which is a constituent of the plural noun 'rocks'. If the corresponding concept ROCKS(_) is the complex concept #ROCK(_)^PLURAL(_)—with #ROCK(_) as an introduced concept that is neutral enough to be coherently conjoined with either PLURAL(_) or ONE(_)—perhaps #ROCK(_) is the complex concept √ROCK(_)^ONE-OR-MORE(_); where √ROCK(_) is an introduced concept that can but need not be conjoined/restricted with a concept that is used to think of stuff-or-things *as* one or more things. But let me stress at once: the idea is *not* to analyze ONE-ROCK(_) in terms of a more basic "root concept" that applies to a certain kind of stuff as opposed to countable things.

I am *not* suggesting that all count concepts are complex concepts that have concepts of stuff as constituents. In my view, thinking of a cow as such is *not* a matter of thinking of something as a certain hunk of beef with some interesting properties not exhibited by beef on the plate. On the contrary, thinking of some beef as such is more likely a matter of thinking about it as meat from one or more cows. Likewise, thinking about a pound of rock, fish, or chicken as such seems to be a matter of thinking about it as a pound of stuff from rocks, fishes, or chickens. We need to leave ample room for the idea that √ROCK(_) is introduced in terms of ONE-ROCK(_), √FISH(_) is introduced in terms of ONE-FISH(_), etc. In my view, a singular count noun can have a lexical root that is used to fetch a number-neutral concept that is introduced via some concept that is equivalent to the concept constructed via the count noun. If 'a fish' is the structured expression [[fish+CT]–PL]—where 'CT' and 'PL' are grammatical features of count and plurality—then the atomic expression 'fish' can used a root concept √FISH(_) that can be conjoined with ONE(_)—or with ~PLURAL(_) and ONE-OR-MORE(_)—to form a concept that is equivalent to an atomic concept, ONE-FISH(_), in terms of which √FISH(_) was introduced.[30]

Talk of a "universal grinder" (see Pelletier 198x) can find its place here, without the usual Fregean apparatus and appeals to shiftable types. I think √FISH(_) can be introduced as a concept that applies to some stuff-or-thing(s) iff ONE-FISH(_) applies to all of them *or* all of it is the result of "grinding" things that ONE-FISH(_) applies to. I'll try to make this idea a bit more precise below. But the details are less important than recognizing that there are cases and cases. To take the usual contrasting example, thinking of some water as such is presumably not a matter of thinking about it as stuff from a water. Our concept of *a water* is presumably introduced via some mass concept of water and some "unitizing" concept.

Given a mass concept WATER(_) that we can use to think about water as such, we can form concepts like SUITABLE-PORTION-OF-WATER(_). I have nothing interesting to say about how to form such concepts; cp. the discussion of EAT(_) and DINE(_) in §3.x above. But given that we can form the concept of a water, I think we can form the neutral concept √WATER(_), which applies to *both* whatever the mass concept WATER(_) applies to *and* to waters, whatever they turn out to be. Though in most contexts, it may be clear that the "peunumbral expansion" of

[30] Chierchia (201x) describes the meanings of (at least many) mass nouns in terms of "uncountable countables," using (un)countability as a *psychological* notion. I'm not sure how this fits with formulating meaning theories as truth theories. But I agree with Chierchia—along with many others, including Gillon (1987), Bale and Barner (2009)—that representing some stuff-or-things *as* countable is an aspect of representation in addition to merely representing the stuff-or-things as such.

WATER(_) is irrelevant, and that one is to ignore any waters that √WATER(_) applies to.

In recognizing that there are cases and cases, we should also be wary of misdiagnosing *beliefs* about the nature of the stuff-or-things we talk about as semantic *implications*. Most of us do not believe that there farms on which tofus are raised and slaughtered. But imagine a child whose parents are conspiracy theorists. They regale the child with stories of tofus running free before letting the child see any tofu, and then tell the child that eating tofu is as bad as eating chicken. Such a child might acquire an atomic concept of a tofu, ONE-TOFU(_), and use it to introduce the more permissive concept √TOFU(_), which can be restricted to form a complex concept that is equivalent to ONE-TOFU(_). But *our* concept ONE-TOFU(_) may be complex and contrived, somehow built up from √TOFU(_), which we introduce via some mass concept that we use to think about tofu without ever thinking of it as countable.

I don't see any reason for thinking that human I-languages somehow encode the facts about tofu(s), chicken(s), fish, sheep, rock, paper, scissors, etc. It seems more likely that we use whatever concepts we start with to introduce more neutral concepts that we can then use in thinking that there are no tofus, and that the poor child has been misled about what tofu is—or that there are tofus, and that the general public has been misled about where tofu comes from. A child might acquire an atomic concept ONE-SKY(_), perhaps after hearing utterances of sentences like (128), and introduce √SKY(_) to accommodate sentences like (129).

(128) The sky is blue today. But skies here are never as blue as skies in Arizona.

(129) There is a lot of blue sky today. There was much less blue sky yesterday.

But given my conception of meaning, another child might acquire an atomic mass concept SKY(_) and introduce √SKY(_) to accommodate sentences like (128). Recall the different perspectives on the concepts NUMBER(_) discussed in §1.x of chapter two: we humans may naturally acquire an atomic monadic concept of numbers that we can use to think about the number/cardinality of a certain set; but ideal Fregean thinkers might view our concept NUMBER(_) as mental shorthand for a more complex higher-order concept.

Correlatively, we can find a place for talk of "coercion" in an account of which concepts children are apt to *lexicalize*, without supposing that these concepts are semantically privileged in a way that requires talk of Fregean types and type-shifting. Moreover, while I take lexical items to be atomic expressions of a human I-language, I do not assume that lexical items are the only meaning-pronunciation pairs that speakers store in some form of long-term memory. I think the meaning of the lexical item 'chicken' is the instruction ***fetch*@'chicken'**, whose execution accesses a neutral concept √CHICKEN(_) that applies to chickens, each chicken, and each portion of chicken. My suspicion is that for most speakers, this neutral concept is introduced by using an atomic count-concept, ONE-CHICKEN(_).[31] Once introduced, √CHICKEN(_) can be combined with other concepts to form a complex concept that is analytically equivalent to ONE-CHICKEN(_). This provides a kind of explanation for how we know that chicken comes from chickens. But this does not imply that each time we hear 'a chicken', we go through the bother of constructing √CHICKEN(_)^ONE(_). It may be more efficient to maintain a "second address" for the complex expression [chicken+CT], so that ONE-CHICKEN(_) can be accessed in contexts where the more complex analytically complex concept is not needed. This is a kind of homophony, but not lexical homophony. And if it seems bizarre to even think of "grinding" dots to create a gallon of

[31] As opposed to some mass-concept and a "unitizer." But I once heard an anecdote from a mother reporting an awkward conversation, in which her daughter at first thought it was funny that the stuff they were eating for dinner "had the same name as" the chickens they had seen the day before at the petting zoo.

dot that can be spread all over the wall, then we may need to be "coerced" into forming √DOT(_) and using fetch this neutral concept to think about some stuff as a gallon of dot.

If we never encounter the expression [dot–CT], but often encounter [dot+CT], we might connect the pronunciation of 'dot' with an instruction of the form ***fetch*@L**, where ONE-DOT(_) is the only concept at **L**. But forcing a speaker to deal with [dot–CT] may trigger a process of creating a new entry for the lexical item itself, so that 'dot' comes to have two related addresses, corresponding to two related instructions: ***fetch*@'dot'**, whose execution accesses √DOT(_); and ***fetch*@'dot+CT'**, whose execution accesses √DOT(_)^ONE(_). As usual, we need to distinguish the elements and operations of an I-language from systems that supplement a basic I-language in ways that facilitate its various uses.

Still, one would like a model for posited concepts like √ROCK(_) and √FISH(_). Boolos showed how the capitalized variable in '∀y[OneOf(y, X) ≡ Fish(y)]' can be interpreted as having one or more values relative to a single assignment of values to variables. The invented formula is true of one or more things, the Xs, iff each thing, y, is such that it_y is one of $them_X$ iff it_y is a fish.[32] To extend this singular/plural neutrality in a way that also yields mass/count neutrality, we need to allow for assignments of *value*—like some mud or some beef—to a variable, without insisting that each assignment assign one or more *values* to each variable. Alas, we 'some fish' is ambiguous. So lets use 'a fish' for the singular expression and use 'sm fish' for the mass/plural expressions. We can disambiguate the latter by speaking of sm fish that is not a fish, or sm fish that are not fish. The idea, then, is to relax the Boolos interpretation of variables so a single assignment of *value-or-values* to variables can assign to a single variable *sm value that is not a value* (e.g. sm fish that is not a fish) or *one or more values* (e.g. one or more fish).

I am *not* suggesting that this is a good idea for purposes of studying the relation between logic and *arithmetic*, which is intimately connected with notions of counting countable things. Correlatively, it may not be a good idea to impose a Fregean logic—rooted in notions of mapping countable things onto truth values—on mass concepts or number-neutral concepts. But in any case, my proposal concerns mental representations, not the stuff-or-things represented. Given a domain of countable things, Boolos showed how 'OneOf(y, X)' can be satisfied by an assignment $\mathcal{A}$ that assigns exactly one thing to 'y' and one or more things to 'X': the thing that $\mathcal{A}$ assigns to 'y' just has to be *one of* the thing(s) that $\mathcal{A}$ assigns to 'X'. But I think we can generalize this idea and say that 'SomeOf(y, X)' is satisfied by assignment $\mathcal{A}$ iff whatever $\mathcal{A}$ assigns to 'y'—be it a value or sm value that is not a value—is *some of* whatever $\mathcal{A}$ assigns to 'X'; where $\mathcal{A}$ may assign to 'X' one or more values, or sm value that is not a value.

The formula '∃X∃y[SomeOf(y, X) ≡ Fish(y)]' can be read as follows, using 'be' instead of 'is or are': there be some (stuff or thing or things), the X, such that there be some (stuff or thing), y, such that it_y is some of it-or-$them_X$ iff it_y be fish. As a special case of this "pirate logic" interpretation, consider an assignment $\mathcal{A}$ such that: $\mathcal{A}$ assigns a certain 2-pound fish, and nothing else, to 'X'; and $\mathcal{A}$ assigns the co-located 2 pounds of fish, and nothing else, to 'y'. Absent good

[32] Likewise '∃X∀Y:Plural(Y)[OneOrMoreOf(Y, X) ≡ Fish(Y)]' is true iff there are one or more things, the Xs, such that any one or more things that are plural, the Ys, are such that $they_Y$ are one or more of $them_X$ iff $they_Y$ are (i.e., iff each of $them_Y$ is a) fish. And '∃X∀Y:~Plural(Y)[OneOrMoreOf(Y, X) ≡ Fish(Y)]' is true iff there are one or more things, the Xs, such that any one or more things that are not plural, the Ys, are such that $they_Y$ are one or more of $them_X$ iff $they_Y$ are fish. So at least in principle, the lower case variables are dispensible.

reasons for distinguishing the fish that $\mathcal{A}$ assigns to 'X' from the fish that $\mathcal{A}$ assigns to 'y', I am inclined to identify the former fish with the latter fish, and diagnose this as a special kind of Hesperus/Phosphorus case. With regard to the domain of things-or-stuff, the (2-pound) fish can *be* the (2 pounds of) fish, even if you think of it *both* as a live animal that can gain some weight *and* as some food, of a certain quantity, that can cease to be a live animal. Likewise, I think (130)

(130) The fish is on the table.

can be used to describe on which there are three (whole) fish, so long as we're thinking about those fishes as dinner—or for some other use where counting is not important, and there might as well have sm fish that are not fish on the table. But if there are exactly three fish on the table, I am inclined to identify the fish that is on the table with the fish that are on the table.

That said, I am not committed to *not* distinguishing the fish from the fish; if there are good reasons for distinguishing the animal(s) from the corresponding matter, so be it. I tend to be suspicious of arguments that purport to extract ontological conclusions from modal premises—or premises concerning how we use 'is' as opposed to 'are'—absent independent reasons for thinking that the premises reflect how things are, as opposed to how we think about them. But I'm no metaphysician. My goal is offer a plausible theory of meaning that can be incorporated into a plausible account of human nature and its place in nature.

Other things equal, I prefer theories of meaning that don't have tendentious implications for what there is in nature, independent of human nature. But for those who want meaning theories to take the form of *truth* theories, it would seem especially important that ontological distinctions be motivated by considerations independent of linguistic meaning or conceptual content. In any case, the question of whether or not whatever $\mathcal{A}$ assigns to 'X' differs from whatever $\mathcal{A}$ assigns to 'y'—like the question of whether or not the morning star differns from the evening star—doesn't seem to be a question about human I-languages. So given that humans are apt to mischaracterize intentional distinctions as ontological distinctions, I think semanticists should be especially cautious when making assumptions about the stuff/things that can be talked about. In any case, my use of 'SomeOf(y, X)' does not presuppose that whatever $\mathcal{A}$ assigns to 'X' differs from whatever $\mathcal{A}$ assigns to 'y'.[33]

4.4 Summary

While the details in this chapter have been many, the moral is simple. Lexical items are remarkably flexible in ways that tell against the idea that lexicalization is process in which linguistic expressions *label* concepts and inherit their adicities. Many different facts suggest that lexicalization is a process in which available concepts are used to *introduce* formally new concepts that can be combined via operations that take monadic or dyadic concepts as inputs.

---

[33] Likewise, my use of '~Plural(Y)^OneOrMoreOf(Y, X)' does not presuppose that that whatever A assigns to 'X' differs from whatever A assigns to 'Y'. And we can, in principle, dispense with lower case variables (see note 33). Let '∃X∀Y[SomeOf(Y, X) ≡ Fish(Y)]' be true iff there be some (stuff or thing or things), the X, such that any (i.e., any stuff or one or more things), the Y, be such that it-or-they$_Y$ be some of it-or-them$_X$ iff it-or-they$_Y$ be fish. Restricting to countable things, '∃X:OneOrMore(X){∀Y:OneOrMore(Y)[SomeOf(Y, X) ≡ Fish(Y)]}' is true iff there be one or more, the Xs, such that any that be one or more, the Ys, be such that it-or-they$_Y$ be among it-or-them$_X$ iff it-or-they$_Y$ be fish. It might be said that since two things could have been assigned to 'X' but not to 'y', A('X') could have been two while A('y') was one, and so A('X') is distinct from A('y'). But this is either (i) a crazy inference, or (ii) a way of introducing a notion of distinctness according to which "things" are individuated relative to a kind of variable. CITE JPL ARTICLE.