Chapter One: Introducing Concepts

In the introductory chapter, I tried to describe meanings and human languages in a way that would make the topic clear. I also advertised a few aspects of my positive proposal, according to which lexical meanings are instructions for how to access concepts of a special sort, and phrasal meanings are instructions for how to build concepts that are monadic and conjunctive. But to explain this proposal, I need to say what concepts are, and then characterize the relevant notions of monadicity and conjunction. The notion of a mental *predicate* will play a key role here.

1. Representation and Constituency

I assume that humans and many other animals enjoy mental representations with two striking properties: they can be used to think about things in certain ways; and they exhibit significant relations of constituency, so that a representation used to think about things in one way (e.g., as cows) can be a part of another representation used to think about things in a distinct yet analytically related way (e.g., as brown cows). I use 'concept' to talk about these mental entities, which can be viewed as expressions generated by languages of thought. In my view, our only remotely adequate models of concepts treat them as symbols that have *adicities*—valences of some sort, on analogy with hooks and eyes, or perhaps interlocking puzzle pieces—by virtue of which the symbols can combine as they do; see, e.g., Frege (1892a), Fodor (1998, 2002), Laurence and Margolis (1999), Rey (200x), Gallistel and Gibbon (2002). In this section and the next, I say that more slowly, stressing the *intentionality* and *composability* of concepts.

1.1 Thinking about Things

Let's not argue about nomenclature. One can use 'concept' more permissively than I do, perhaps to include images or other representations that are not composable constituents of thoughts.[1] One can also use 'concept' less permissively, perhaps to exclude representations that fail to meet certain normative constraints. Or might reserve the term for certain representable "contents," as opposed to physically instantiated symbols. But I want to talk about a kind of physically instantiated composition that is important for cognitive science, along with a correspondingly demanding but non-normative notion of constituent, without denying that contents and nonconceptual representations are also important. In the end, one can replace 'concept' with a technical term characterized by its role in my account of how meaning is related to cognition. Until then, I use 'concept' in a restricted way; cp. Fodor (2008). Still, some readers may be unfamiliar or uncomfortable with any talk of using representations to think about things in certain ways. So some clarification, regarding aboutness and ways, is in order.

For present purposes, the relevant notion of *thinking about* is intentional. We can think about unicorns, even though there are no unicorns to think about.[2] Indeed, one can posit unicorns and spend a lot of time thinking about them, constructing theories of them, diagnosing various

---

1 A camera may contain images of a fish, and a single image can be the result of stacking shots. Given disparate images, one can also form a less detailed "abstraction" that preserves certain features shared by inputs. But as Fodor (1998, 2003) notes, it is hard to see how any image can be used to think about fish *in general*, or to think about things *as fish*; see also Cummins (1989). An image can be associated with a concept. But a typical pet image, whatever that would be, cannot combine with a typical fish image to form a third image with which one could think about pet fish. Combining a pet image with a fish image will typically yield a mess. So it's not plausible that images are concepts in the sense described here; cp. Prinz (2002). Likewise for prototypes; see Fodor and Lepore (2002).

2 Or at least no real unicorns. In a toy store, on in a picture book, the unicorns might be next to the teddy bears. Even if every possible world is devoid of real unicorns, (see Kripke 1980), there are actual simulacra. And in constructing certain things—toys, pictures, arguments, etc.—one may well think about unicorns. Perhaps one needs words to think about unicorns. But I see no reason for thinking that every "empty concept" reflects some word.

observations as symptoms of unicorns, etc. Similarly, one might hypothesize that a certain planet, Vulcan, passes between Mercury and the sun, and then think about this alleged planet—to estimate its mass, or wonder whether it is habitable—much as one might hypothesize that Uranus passes between a certain planet, Neptune, and the sun. A mental episode can be one of thinking about something, even if nothing is such that the episode is one of thinking about it.

One can think about countable spatiotemporal things. But one can also think about stuff, like gold or phlogiston. And the things/stuff we think about include abstracta, like numbers or intelligence. Though whenever one thinks about some stuff, or one or more things, one thinks about it or them in some way; cp. Evans' (1982) discussion of Frege (1892b). And things/stuff can be thought about in various ways. Thinking about Hesperus can differ from thinking about Phosphorus, even though Hesperus is Phosphorus. For while the planet Venus is always itself, thinking about it as the evening star differs from thinking about it as the morning star. Thinking about woodchucks can differ from thinking about groundhogs, despite woodchucks being groundhogs; likewise for cilantro and coriander, eggplant and aubergine. Thinking about the aces can differ from thinking about the marked cards, even given a rigged deck. I take concepts to be mental representations that can be used to think about things in this intentional sense that allows for thoughts of unicorns and for Hesperus-thoughts that are not Phosphorus-thoughts.

We can still say that many concepts *apply* to mind-independent entities, like animals and planets, where the notion of application is not intentional. Consider an ancient astronomer who used exactly one concept C to think about Hesperus, in the intentional sense of thinking about the evening star *as such*, and a distinct concept C* to think about Phosphorus as such. We can say that both concepts applied to Venus: even if the astronomer would have denied it, both of his concepts applied to Hesperus, and both applied to Phosphorus. Likewise, a concept used to think about groundhogs (as such) can apply to woodchucks, which are such that any concept used to think about them applies to groundhogs. Distinct concepts, which let thinkers think about things in different ways, may apply to the same things. But thinkers can also have "empty" concepts that do not apply to anything; cp. Husserl (189x), Dummett (1981, 1994), Rey (201x). And as we'll see, such concepts can be enormously valuable.

Let me stress that these familiar points concern thought, not how we talk about it. I grant that 'think about' can also be used to talk about a relation that thinkers can bear to entities/stuff to which concepts can apply; see, e.g., Burge (197x, 2010). In this "*de re*" sense, one can think about bosons and dark matter only if the world includes bosons and dark matter, about which one can think. Any episode of thinking *de re* about Hesperus is an episode of thinking *de re* about Phosphorus. This extensional/externalistic notion has utility, in ordinary conversations and perhaps in cognitive science, when the task is to describe representations of a shared environment for an audience who may represent that environment differently. But however this notion is related to words like 'think' and 'about', many animals have concepts that let them think about things in ways that are individuated intentionally. We animals may also have representations that are more heavily anchored in reality; see, e.g., Pylyshyn (2008). But a lot of thought is intentional, however we describe it.[3]

---

3 If any stuff we think about is a collection of things, any episode of thinking *de re* about some stuff is an episode of thinking *de re* about a collection of countable things. But prima facie, thinking about mud (or space) as such differs from thinking about any collection of countables; see chapter five. One can be professionally suspicious of the intentional; see, e.g., Quine (1960). But prima facie, human and many other animals can think about a certain card in one way when it is face up, and another way when the same card is face down. And I don't see how skepticism about this can be justified by reflection on how humans *use phrases* like 'thinks the ace is on the left'.

In chapters four and six, I briefly discuss phrases like 'think about Venus', 'concept of groundhogs', and 'doubts that Hesperus is Phosphorus'. Inten*t*ional idioms, not to be confused with intensions (i.e., procedures) or intensional logics, tell us something about meaning. But studying ordinary thought reports and concept descriptions may reveal little about the mental representations described, much as studying 'It rained' may reveal little that interests meteorologists. There are interesting puzzles in the vicinity; see Kripke (1979a). One wants to know how two concepts *can* be used to think about something in different ways: if C and C* both represent Venus, as opposed to any way of thinking about Venus, why don't both represent Venus as such? But we face these questions about thought, whatever we say about language. And distinguishing contents from meanings may help demystify both. If lexical meanings are instructions for how to access concepts from lexical addresses, then two words can differ in meaning even if they are (or come to be) linked to a single concept/content.

In any case, however we talk about thought, we can think about things. I have said that concepts are mental symbols that can be used to think about things. So let me turn to the other key feature of concepts—viz., they exhibit significant relations of constituency.

1.2 Productive Atoms

Following custom, let's indicate concepts with small capitals. And let's say that COW can be used to think about cows as such. A thinker can use the concept COW to think about cows *as* cows.

There may well be more than one such concept, depending on what is required for thinking about things *as* cows.[4] We can also think about cow-stuff (some of which we call 'beef'), herds of cows, the property of being a cow, etc. So perhaps we should say that for each of many indices *n*, COW_n is a cow-concept. But typological distinctions will soon be needed in any case; see section two. For now, let's pretend that COW is the only concept used to think about cows as such—i.e., the only concept type whose spatiotemporally instances are so used. Likewise, let's pretend that VENUS is the concept used to think about Venus as such, without yet encoding the idea that VENUS is about a *particular* thing, while COW is a more *general* concept. Let's also bracket polysemy, as if BOOK is *the* concept used to think about books as such, even though books can be shelved or written. Before turning to these and other complications, I want to stress the crucial point that many concepts are somehow built up from others.

Thinkers who can think about brown things as such have a concept BROWN. Those who can think about brown cows as such have a concept BROWN-COW. This notation leaves room for the hypothesis that both concepts, along with COW, are conceptual atoms with no conceptual parts. (Treat the hyphen as a character, like 'O', that can appear in concept-designators.) The notation also allows for the hypothesis that all three concepts are composite. Perhaps COW has BEEF as a constituent, or vice versa, or the concept we use to think about brown things shares a part with the concept we use to think about brown stuff. One can speculate that COW is complex in a more substantive way, with a constituent like ANIMAL; though if COW has any such analysis, the details remain obscure.[5]

---

4 Can children think about cows this way on the basis of encountering just a few cows? How about cattle dogs, or cattle? Under what conditions would encountering a few pictures of cows suffice? As discussed below, my claims about meaning can be evaluated without taking a stand on these issues, which provide ample opportunity for terminological debate about which pairs of individuals share a given concept/word/language.

5 And *pace* Kantian enthusiasm for genera/differentia, GOLD cannot be decomposed into METAL and YELLOW, even if YELLOW are COLOR are analytically related. Like Fodor and Lepore (2002), I think many putative decompositions of word-level concepts are implausible; see note X. Though if formally atomic concepts can be *introduced* in the course of lexicalization, then *some* analyses may be sustainable; cp. Pietroski (2003b). In chapter two, I return to this point and its relation to Frege's (1884) notion of analysis, illustrated with his discussion of numbers.

By contrast, BROWN-COW does seem to be built up from BROWN and COW, and not just because 'brown cow' is built up from 'brown' and 'cow'. Prima facie, one cannot think about things as brown cows without thinking about them as cows. A thinker who lacks COW might have a concept that applies to all and only the brown cows. But it seems that no such concept could be used to think about brown cows *as such*. In traditional terms, it seems analytic that brown cows are cows. Simple logic, of a sort that ordinary humans appreciate, appears to ensure that COW applies to whatever BROWN-COW applies to. However this intuition is described, it bolsters the idea that BROWN-COW has simpler concepts as parts. To a first approximation, an episode of thinking about things as brown cows is an episode of thinking about things as cows *and* as brown things—or perhaps, to a second approximation, as cows and brown ones. Having BROWN-COW requires having both COW and BROWN, along with some capacity to combine these two concepts conjunctively; cp. Fodor and Lepore (2002) on "reverse" compositionality; see

By contrast, even if one discovers or defines a concept DIFF that applies to all cows and differentiates them from other animals, it won't be plausible that COW is a conjunction of ANIMAL with DIFF. It won't be plausible, even to a first approximation, that to think of things as cows is to think of them *as* animals to which DIFF applies. For one can think about things as cows *without* thinking about them as animals. Indeed, one can suspect (without contradiction) that cows are cleverly disguised Martian robots; see Putnam (197x). While cows are animals, and any evidence to the contrary is misleading, it remains logically/conceptually possible that cows are not animals. But it is not even logically possible that brown cows are not cows.[6]

If BROWN-COW is indeed a conjunction whose constituents include COW and BROWN, then these concepts are somehow combinable, and likewise for many pairs of concepts. Not all complex concepts are conjunctive, even if all phrases direct construction of conjunctive concepts; see §2.1 below. But whatever modes of composition are available to humans, we can think about cows that danced, brown cows that danced on Tuesdays, dogs that chased a cow, dogs that chased every dancing cow, cows a dog chased, etc. Humans can think about things in endlessly many ways. One might suspect, as I do, that this capacity is interwoven with our capacity to generate expressions like 'cow that danced' and 'cow a dog chased'. Though however the capacities emerged in natural history, humans can form boundlessly many concepts that have analyses, even if the details are unclear.

Our atomic concepts may outnumber the lexical expressions of English, even ignoring polysemy, if only because many concepts may go unexpressed. But each human has only finitely many concepts that have no conceptual constituents. Still, we can form boundlessly concepts that have distinct but *systematically related* contents. Given DOCTOR, LAWYER, and NEED, we can add DOCTOR-WHO-NEEDS-A-LAWYER and LAWYER-WHO-NEEDS-A-DOCTOR. Given COW and DANCED, we can add COW-THAT-NEEDS-A-DOCTOR, DOCTOR-WHO-DANCED, and so on.

Moreover, this kind of concept composition is *cognitively productive* in a way that goes beyond mere recursive generation of representations. A thinker might be able to generate formally new concepts *ad nauseum*, without increasing the number of ways it can think about things. Imagine a mind that can conjoin COW with BROWN, conjoin the result with COW, conjoin the result with BROWN, etc. To be concrete, suppose that BROWN-COW is the complex concept

6 Ambiguity complicates. Sometimes, a brown cow is a root beer float. But at least one concept we can express with 'brown cow' is built from concepts we can express with 'brown' and 'cow'. The stuff/thing contrast suggests more complexity. Like a brown house, but unlike brown paint (or root beer), a live brown cow needn't be brown throughout; cp. Chomsky (2000b). Concepts like CHICKEN-EGG present further twists, given that at least one chicken emerged from an egg not laid by a chicken; though perhaps a chicken egg is both an egg and chicken.

+{COW, BROWN}; where '+' stands for some form of conjunction, and brackets indicate a list of conjuncts. If we focus on concept types, and adopt a set-theoretic perspective, {COW, BROWN} is identical to {COW, BROWN, COW, BROWN}. But a list can include more than one token of a concept. By distinguishing first tokens from seconds, we can distinguish {`COW, `BROWN} from {`COW, `BROWN, ``COW, ``BROWN}. Though using concepts twice, in this way, does not provide a substantively distinct way of thinking about things.

One can insist that *any* distinction matters, and that thinking of things as brown cows *differs* from thinking of things as brown cows that are brown cows. But this difference is, to borrow Frege's (1892b) phrase, cognitively insignificant: it is trivially analytic that brown cows are brown cows that are brown cows. By contrast, thinking about doctors who need lawyers differs from thinking about lawyers who need doctors, even if the lawyers and the doctors turn out to be the same individuals; cp. Mates (195x).

I have stressed the unbounded cognitive productivity of human concepts for two reasons. First, mere conjunction of atomic concepts cannot explain it. Second, other animals may have concepts that exhibit only more restricted kinds of productivity; see Spelke (200x), Carruthers (200x). But whatever we say about other animals, humans enjoy mode(s) of conceptual composition that yield unbounded cognitive productivity given suitable atomic concepts. This raises the question of *which* composition mode(s) we employ to generate phrases like 'doctor who needs a lawyer' and the corresponding concepts. I think we use relatively simple capacities to conjoin concepts, along with a few restricted tricks for creating mental predicates, which I take to be concepts of a special sort. But thus far, I have spoken of concepts without categorizing them. And not only are there distinctions to make, the notion of a mental predicate lies at the cusp of two historically important descriptions of thoughts and their constituents.

Let's say that *atomic* thoughts have no other thoughts as parts. One classical view is that atomic thoughts are mental sentences that exhibit Subject-Copula-Predicate structure. On this view, articulated by Aristotle and the medieval logicians, predicate-*conjunction* plays a central role in logic. As discussed in section two and again in chapter two, I think that atomic thoughts are more varied, but that the classical tradition was importantly right about a kind of thought that is intimately connected with human I-languages. Frege (1879, 1892a)—or at least Frege as reconstructed by Church (1941, 195x) and others—offered an alternative picture of thoughts as exhibiting Function-Argument structure. As discussed in section three, this led to a much better appreciation of how *relational* concepts are logically related. The idea is that a unary/binary/*n*-ary function (from entities to truth values) can be represented with monadic/dyadic/*n*-adic concept that can be "saturated" by the corresponding number of "saturating" concepts, each of which indicates an entity in the function's domain; a thought can thus indicate the (truth) value of the function given the indicated *n*-tuple of arguments. Monadic concepts are the analogs of classical predicates, which can be viewed as special cases of "unsaturated" thought components.[7]

This makes it possible to explain the notion of a mental predicate via classical and Fregean accounts of how thoughts are structured and logically related. And with a tolerably clear notion of a mental predicate in hand, I can turn (in section four) to the idea that phrasal meanings are instructions for how to build mental predicates via certain operations that severely constrain what lexical meanings can be.

[7] There is a technical use of 'predicate' that applies to *any* unsaturated thought component, or least any that can be saturated by an *n*-tuple of unsaturable concepts. But I will avoid this usage and distinguish predicates from *relations*.

2. An Old Typology

Let's say that a *general* concept can be used to think about one or more things (or some stuff) in a general way, while a *denoting* concept can be used to think about something as a particular thing (or quantity of some stuff). There are also *logical* concepts, including *quantificational* concepts that can be used with general concepts to form generalizations. In chapters five and six, I discuss *plural* concepts, *count* versus *mass* concepts, and the possibility of *neutral* concepts. But my aim is not to provide a complete typology or a substantive theory of concepts. My goal is to specify the types such that lexical and phrasal meanings are instructions for how to access and assemble concepts of those types, and to offer an account of why human languages privilege those conceptual types as opposed to others. For these purposes, it is especially important that general concepts can be *classificatory/monadic* or *relational/polyadic*: the former facilitate acts of predication, in which subjects of thoughts are represented as meeting certain conditions; the latter facilitate acts of relating things in certain ways. In thinking that Venus is a planet that orbits the sun, one might use a denoting concept to think about Venus, a classificatory concept to think about planets, and a relational concept to think about one thing orbiting another.

Ordinary language provides only a rough guide as to whether concepts are classificatory or relational. (Fathers father children; and mortals can suffer mortal wounds, as if 'mortal' implies a relation to a death.) Yet the distinction seems clear enough, even if application to cases is hard. This leaves it open whether any species of general concepts is more basic than any other. But there is an ancient tradition of focusing initially on classification.

2.1 Sentences, Subjects, and Predicates

I assume that there are *natural* sentences, distinct from formulae that theorists invent. The natural sentences may include some pronounceable expressions of human I-languages. But for reasons that will emerge, the classical view was that grammatical structure imperfectly reflects the Subject-Copula-Predicate structure of "deeper" sentences.[8] Many current theories of grammatical form go farther, not even employing the notions of sentence, subject, and predicate: expressions are described as phrasal projections of lexical items, which include "functional" items (e.g., complementizers, tense morphemes, and prepositions) along with "open class" vocabualary (e.g., nouns and verbs); and no lexical category is the obvious correspondent of "Sentence" in older theories. I'll return to the idea that a human language does not generate sentences as such, even if it generates phrases that we often call sentences. But in any case, one can grant that many *thoughts* have subjects and predicates, even if pronounceable expressions do not. There can be *languages of thought* that generate sentences in accord with procedures like the following.

Sentence ➔ Subject copula Predicate
Predicate ➔ PLANET, BRIGHT, COW, BROWN, PHILOSOPHER
Copula ➔ *IS*, *WAS*
Subject ➔ Quantifier Predicate
Quantifier ➔ **EVERY**, **SOME**, **NO**
Subject ➔ Denoter
Denoter ➔ VENUS, SOCRATES

Such a procedure generates mental sentences like (1-3).

(1) [**EVERY**$_{\text{Quant}}$ PLANET$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ BRIGHT$_{\text{Pred}}$

[8] Aristotle spoke of "affectations of the soul." Ockham and others explicitly posited a mental language, which would have no need for Latin's declensions; see Pannacio (2007).

(2) [VENUS$_{Den}$]$_{Subj}$ *IS*$_{Cop}$ PLANET$_{Pred}$

(3) [VENUS$_{Den}$]$_{Subj}$ *IS*$_{Cop}$ BRIGHT$_{Pred}$

Subscripts are reminders of the typology, and outermost brackets of the form […]$_{Sent}$ are suppressed for simplicity. Mental symbols may connect certain *contents* with biochemically realized signals that are related to perception, inference, and action. Recall that languages connect interpretations of some kind with signals of some kind; meanings and pronunciations are special cases. An animal with no capacity for external communication may have thoughts of the "categorical" forms exhibited with (1-3). Humans can generate pronounceable analogs of these sentences, despite quirks of grammar—e.g., that English permits 'is bright' and 'are planets', but not 'is planet'.[9] Though even if we can use (4) to express (3), the meaning of (4) may be related to (4) only indirectly.

(4) Venus is bright.

More specifically, the meaning of (4) may be an instruction for how to build a sentential concept, and executing the instruction may yield a thought that is equivalent to (3) yet formally different. As discussed in chapter two, Frege showed how a string like 'every integer is rational' might be used to build a thought that can be viewed as a simplification of a more ideal thought. And at least in principle, (4) might be used to build a quantificational thought like (4a) or (4b);

(4a) [**EVERY**$_{Quant}$ VENUSIZER$_{Pred}$]$_{Subj}$ *IS*$_{Cop}$ BRIGHT$_{Pred}$

(4b) [**SOME**$_{Quant}$ BRIGHTNESS-STATE$_{Pred}$]$_{Subj}$ *IS*$_{Cop}$ EXHIBITED-BY-VENUS$_{Pred}$

where VENUSIZER applies to whatever VENUS denotes, and states can be viewed as boring events; cp. Quine (195x), Parsons (1990). Similarly, (5) might be used to build (5a) or (5b)

(5) Every planet rose.

(5a) [**EVERY**$_{Quant}$ PLANET$_{Pred}$]$_{Subj}$ *IS*$_{Cop}$ ONE-THAT-ROSE$_{Pred}$

(5b) [**EVERY**$_{Quant}$ PLANET$_{Pred}$]$_{Subj}$ *WAS*$_{Cop}$ RISER-IN-A-RISING$_{Pred}$

via some procedure that underlies a distinctively human form of unbounded cognitive productivity. For now, though, it will do no harm to pretend that (4) is a simple case in which grammatical structure directly reflects conceptual structure.

From this perspective, (4) has as a subject in the derivative sense of having a constituent that indicates the subject of (3). But thoughts, or least episodes of thinking, can have subjects in a sense that is not derivative or merely formal. One can think that Venus is bright by perceiving or remembering Venus, and thereby thinking about it in a certain way, while classifying it as bright. A mental act of this sort has Venus as its subject matter, and not just because VENUS$_{Den}$ and BRIGHT$_{Pred}$ differ formally. Using classical terminology, we can say that the judgment is one in which brightness is predicated of Venus, which can be thought about independently of any predication.[10] But this might suggest that BRIGHT$_{Pred}$ *denotes* brightness, and that predicates are special cases of denoters; cp. Frege (1892a), discussed in chapter two. To avoid this suggestion, it may be better to say that the judgment is *predicational*, with brightness as the content of the predication, and *of* Venus; where there can also be thoughts *of* Vulcan (see §1.x above).

I don't deny that sentences like (4) can be used in mental acts, say of hypothetical reasoning, that are less directly about Venus. And a quantificational sentence may have its

---

[9] At least not if 'planet' is used as a count noun; though a bumbling creator might say, 'There is mud all over the floor, and there is planet all over the yard'. Many other human languages, including ancient Greek, do not require that an unplural count noun appear with an article.

[10] *Pace* Russell (1912). I don't think there were ever any good reasons for denying that humans and other animals have singular concepts of distal objects. But in light of Kripke (1980), Evans (1982), and many others, it seems that intuitive considerations *bolster* other reasons for positing such concepts. Citations to Pylyshyn and Burge.

subject in a derivative way; perhaps (1) is *about* the planets only because it predicates brightness of each planet and nothing else. On the other hand, if (1) restricts the *domain* of quantification to the planets—with brightness predicated of everything in the relevant domain—then the planets may be the subject matter of (1), just as Venus is the subject matter of (2) and (3). But even if some mental sentences have subjects in a derivative sense, any grammatical notion of subject seems to be more removed from the notion of a potential target for predication.

One can say that the grammatical subjects of (6) and (7) are pleonastic—

(6) It rained yesterday.

(7) It was easy to please the guest.

i.e., they do not indicate subjects of corresponding thoughts. One can also say that in (8),

(8) The guest was easy to please.

'the guest' is both the non-pleonastic subject of (8) and part of its tensed predicate. But then talk of grammatical subjects seems to be (at best) a way of saying that tensed clauses like (4-8) have a "left edge constituent" that somehow makes them suitable for expressing complete thoughts. By contrast, "mere phrases" like (9) and (10) feel incomplete if used to express thoughts.

(9) rained yesterday

(10) easy to please the guest

Correlatively, 'it' seems to be logically irrelevant in (7), which is somehow equivalent to (8). And while (11) somehow implies that something barked yesterday,

(11) It barked yesterday.

(6) does not likewise imply that something rained yesterday. The phenomenon is real. But the notion of a grammatical subject, like the notion of a complete sentence, seems to depend on the idea of (completely) expressing a thought.

More generally, it has long been clear that grammatical forms do not perfectly reflect the logical forms of thoughts, whatever these forms are. The ancient Greeks were aware of sophisms like the following: that dog is a father, and that dog is yours; so that dog is your father. This inference, obviously bad, differs somehow from a superficially parallel but good inference: that dog is a mutt, and that mutt is yours; so that dog is your mutt. (See Plato, *Euthydemus* 298 d-e.) This raises the question of what makes good inferences good, and Aristotle's partial answer seems fundamentally correct. *Sentences* exhibit valid *forms* of inference.

The relevant form exhibited by (12-14) can be described in various ways.

(12) $[\textbf{EVERY}_{\text{Quant}}\ \text{HUMAN}_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{ANIMAL}_{\text{Pred}}$

(13) $[\textbf{EVERY}_{\text{Quant}}\ \text{ANIMAL}_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{MORTAL}_{\text{Pred}}$

(14) $[\textbf{EVERY}_{\text{Quant}}\ \text{HUMAN}_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{MORTAL}_{\text{Pred}}$

Using 'Φ', 'Ψ', and 'Ω' as placeholders for predicates, we might say that any instance of the ordered triple (15) is such that the first two sentences jointly imply the third.

(15) $\langle[[\textbf{EVERY}_{\text{Quant}}\ \Phi]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \Psi]_{\text{Sent}}\ ,\ [[\textbf{EVERY}_{\text{Quant}}\ \Psi]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \Omega]]_{\text{Sent}}$

$[[\textbf{EVERY}_{\text{Quant}}\ \Phi]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \Omega]_{\text{Sent}}\rangle$

But let's assume that sentences can have sentential constituents as indicated below.

Sentence ➔ **NOT** Sentence

Sentence ➔ Sentence ***AND*** Sentence

Sentence ➔ Sentence ***IF*** Sentence

Then one can say that instances of (16) are truisms.[11]

[11] This also makes it possible to formulate schematic principles of propositional logic. S1 and S2, so S2 and S1;

(16) $[[\textbf{EVERY}_{\text{Quant}}\ \Phi]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \Omega]_{\text{Sent}}$ ***IF***
$[[[\textbf{EVERY}_{\text{Quant}}\ \Phi]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \Psi]_{\text{Sent}}$ ***AND*** $[[\textbf{EVERY}_{\text{Quant}}\ \Psi]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \Omega]_{\text{Sent}}]_{\text{Sent}}$

By contrast, typical instances of (17) are not truistic.

(17)] $[[[\textbf{EVERY}_{\text{Quant}}\ \Phi]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \Omega]_{\text{Sent}}$ ***IF*** $[[\textbf{EVERY}_{\text{Quant}}\ \Phi]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \Psi\ ]_{\text{Sent}}]_{\text{Sent}}$
AND $[\textbf{EVERY}_{\text{Quant}}\ \Psi]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \Omega]_{\text{Sent}}]_{\text{Sent}}$

Instances of (16) are also instances of the following form, many instances of which are not truistic: S1 ***IF*** [S2 ***AND*** S3]. So evidently, sentences with no sentential constituents can have logically relevant structure. If one is concerned with the nature of logic, rather than psychology, one might not want to employ placeholders for mental predicates. It might be better to use (second-order) variables ranging over things of which predicates might be predicated. Then schemata like (18), read as conditionals rather than conjunctions,

(18) every Φ is Ω if every Φ is Ψ, and every Ψ is Ω

can be viewed as truistic generalizations that are not restricted to expressions of some language. As we might now put it, with help from Boolos (1998): whatever the Φs, Ψs, and Ωs may be—cows, grazers, brown cows, brown animals that graze, or whatever—every one of the Φs is one of the Ωs *if* every one of the Φs is one of the Ψs, and every one of the Ψs is one of the Ωs. But let's not worry here about the difference between (18) and (16). Let's also ignore the difference between *conditional* schemata like (18) and *argument* schemata like (19).

(19) every Φ is Ψ
<u>every Ψ is Ω</u>
every Φ is Ω

The important point is that the variables in such generalizations correspond to predicates. Moreover, Aristotle did not merely identify a few valid forms like (19-21).

(20) every Φ is Ψ
<u>no Ψ is Ω</u>
no Φ is Ω

(21) some Φ is Ψ
<u>every Ψ is Ω</u>
some Φ is Ω

He offered an interesting description of the many valid syllogisms that exhibit this "categorical" form, including those in which one of the three propositions has a subject (of the form 'Every/Some/No Φ') that combines with its predicate *negatively*, as in the conclusion of (22).

(22) some Φ is Ψ
<u>no Ψ is Ω</u>
some Φ isn't Ω

As he also noted, an instance of (23) contradicts the corresponding instance of (24),

(23) some Φ isn't Ω
(24) every Φ is Ω

much as instance of (25) contradicts the corresponding instance of (26).

(25) no Φ is Ω
(26) some Φ is Ω

In this sense, sentences (27) and (28) are equivalent.

(27) $[\textbf{SOME}_{\text{Quant}}\ \text{COW}_{\text{Pred}}]_{\text{Subj}}\ \textit{ISN'T}_{\text{cop}}\ \text{BROWN}_{\text{Pred}}$
(28) $[\textbf{NOT}\ [[\textbf{EVERY}_{\text{Quant}}\ \text{COW}_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \text{BROWN}_{\text{Pred}}]_{\text{Sent}}]_{\text{Sent}}$

S2 if S1, and not-S2, so not-S1; etc. Stoic logicians also discussed structural ambiguities, the differences between sentences and truth evaluable judgments/claims, and the relevance of various contextual factors—including deaths, which raise the question of how one can refer to those who no longer exist; see Bobzien (1998).

By exploiting such equivalences, and sometimes proving a conditional by *reductio*—showing that the antecedent cannot be consistently conjoined with the negation of the consequent—Aristotle showed that the valid categorical syllogisms can be derived from (19-22), which came to be known as Barbara, Celarent, Darii, and Ferio; see Parsons (2014) for detailed discussion.

To take a trivial example, given that no Ω is Ψ if and only if no Ψ is Ω, (29) can be derived from (30).

(29) every Φ is Ψ
<u>no Ω is Ψ</u>
no Φ is Ω

(30) every Φ is Ψ
<u>no Ψ is Ω</u>
no Φ is Ω

More interestingly, given that (23) contradicts (24), (31) can be derived by *reductio* from (19);

(31) some Φ isn't Ψ
<u>every Ω is Ψ</u>
some Φ isn't Ω

To affirm the premises of (31) while denying its conclusion is to affirm each of the following: some Φ isn't Ψ; every Ω is Ψ; every Φ is Ω. Or replacing the first element of this trio with its equivalent, and the reversing the order: every Φ is Ω; every Ω is Ψ; not every Ω is Ψ. But any trio of this form would violate (19). So we need not view (31) as a basic form of valid inference.

2.2 Predicates, Conjunction, and Reduction

I am reviewing some classical logic, in order to provide some initial grip on the notion of predicate I have in mind when I say that meanings are instructions to build mental predicates.

Aristotle's logic suggested that predicates figure in certain patterns of inference that humans find naturally compelling. He was especially interested in logical relations exhibited by two or more *generalizations*. But as he recognized, (32) is also compelling;

(32) δ is Ψ if every Φ is Ψ, and δ is Φ

where 'δ' corresponds to denoting concepts like VENUS. This is, of course, the syllogism relevant to the inference from (1) and (2) to (3).

(1) $[\textbf{EVERY}_{\text{Quant}}\ \text{PLANET}_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{BRIGHT}_{\text{Pred}}$
(2) $[\text{VENUS}_{\text{Den}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{PLANET}_{\text{Pred}}$
(3) $[\text{VENUS}_{\text{Den}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{BRIGHT}_{\text{Pred}}$

There also seem to be cases in which two predicates are *joined* to form a third. Prima facie, we can use 'Venus is a brown cow' to express an analog of (33), which implies (34).

(33) $[\text{VENUS}_{\text{Den}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ [\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}$
(34) $[\text{VENUS}_{\text{Den}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{COW}_{\text{Pred}}$

Relatedly, it seems that (35) is a truism, while (36) is not.

(35) $[\textbf{EVERY}_{\text{Quant}}\ [\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \text{COW}_{\text{Pred}}$
(36) $[\textbf{EVERY}_{\text{Quant}}\ \text{COW}_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ [\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}$

As noted in section one, $[\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}$ seems to be a conjunction of its constituents. And as medieval logicians knew, such predicates interact with quantifiers in interesting ways.

Suppose that some brown cow is a dumb animal. Then some cow is a dumb animal, and some brown cow is an animal. Conversely, if no cow is dumb, then no brown cow is dumb, and no cow is a dumb animal. In a sentence of the categorical form (37),

(37) $[\ldots_{\text{Quant}}\ \ldots_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{Cop}}\ \ldots_{\text{Pred}}$

the indefinite quantifier **SOME** licenses *reduction* of a conjunctive predicate—e.g., replacing $[\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}$ with COW_{Pred}—in both the "intrasubject" position and the "copularized" position, while the negative quantifier **NO** licenses *expansion* of a predicate in both positions.

Intuitively, reduction creates a logically weaker predicate, while expansion creates a logically stronger predicate: [**SOME**$_{\text{Quant}}$ [BROWN$_{\text{Pred}}$ COW$_{\text{Pred}}$]$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ DUMB$_{\text{Pred}}$ is logically stronger than [**SOME**$_{\text{Quant}}$ COW$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ DUMB$_{\text{Pred}}$; but [**NO**$_{\text{Quant}}$ COW$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ DUMB$_{\text{Pred}}$ is logically weaker than [**NO**$_{\text{Quant}}$ [BROWN$_{\text{Pred}}$ COW$_{\text{Pred}}$]$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ DUMB$_{\text{Pred}}$.

As the Stoics (see note 11) had observed, conjunct reduction seems to be the default direction of inference for *sentences*: S1 and S2; so S1. When sentences are offered without any explicit connectives, as in a story, the default expectation is that each claim is offered as an *addition* to those already presented. But this pattern can be reversed with (classical) disjunction: S1; so either S1 or S2. Likewise, sentential negation licenses expansion in its scope: not S1; so not both S1 and S2. This pattern apparently extends to predication. If one of the brown cows is one of the dumb animals, then one of the cows is one of the animals; but if Venus isn't a cow, then Venus isn't a brown cow. In one sense, this is hardly surprising. Who would want the expectation to be that lengthening a story (or a list of premises) reduces the content? But thinking about how quantifiers relate to predicate reductions/expansions, thereby focusing on minimally different *pairs* of sentence, turned out to be a good idea.

One can view the so-called *dicta de omni et nullo* in this light: the default (all-purpose) rule is to permit predicate reduction; though in special cases, with negation as the prototype, the direction of valid inference is inverted. We can imagine mental languages that work differently. But if it is analytic that [COW$_{\text{Pred}}$] applies to whatever [BROWN$_{\text{Pred}}$ COW$_{\text{Pred}}$]$_{\text{Pred}}$ applies to—while the converse assumption is substantive, and equivalent to (36)—then minds that deploy such concepts may well be disposed to treat reduction as the "normal" direction of inference. Correlatively, one might expect minds to treat junctions of predicates as *con*junctions. In any case, it is interesting that in terms of implications involving complex predicates, the universal quantifier **EVERY** is like **NO** with regard to the subject position, but like **SOME** with regard to the copularized position: if every cow is an animal, then every brown cow is an animal; yet if every cow is a dumb animal, then every cow is an animal.

Let concatenation of predicate variables as in 'some ΩΦ' signify a complex predicate like 'brown cow', in which two predicates have been suitably conjoined; and ignore any effects of linear order, like the ungrammaticality (in English) of 'some cow brown is an animal'. Then we can list six generalizations like the following: if some ΩΦ is Ψ, then some Φ is Ψ; etc. But these generalizations can be perspicuously summarized with the table below.[12]

[12] See Ludlow's (2002) summary and development of Sanchez (19xx), which reviews relevant medieval achievements. I am greatly indebted to Ludlow's discussion. Within linguistics, Ladusaw (1980) animated discussions of these quantifier properties by observing that so-called "negative polarity items"—expressions like 'ever' and the idiom 'give … a plug nickel', which seem to require a negative context (as in 'I wouldn't/*would give you a plug nickel for this book, and you should never/*ever read it')—are licensed in the "inverting" environments (now described as "downwardly entailing"), but not in the default environments that license predicate reduction.

No/Every/*Some *cow that ever grazed* is an animal.

No/*Every/*Some cow *is an animal that ever grazed*.

Though as Ludlow stresses, the quantifier **MOST** does not license reduction or restriction in the subject position. It may be that most (of the) brown cows are dumb, while most cows aren't dumb; and most cows can be dumb, even if most brown ones aren't. Still, 'Most cows that ever grazed are dumb' is fine. Ludlow offers a specific proposal according to which a quantifier can mark one of its predicate positions in the non-default way, so that reduction is not licensed, without thereby *inverting* the default direction of inference; cp, Ladusaw (1996).

	Predicate in Subject	**Copularized Predicate**
Some	default	default
No	*inverted*	*inverted*
Every	*inverted*	default

From this perspective, Aristotle's universal quantifier is governed by two inferential principles that characterize its distinctive logical role.

(38) $\frac{\text{every } \Phi \text{ is } \Omega\Psi}{\text{every } \Phi \text{ is } \Psi}$ (39) $\frac{\text{every } \Phi \text{ is } \Psi}{\text{every } \Omega\Phi \text{ is } \Psi}$

Principle (38) reflects the fact that **EVERY** leaves the copularized ("external") predicate position, corresponding to 'Ψ', in the default state of licensing *reduction*—e.g., $[\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}$ to COW_{Pred}. Principle, (39), reflects the fact that with regard to intra-subject ("internal") predicate position, corresponding to 'Φ', **EVERY** inverts the normal inferential polarity and licenses *expansion*—e.g., COW_{Pred} to $[\text{BROWN}_{\text{Plred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}$. We can view (38) as an "elimination rule" that can go unstated in a system that permits reduction in the absence of any quantfier-specific rule. But in such a system, (39) is a substantive "introduction" rule, which captures the idea that instances of (40) permit the inverted direction of inference for predicates in the 'Φ'-position.[13]

(40) every Φ is Ψ

I'll return to another sense in which an instance of (40) warrants addition of its external predicate to its internal predicate in *default* positions, as in (41) and the more interesting (42).

(41) $\frac{\text{every } \Phi \text{ is } \Psi}{\text{every } \Phi \text{ is } \Psi\Phi}$ (42) $\frac{\text{some } \Phi \text{ is } \Psi\text{, and every } \Psi \text{ is } \Omega}{\text{some } \Phi \text{ is } \Omega\Psi}$

But first, I want to stress the relevance of (38) and (39) for reducing many valid inference patterns to a small stock of basic principles governing conjunctive predicates.

(38) $\frac{\text{every } \Phi \text{ is } \Omega\Psi}{\text{every } \Phi \text{ is } \Psi}$ (39) $\frac{\text{every } \Phi \text{ is } \Psi}{\text{every } \Omega\Phi \text{ is } \Psi}$

For it turns out that the predicates figure centrally in a rather elegant and arguably natural logic that philosophers of language often forget about; though see Sommers (198x), Parsons (2014).

According to Aristotle, (43) implies that there is at least one Φ;

(43) every Φ is Φ

and perhaps The Philosopher was correct. If (44) likewise implies that there is some $\Psi\Phi$,

(44) every $\Psi\Phi$ is Φ

then (43) and (44) are not strictly equivalent. Still, instances of (44) seem truistic; recall (35).

(35) $\{[\textbf{EVERY}_{\text{Quant}}\ [\text{BROWN}_{\text{Pred}}\ \text{COW}_{\text{Pred}}]_{\text{Pred}}]_{\text{Subj}}\ \textit{IS}_{\text{cop}}\ \text{COW}_{\text{Pred}}\}$

And in any case, (44) can be derived from (45), given (38) and (39).

(45) every Φ is $\Psi\Phi$

I take it that (46) is even more obviously truistic;

(46) some Φ is $\Psi\Phi$ if some Φ is Ψ

---

[13] We can view **SOME** as governed by two default elimination rules, and **NO** as governed by two introduction rules, both of which need to be stated explicitly. But my suggestion is not that contents of quantificational concepts (much less the meanings of corresponding words) are fully characterized by any Gentzen-style rules of inference. For one thing, there are various indefinite quantifiers; **TWO**, **SEVERAL**, etc. And while the contents and inferential roles of logical concepts are clearly related, I take no stand on *how* the content of **EVERY** is related to its effect/non-effect on the default inferential polarities of predicate positions, not even for instances of (40).

and given (46), (21) can be used to derive the intuitively similar (21+).[14]

(21) some Φ is Ψ	(21+) some Φ is Ψ
<u>every Ψ is Ω</u>	<u>every Φ is Ω</u>
some Φ is Ω	some Φ is Ω

This permits reduction (20) and (22), Celarent and Ferio, to (21).

(20) every Φ is Ψ	(22) some Φ is Ψ
<u>no Ψ is Ω</u>	<u>no Ψ is Ω</u>
no Φ is Ω	some Φ isn't Ω

Proofs by *reductio* are sketched below, with line (c) reflecting negations of the conclusions.

(a) every Φ is Ψ		(a) some Φ is Ψ	
(b) no Ψ is Ω		(b) no Ψ is Ω	
(c) some Φ is Ω		(c) every Φ is Ω	
(d) some Φ is ΩΦ	[c, 46]	(d) some Φ is ΨΦ	[a]
(e) some ΩΦ is Φ	[d]	(e) some ΨΦ is Φ	[d]
(f) some ΩΦ is Ψ	[e, a, 21]	(f) some ΨΦ is Ω	[d, c, 21+]
(g) some Ψ is ΩΦ	[f]	(g) no ΨΦ is Ω	[b]
(h) some Ψ is Ω	[h]	(h) ⊥	[f, g]
(i) ⊥	[h, b]		

So given the rules for reducing/adding predicates, Aristotle's syllogisms can be derived from (19) and (21), Barbara and Darii.

(19) every Φ is Ψ	(21) some Φ is Ψ
<u>every Ψ is Ω</u>	<u>every Ψ is Ω</u>
every Φ is Ω	some Φ is Ω

Moreover, (19) and (21) can themselves be viewed as applications of the special license provided by a universal generalization to *add* its external predicate to its internal predicate in default positions. Recall (41) and (42).

(41) <u>every Φ is Ψ</u>	(42) <u>some Φ is Ψ, and every Ψ is Ω</u>
every Φ is ΨΦ	some Φ is ΩΨ

And consider the following chains of reasoning, which are intimately related to (21).

(a) some Φ is Ψ	[premise]	(a) some Φ is Ψ	[premise]
(b) every Ψ is Ω	[premise]	(b) every Ψ is Ω	[premise]
(c) every Ψ is ΩΨ	[b, 41]		
(d) some Φ is ΩΨ	[a, c, 21]	(d) some Φ is ΩΨ	[a, b, 42]
(e) some Φ is Ω	[d]	(e) some Φ is Ω	[d]

The move from (b) to (d) on the right looks like a variant of (21) itself, and in one sense it is. But instead of viewing this step as warranted by a syllogism, we can view it as an application of a conversion licensed by (b): adding Ω to Ψ is permissible, even in a default position where predicate expansion is not licensed in general.[15] Similarly, the chain of reasoning below can be viewed as an elaborated—or perhaps fully spelled out—variant of (19).

---

[14] Given the first premise of (21+) and (46), it follows that **some Φ is ΨΦ**. Given the second premise of (21+) and (39), it follows that **every ΨΦ is Ω**. So given (21), it follows that **some Φ is Ω**.

[15] If it helps, think of (42) as a limited rule of "existential instantiation" that permits conversion of 'Φa & Ψa' into 'Φa & Ωa' given (b).

(a) every Φ is Ψ [premise]
(b) every Ψ is Ω [premise]
(c) every Φ is ΨΦ [a, 41]
(d) every Ψ is ΩΨ [b, 41]
(e) every Φ is ΩΨΦ [c, d]
(f) every Φ is ΩΨ [e]
(g) every Φ is Ω [f]

Here, line (c) reflects the very license that (a) provides: Ψ can be added to Φ, even in a default position where predicate expansion is not licensed in general. Likewise, (d) reflects the license that (b) provides: Ω can be added to Ψ. Line (e) simply reflects the fact that these licenses are *agglomerative*: if you can add Ψ to Φ, and you can add Ω to Ψ, you can do both. Not all permissions are agglomerative. Given diet or manners, it may be permissible to have cake, and permissible to have pie, but not permissible to have both. Given law and an empty gun, it may be permissible to load the gun with bullets, and permissible to be pull the trigger, but not permissible to do both (in that order). Put another way, some permissions/obligations are *defeasible*; and as Horty (201x) shows, there may be governed by an interesting logic. But one can view the medieval logicians as getting at a less nuanced system in which (47) doesn't imply (48), but (48) does imply (49).

(47) [**EVERY**$_{\text{Quant}}$ BIRD$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ WINGED$_{\text{Pred}}$
(48) [**EVERY**$_{\text{Quant}}$ BIRD$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ ONE-THAT-FLIES$_{\text{Pred}}$
(49) [**EVERY**$_{\text{Quant}}$ [FIFTY-POUND$_{\text{Pred}}$ BIRD$_{\text{Pred}}$]$_{\text{Pred}}$]$_{\text{Subj}}$ *IS*$_{\text{Cop}}$ ONE-THAT-FLIES$_{\text{Pred}}$

This is not especially interesting as logic. But given the intuitive force of Aristotelian syllogisms, which form a *network* that can be reduced to small number intuitively compelling principles, it seems that predicates play a central role in an interesting form of thought that is governed by a "natural logic" that treats predicate-*reduction* as a privileged form of inference. At a minimum, we can draw on this ancient tradition to characterize a notion of predicate that is independent of modern logic and contemporary conceptions grammatical structure. This is what I have in mind when I say that phrasal meanings are instructions for how to build mental predicates, One can imagine animals who are able to generate categorical thoughts, in a limited way, but then acquire a capacity to form predicates in a cognitively productive way—thereby becoming able to construct boundlessly many instances of (40) that are systematically related.

(40) every Φ is Ψ

2.3 Predicates and Relations

It is, however, no part of my view that all thoughts are categorical. On the contrary, I assume that humans and many other animals have thoughts that involve irreducibly relational concepts. We may have many ways of thinking that one thing caused another. But in thinking that a certain flash caused a certain bang, we presumably use a relational concept of causation. Likewise, I assume that we use relational concepts in thinking that one thing is *above* another, that Venus is *between* Mercury and Earth, etc. Let's indicate relational concepts with subscripts that reflect the relevant number of relata—the "arity" of the relation, ignoring tense for simplicity—as follows: ABOVE$_2$, BETWEEN$_3$, etc. And let's assume that a relational concept can combine with the relevant number of singular/denoting concepts to form a complete thought. In section three, I cash this out in terms of Frege's (1892a) notion of *saturation*. Though whatever the details, the point is that there are many noncategorical thoughts, which we can represent with a famliar kind of notation: ABOVE$_2$(VENUS$_{\text{Den}}$, BESSIE$_{\text{Den}}$); BETWEEN$_3$(VENUS$_{\text{Den}}$, MERCURY$_{\text{Den}}$, EARTH$_{\text{Den}}$); etc.

A thinker might be able to use some relational concepts, in limited ways, to form complex predicates like ONE-THAT-IS-ABOVE-BESSIE$_{Pred}$ and ONE-THAT-CAUSED-A-BANG$_{Pred}$. As discussed in section three, such use is not trivial; forming "partially saturated" concepts like ABOVE$_2$(_, BESSIE$_{Den}$) is no small matter. But in any case, these complex predicates are not mere conjunctions of simpler ones. They presumably have relational concepts as constituents.

It may be that some of these relational concepts are reducible to others. Perhaps some analog of CAUSED$_2$(_, _) can be reduced to a cluster of concepts that includes BEFORE$_2$ and a "second-order" concept that can combine with a pair of predicates. One can imagine a mind that uses the classical quantifier **EVERY**$_{Quant}$ to introduce a concept **EACH**$_2$ that can appear in thoughts like **EACH**$_2$(COW$_{Pred}$, BROWN$_{Pred}$). One can also imagine a mind that uses **EACH**$_2$ to introduce **EVERY**$_{Quant}$. And a sophisticated mind might be able to use a second-order concept to introduce a first-order concept HUME$_2$(_, _) that can appear in thoughts like HUME$_2$(FLASH$_{Den}$, BANG$_{Den}$); where the relational concept applies to pairs of things such that for *some* predicates Φ and Ψ, the first thing is Φ, the second thing is Ψ, and every Φ is Ψ. But any such reduction of CAUSED$_2$(_,_) requires a stock of basic relational concepts and the resources required to introduce/define others.

As discussed in chapter two, Frege offered a framework for such introduction/definition. He also showed how a certain kind of thinker could use relational concepts, along with concepts of truth values, to introduce *analogs* of classical predicates. A "deeply relational" thinker, who has no categorical thoughts, can mimic a thinker who has such thoughts and reasons with them in an Aristotelian fashion. And we view classical logic as a (very) limiting case of Frege's far more powerful logic, which was designed to accommodate higher-order relational thought and arguments involving arithmetic *induction*, as opposed to mere syllogisms. But this doesn't show, or even suggest, that humans are deeply relational thinkers.

Frege assumed that we naturally think and talk in a subject-predicate format, and that we need help—of the sort provided by his invented Begriffsschrift—in order to use our rudimentary capacities for relational thought in systematic ways. I think this is basically right: our categorical thoughts are governed by a natural logic that lets us appreciate certain implication relations among predicates; but our relational concepts are related in less systematic ways. We use relational concepts in natural modes of thought; but we do not reason with such concepts in the manner of an ideal Fregean thinker. Indeed, it seems that given relational concepts like BETWEEN$_3$, we go the trouble of making predicates like BETWEEN-MERCURY-AND-EARTH$_{Pred}$. As discussed in chapter five, many facts suggest that we use human languages as tools for connecting relational concepts with formally distinct but analytically related predicates that we reason with in natural systematic ways. In my view, event analyses often reflect this.

If (6) is used to assemble and expression a thought of sort indicated with (6a),

(6) It rained yesterday.

(6a) [**SOME**$_{Quant}$ THING$_{Pred}$]$_{Subj}$ *WAS*$_{Cop}$ [RAIN-EVENT$_{Pred}$ YESTERDAY-EVENT$_{Pred}$]$_{Pred}$

then 'yesterday' corresponds to a predicate that can be validly eliminated. And perhaps (50)

(50) Sadie chased Bessie.

is used to assemble and expression a thought of sort indicated with (50a),

(50a) [SADIE$_{Den}$]$_{Subj}$ *WAS*$_{Cop}$ CHASER-IN-A-CHASING-OF-BESSIE$_{Pred}$

even if the predicate in (50a) is formed by using the relational concept in (50b).

(50b) CHASED$_2$(SADIE$_{Den}$, BESSIE$_{Den}$)

I'll have much more to say about this below. But to repeat, my view is not that all thoughts are categorical. On the contrary, I think humans use meanings to *make* predicates from ingredients

that include many relational concepts. And while I think classical logic has a natural basis that Frege transcended, I don't think classical logic provides an adequate theory of natural thought.

Towards the end of the eighteenth century, Kant could say (without much exaggeration) that logic had followed a single path since its inception, and that 'since Aristotle it has not had to retrace a single step'. But he just wrong to say that classical logic was 'to all appearance complete and perfect'.[16] As the medieval logicians knew, relations presented difficulties. For example argument (51) is intuitively valid.

(51) Every horse chased some cow.
<u>Every cow is brown.</u>
Every horse chased some brown cow.

One wants to say that the second premise licenses extension of 'cow' (in the first premise) into 'brown cow', even though 'cow' is part of a quantifier that is in turn part of a complex predicate that has been combined with a quantificational subject, along lines suggested by (51a).

(51a) [(every Φ) is (one that R some Ψ)]
<u>[(every Ψ) (is Ω)]</u>
[(every Φ) is (one that R some Ω^Ψ)]

where 'R' somehow corresponds to relations, but (51a) is not taken as a *basic* form of valid inference. But this highlights the need for a systematic way of accommodating relational notions, without spoiling the idea that quantifiers determine the inferential polarity of the predicates they combine with; see note (12) above.

By itself, (51) is not so troubling, since one can say that the second premise provides a general license to replace 'cow' with 'brown cow. But consider (52).

(52) <u>Every horse that saw some *dog* chased some cow that some *yellow dog* saw.</u>
Every horse that saw some *yellow dog* chased some cow that some *dog* saw.

Each occurrence of 'dog' appears in a constituent of the form 'some Φ'. Though in the first case, 'dog' can be validly replaced with the extended predicate 'yellow dog'; here, the default direction of inferential polarity has been inverted. And of course, (52) is asymmetrically valid, in that its premise does not follow from its conclusion. But it's not that the "main" quantifier determines inferential polarity, as shown by the equally valid (53) and (54).

(53) <u>Some horse that saw every *dog* chased a cow that every *dog* saw.</u>
Some horse that saw every *yellow dog* chased a cow that every *yellow dog* saw.

(54) <u>Every horse that saw no *yellow dog* chased a cow that no *dog* saw.</u>
Every horse that saw no *dog* chased a cow that no *yellow dog* saw.

Note that in the premise of (54), 'yellow dog' has the default inferential polarity—predicate reduction is licensed—as if the inverting effects of 'no' and 'every' cancel out.

To account for such facts, as opposed to merely listing them, one needs to do more than simply allow for predicates of the form 'Φ that R no ΩΨ' and 'R some Δ that no Φ R'. Even bracketing the further complexities illustrated with (55),

(55) Every farmer who owns a horse feeds it, himself, and his children.

[16] It is worth remembering that Frege disagreed with Kant about the nature of logic. They agreed that arithmetic could not be reduced to Aristotelian syllogisms.

we need a systematic account of how predicates can have quantifiers and relational concepts as constituents, in ways suggested by the relative clauses in (52-54). One way or another, this requires expansion of the classical conception of thoughts. Predicates can be complex, in ways that are not merely conjunctive, without containing complete *thoughts* as constituents. Relational concepts can figure systematically in ways of generating cognitively productively predicates.

3. Composition by Adicity

Frege (1879, 188x, 1892a) helped a lot. But it is worth going slowly here—even, and perhaps especially, for those who have seen Frege's notation used to theorize about words—because Frege was offering a proposal about the concepts and overt symbolism we *ought* to use in scientific inquiry. This proposal was intimately connected with Frege's attempt to reduce arithmetic to a logic that treated monadicity as a special case of polyadicity involving relations that countable things can bear to truth values; see chapter two. So as Frege stressed, ordinary concepts/words need not and apparently do not conform to his ideal typology. One can speculate that *pace* Frege, his proposal was prescient cognitive science. My own view is that while Frege's insights about conceptual architecture deserve a central place in psychology, we can live without truth values as relata, and natural lexical items do not exhibit Fregean types. Though in any case, theorists need to distinguish (i) claims about the concepts that are naturally available to humans, without special training or the invention of formal languages like Frege's Begriffsschrift, from (ii) claims about the concepts that theorists ought to deploy for certain intellectual purposes.

3.1 Thoughts, Saturaters, and Unsaturateds

One of Frege's leading ideas was that complex concepts, including thoughts as special cases, can be analyzed as the results of *saturating* some concepts with others.[17] Saturated concepts vary, as do their saturaters. But denoting concepts and complete thoughts cannot be saturated, while general and quantificational concepts are saturatable. To encode this crucial asymmetry, between thought components that get saturated and those that saturate others, Frege used notation that was designed to represent functions and their arguments: a monadic concept like COW(_) can be used to think about some things in a general way, while a denoting concept like BESSIE can be used to think about something as a particular; and combining COW(_) with BESSIE yields a complete thought, COW(BESSIE), according to which the individual denoted is one of the things thought about more generally. One can also view COW(_) as a "mental sentence frame" that is the result of abstracting away from the singular concept in COW(BESSIE).

Let's say that thoughts are representations of a special type <**t**>, and that denoting concepts are "entity-labels" of type <e>. Monadic concepts can then be described as instances of type <e, **t**>—unsaturated representations that can be "filled" by a singular concept to form a complete thought. Since COW(BESSIE) is of type <**t**>, COW(_) is of type <e, **t**>. As stressed below, this does not yet require appeal to truth *values*, much less the Fregean idea that each ideal thought *denotes* a truth value. The typology here concerns concepts, not any posited denotations of ideal concepts. I assume that the concept VULCAN can be of type <e> without there being any entity such that the concept applies to it, and that COW(BESSIE) can be of type <**t**> without there being any entity such that the thought applies to it.

17 The terminology in this field (intentional/intensional, concept/conception, etc.) conspires against us. Frege used 'Begriff' ('Concept') to talk about functions from entities to truth values; see §1 of chapter two. And given his concern with ideal representables, rather than actual representations, his notion of analysis was partly normative; cp. Quine's (1963) notion of regimentation. But in my view, and I believe Frege's, natural concepts can exhibit a saturated/unsaturated distinction without the unsaturated concepts actually representing (or being) Begriffs.

We can go on to describe a concept of negation that can be saturated by one thought to form another: ~[_], of type <**t**, **t**>. Such a concept is formally monadic; though intuitively, it serves to classify thoughts with regard to correctness, rather than classifying entities in terms of some property or condition like being a cow. For example, ~[COW(BESSIE)] is to be classified as true if and only if COW(BESSIE) is to be classified as false. We can also describe a concept of existential closure, ∃[Φ(_)], that can be saturated by a monadic concept to form a complete thought. Here, 'Φ(_)' is a placeholder for concepts of type <e, **t**>. So ∃[Φ(_)] is a concept of type <<e, **t**>, **t**> that serves to classify monadic concepts according to whether or not they apply to something. For example, ∃[COW(_)] is to be classified as true if and only if COW(_) is such that it applies to something—i.e., iff there is at least one cow. So '∃' does not correspond to the classical **SOME**$_{\text{QUANT}}$; it is more like [**SOME**$_{\text{QUANT}}$ **THING**$_{\text{PRED}}$]$_{\text{SUBJ}}$.[18]

I'll return to the role of variables in representations like COW(X) and ∃X[COW(X)]. But let's distinguish the mere idea of an unsaturated conceptual slot, as it figures in a Fregean taxonomy of concepts, from the idea of a slot that is linked to another via some variable.

3.2 Encoding Polyadicity

I am rehearsing familiar claims about concepts with meaning aforethought. It is widely held that human linguistic meanings recapitulate key aspects of the Fregean typology. For example, names are often treated as expressions of type <e>, as if 'Bessie' is a pronounceable denoting concept; and the semantic effect of combining a name with a predicate is often described in terms of saturation. By contrast, I will eschew the usual semantic typology for linguistic expressions, and describe meaning composition in terms of conjunction. But I grant that thinkers often saturate one concept with another. Indeed, on my view, humans use concepts of various Fregean types—including polyadic concepts—to introduce conjoinable monadic concepts. And since I happily appeal to polyadic concepts, I need to represent them.

Frege offers two ways of representing a polyadic concept: as a concept that can be *fully* saturated by one ordered *n*-tuple of singular concepts; or as a concept that can be *partly* saturated by one denoting concept. For many purposes, this distinction does not matter. But it is worth being explicit about both options, especially in the context of explaining and evaluating the hypothesis that human linguistic meanings are concept assembly instructions.

Let's say that the dyadic concept ABOVE(_, _) can be used to think about one thing being above another, and that this concept can be partly saturated by a singular concept to form a complex concept like ABOVE(_, BESSIE); where this monadic concept can be used to think about things above Bessie. The dyadic concept is of type <e, <e, **t**>>; saturating it with a concept of type <e> yields a concept of type <e, **t**>. Let ABOVE(⟨_, _⟩) be a concept that can be saturated by an *ordered pair* of denoting concepts to form a thought like ABOVE(⟨VENUS, BESSIE⟩). Here, we needn't worry about whether such saturation requires a *concept of* ordered pairs. The idea is that ABOVE(⟨VENUS, BESSIE⟩) has a monadic constituent whose saturater is complex in a way that

---

18. While denoting concepts cannot be saturated, predicative concepts can saturate "higher order" concepts of type <<e, **t**>, **t**>. I use regular angle brackets for conceptual types and steeper brackets for ordered pairs. So ⟨<e>, <**t**>⟩ is an ordered pair of types that can be identified with the set {<e>, {<e>, <**t**>}}, just as {1, {1, 0}} can be identified with the ordered pair of numbers ⟨1, 0⟩. One can describe a function F with a domain/range of two entities, ⊤ and ⊥: F(T) = ⊥, and F(T) = ⊥. One can also stipulate that **TRUE** and **FALSE**, special concepts of type <**e**>, denote ⊤ and ⊥. Then **TVC**(_) can be a "truth value conversion" concept of type <**e**, **e**> such that: **TVC**(**TRUE**) is a way of thinking about ⊥, and **TVC**(**FALSE**) is a way of thinking about ⊤. But if natural thoughts are not concepts of truth values, then **~[False]** is conceptual gibberish, as is **TVC**(COW(BESSIE)). For while ~[_] and **TVC**(_) are both negation-concepts, the first requires saturation by a thought, and the second requires saturation by a concept of a truth value.

reflects the relational character of the thought. We can say that ABOVE(⟨_, _⟩) is of type <e^2, **t**>. But whatever our notation, we can stipulate that the formally distinct concepts are cognitively equivalent: for any concepts α and β of type <e>, ABOVE(α, β) iff ABOVE(⟨α, β⟩).

Still, for some purposes, the formal distinction might matter. One can imagine a mind that can assemble the thought ABOVE(⟨VENUS, BESSIE⟩) *without* being able to form ABOVE(⟨_, BESSIE⟩). Such a mind might need to saturate ABOVE(⟨_, _⟩) all at once, with an ordered pair of concepts, because it cannot generate partly saturated structures like ⟨_, BESSIE⟩. One can also imagine a thinker who *can* form ABOVE(⟨_, BESSIE⟩) and ABOVE(⟨VENUS, _⟩): concepts that can be used to think about, respectively, things above Bessie and things Venus is above. This is the kind of thinker Frege had in mind. Likewise, Frege was not worrying about a mind that could form ABOVE(_, BESSIE) but not ABOVE(VENUS, _). For his idealized purposes, ABOVE(_, _) and ABOVE(⟨_, _⟩) count as the same concept depicted in two slightly different ways: one that highlights the relational character of the concept; and one that highlights the possibility of representing relationality in terms of saturation by *n*-tuples. But if part of our task is to explain the unbounded productivity of human concepts, then the human capacity to form concepts like THING-VENUS-IS-ABOVE is one of our explananda. So we need to think about the operations required to *build* concepts like ABOVE(VENUS, _) from atomic constituents.

Perhaps we can build ABOVE(VENUS, BESSIE) and extract BESSIE. But then one wants to know why the result is a saturatable representation that can be used to think about things that Venus is above, as opposed to mere residue of a destructive operation. (Extracting metals from ores, for example, leaves a mess.) I grant that concept abstraction is often possible. But it seems unlikely that *every* potential extraction/abstraction from ABOVE(VENUS, BESSIE) yields a concept that humans can naturally form: __(VENUS, BESSIE) would be of type <<e, <e, **t**>>, **t**>; __(VENUS, _) would be of type <<e, <e, **t**>>, <e, **t**>> or <e, <<e, <e, **t**>>, **t**>>. One can stipulate that any thinker who has ABOVE(_, _) and VENUS can form ABOVE(VENUS, _). But then it is an empirical question whether ordinary humans acquire concepts like ABOVE(_, _), as opposed to simulacra that allow for a more limited kind of composition/abstraction.

This point multiplies given other polyadic concepts. The technical and empirical details, which turn out to be interesting, are for later chapters. But I assume that humans and other animals have triadic concepts like GIVE(_, _, _) and BETWEEN(_, _, _)—or GIVE(⟨_, _, _⟩) and BETWEEN(⟨_, _, _⟩)—that can be used to think about things that exhibit a certain ternary relation. Such concepts can be described as instances of type <e, <e, <e, **t**>>> or <e^3, **t**>. Somehow, humans can form concepts like THING-VENUS-GAVE-TO-BESSIE, THING-VENUS-GAVE-BESSIE-TO, and THING-THAT-GAVE-BESSIE-TO-VENUS. We can even think about pairs of things such that the first gave Bessie to the second. Ignoring tense for simplicity, we can form concepts equivalent to: GIVE(VENUS, _, BESSIE); GIVE(VENUS, BESSIE, _); GIVE(_, BESSIE, VENUS); GIVE(_, BESSIE, _); etc. But children may *start* with composable concepts that are *less* cognitively productive.

Perhaps even adults have only limited capacities for partially saturating natural concepts. We may have VENUS, BESSIE, and GIVE(_, _, _) but no capacity form to GIVE(VENUS, _, BESSIE). We may instead have a capacity to generate a formally distinct concept—built conjunctively from concepts like GIVE(_) and AGENT(_, VENUS)—that is equivalent for many purposes. And we may think about one thing *giving* a second to a third without thinking about the three things as exhibiting a ternary relation. But this is getting ahead. In this chapter, my goal is to explain and motivate the idea that phrasal meanings are instructions for how to build mental predicates, which can also be described as monadic concepts. So let me return to discussing adicity.

3.3 Analysis, Polyadicity and Junction

Among natural concepts, there may be a highest usable adicity. I have no idea what the limit is. But tetradic concepts may be common. Think about the vertices of a square, or about the difference between giving and selling. One can speculate that the concept SELL(_, _, _, _)—with a slot for what a buyer transfers to a seller in exchange for a thing bought—is complex, and that its component parts all have lower adicities. One can also claim that TRANSFER(_, _, _) is analyzable in terms of dyadic and monadic concepts of causation, things, events of motion, and relational states of "having" things. But I don't know how to provide such analyses.[19]

On the other hand, we face unpleasant questions if the verb in 'Venus sold Bessie' expresses a tetradic concept. Most saliently, how can 'sold' combine with *two* grammatical arguments to form a sentence that can be used to express a complete thought? Part of my goal is to answer such questions, without positing covert grammatical arguments, by saying that humans use concepts like SELL(_, _, _, _) to *introduce* monadic concepts of events like SELL(_), which can be combined with concepts like AGENT(_, _) and PATIENT(_, _); where these dyadic concepts, used to think about participation relations that events can bear to their participants, can be linked to grammatical relations that verbs can bear to subjects and objects. This raises questions, addressed below, about *how* concepts like SELL(_) and PATIENT(_, _) are combined. But in my view, we should resist the idea that SELL(_, _, _, _) is composed of concepts that have lower adicities, and say that SELL(_) is introduced in terms of concepts that have higher adicities. Likewise, I claim, for GIVE(_, _, _)/GIVE(_) and CHASE(_, _)/CHASE(_).

As Frege recognized, introducing atomic symbols can be useful when symbols can be manipulated in ways that depend on their formal properties. Correlatively, "fruitful analyses" can reflect introduction of formally useful concepts, as opposed to decomposition of extant concepts; see Horty (2007). In chapter two, I discuss Frege's illustration in terms of arithmetic concepts like NUMBER(_), THREE, and PRECEDES(_, _). But utility depends on what you want. Frege wanted to reduce arithmetic to logic. So he invented a new logic that was geared for relation notions. For other purposes, one might want to exploit a natural logic of predicates.

Suppose that BROWN(_) can be joined with COW(_) to form BROWN^COW(_), a complex predicate that can combine with a subject to form a thought while being used to think about things as cows *and* as brown. Imagine the unfilled slot of BROWN(_) being identified with that of COW(_), making the resulting junction monadic.[20] Call this mode of combination "M-junction,"

---

19 Like Fodor and Lepore (2002), I think extant proposals fail if offered as decompositions of natural concepts. Even if transferring y to z is a matter of causing y's possession of z, I don't think humans have a concept DIFF such that the concept with which we think about transferring y to z has constituents that include CAUSE, POSSESS, and DIFF. One can *posit* a concept that distinguishes transfers from mere causings-to-have. But this is not yet to *provide* a plausible analysis, much less a plausible adicity-reducing analysis. And as we'll see in chapters two and five, accounting for "one-way" implications (e.g, kill/die) does not require *decomposition* if concepts can be introduced.

20 If it helps, think of BROWN^COW(_) as the result of *ad*joining BROWN(_) to COW(_). In any case, '^' does not signify mere concatentation of mental symbols. To be a complex *concept*, BROWN^COW(_) has to differ from {BROWN(_), COW(_)} and be relevantly like COW(_): saturatable by Bessie, a potential saturater of ∃_[Φ(_)], and joinable with other concepts. Kant and Frege stressed that a thought is not a mere list of ideas, echoing Plato's observation that facts are not mere lists of objects. Likewise, a junction is not a mere list of concepts. Indeed, the notation leaves room for the idea that BROWN^COW(_) is the result of saturating a *concept* of conjunction, as in +[BROWN(_), COW(_)]; where +[Φ(_), Ψ(_)] is a concept of type <<e, **t**>, <<e, **t**>, <e, **t**>>>. This would make conjunction an "indirect" mode of conceptual combination that operates via the "direct" mode of saturation. But one can also imagine a primitive joining operation that yields BROWN^COW(_), given COW(_) and BROWN(_) as inputs; where COW(_) can be derived from BROWN^COW(_) in "default" contexts, as discussed in section two. Similarly,

without yet assuming that humans employ it. Given some atomic monadic concepts, M-junction can generate endlessly many more. But this kind of recursion is pretty boring. As noted above, the number of atomic monadic concepts puts a bound on the *cognitive* productivity of conjoining them. Generating COW^BROWN^COW(_), and so on, is a cheap trick that earns little applause.

Let me note, digressively, that mere recursive saturation is similarly soporific: ~[COW(BESSIE)], ~[~[COW(BESSIE)]], ~[~[~[COW(BESSIE)]]], etc. Concepts like *DENY*[_, _]—saturatable by a *thought* like COW(BESSIE) and a singular concept like VENUS to form a complex thought like *DENY*[VENUS, COW(BESSIE)], according to which Venus denied that Bessie is a cow—do yield a kind of unbounded cognitive productivity: *DENY*[BESSIE, COW(VENUS)], *DENY*[VENUS, *DENY*[BESSIE, COW(VENUS)]], etc. But human thought is unboundedly cognitively productive even setting aside concepts with which we think about relations that thinkers bear to representations. And a capacity to saturate concepts does not yet explain how we form concepts like COW-THAT-CHASED-VENUS, COW-VENUS-CHASED, COW-THAT-DENIED-VENUS-IS-A-COW, COW-THAT-VENUS-GAVE-TO-BESSIE, etc. Intuitively, construction of these concepts requires conjunction of COW(_) with complex monadic concepts: CHASE(_, VENUS); CHASE(VENUS, _); *DENY*[_, COW(VENUS)]; GIVE(VENUS, _, BESSIE); etc. But to reiterate a point from §2.2, part of the task is to explain the human capacity to form concepts like THING-VENUS-GAVE-TO-BESSIE.

One can say that given any polyadic concept with *n* unsaturated "slots," a human thinker can use *n*-1 saturaters to create a monadic concept, leaving any one of the slots unfilled. But that leaves the question of how we came to have this impressive capacity. And in chapter five, I offer evidence that a simple form of conjunction lies at the core of unbounded cognitive productivity. Our natural capacities to combine concepts are impressive, but constrained in ways that suggest less than an ideal Fregean mind. I think we make clever use of monadic concept conjunction, as medieval logicians suspected, along with a few other tricks: the introduction of concepts like GIVE(_); appeal to thematic concepts like AGENT(_, _), PATIENT(_, _), and RECIPIENT(_, _); a *very* limited operation of existential closure; and one other operation, described in chapter six, that is required on any account. For specialists, let me say at once: the additional operation is a mental analog of relative clause formation, but *much* less powerful than lambda abstraction, and limited to monadic concepts of a special sort. The idea will be that with help from human language syntax, we can generate an analog of GIVE(VENUS, _, BESSIE) without saturating GIVE(_, _, _)—much less saturating it twice, or thrice and then desaturating once.

Before the digression/advertisement, we were considering the old idea that thinkers can join concepts, thereby forming conjunctions like BROWN^COW(_). And to recapitulate: a mere capacity to join monadic concepts cannot deliver THING-VENUS-GAVE-TO-BESSIE; though neither can a mere capacity to saturate concepts, without a further capacity (that humans may not have) to extract any argument from a multiply saturated concept. Nonetheless, I think a variant of the old idea can help explain both the cognitive productivity of human concepts and the semantic character of human linguistic expressions, at least if the variant is embedded in a broader proposal according to which meanings are concept assembly instructions.

4. Add a *Little* Dyadicity

To get beyond M-junction, which can only generate concepts like BROWN^COW(_), I posit a joining operation that allows for conjunction of concepts like PATIENT(_, _) and COW(_). But the posited operation allows for only a smidgeon of polyadicity. An *atomic dyadic* concept can be

---

saturating one concept with another does not require a concept of saturation. But for now, it doesn't matter if BROWN^COW(_) really has three conceptual constituents, like +[BROWN(_), COW(_)].

combined with a monadic concept subject to a severe constraint: the *first* slot of the dyadic concept must link to the monadic concept, which must be existentially closed as shown below.

```
                    _________
                   |         |
∃[PATIENT(_, _)^COW(_)]
|_____________________|
```

This yields a complex monadic concept whose sole open slot is the *second* slot of the dyadic concept—counting from the right, so that the first (or "most internal") slot of any predicative concept is the one that would be saturated first. Recall that saturating PATIENT(_, _) with BESSIE would yield PATIENT(_, BESSIE), which is a concept of things that have Bessie as their patient. Likewise, the concept indicated above applies to something if and only if its patient is a cow.

4.1 Very Limited Conjunction

The idea is to characterize a form of conjunction that recursively generates monadic concepts, and no others, but *temporarily* tolerates an atomic dyadic conjunct. In terms of the example: the first slot of PATIENT(_, _) is linked to the adjacent monadic conjunct, as if two monadic concepts were being M-conjoined; but the presence of a second slot requires closure of the first, since conjunctions must be monadic. Put another way, while the posited system permits conjunction of an atomic dyadic concept with a complex (i.e., generated) concept, this system cannot conjoin two polyadic concepts of any kind. So no generated conjunct can be dyadic.

As discussed in section three, conjoining CHASED(_, _) with SAW(_, _) would require a determination of how the four conceptual slots are to be linked or left unlinked. And it is all too easy to stipulate that polyadic concepts can be combined in productive ways, as if humans are born with ideal Fregean minds. So I want to start by characterizing a recursive form of conjunction that generates monadic concepts, in a way that goes beyond M-junction, but *without* assuming that polyadic concepts can be conjoined subject to further constraints that need to be theoretically motivated and biologically implemented. More specifically, I want to ask if a variant of M-junction that allows for a smidgeon of polyadicity can serve as the core recursive operation of semantic composition for human languages and thereby help *explain* human cognitive productivity. To operationalize 'smidgeon', I am restricting polyadicity to a single dyadic conjunct that must be (i) atomic rather than generated, and (ii) immediately made part of a complex monadic concept, as opposed to leaving both slots open to further processing.

Call the envisioned form of conceptual combination "Θ-junction," as a reminder that the dyadic concept can be thematic and indicated by a *grammatical relation*.[21] For some purposes, it would do no harm to represent results of Θ-junction with an ampersand and variables, as in ∃X[PATIENT(E, X) & COW(X)]. But as we'll see, given standard interpretations for the ampersand and variable, the concept ∃X[PATIENT(E, X) & COW(X)] has sophisticated constituents that support a kind of cognitive productivity that ordinary humans may not enjoy. Note that ∃X[COW(X) & CHASED(X, VENUS) & GIVE(VENUS, X, BESSIE) & BROWN(BESSIE)] may not be a natural concept. So the linking notation, used above, is not a mere typographic alternative to using variable letters and an ampersand. The hypothesis is that humans can employ a certain mode of concept junction that permits a *little* more than M-junction. The links indicate how the posited form of junction

21 One can characterize other "minimally nonmonadic" forms of junction. But Θ-junction has an obvious appeal for neo-Davidsonian semanticists who regularly appeal to event predicates and thematic relations. Cp. Higginbotham's (1985) discussion of theta-binding and theta-linking.

(deterministically) accommodates a dyadic input. The links do not indicate a particular choice about how to use three independent conceptual slots in forming a complex concept.

I'll say more about this, and the posited form of closure, in section five. But since the linking notation is typographically inconvenient, let me depict the Θ-junction of COW(_) with PATIENT(_, _) as follows: ∃[PATIENT(_, _)^COW(_)]. Despite the absence of variable letters, this is unambiguously a monadic concept that applies to an entity—with events as special cases of entities—if and only if that entity has a cow as its patient: the first/rightmost slot of PATIENT(_, _) must link to *the* slot of COW(_), which must be the target of closure. In general, the Θ-junction of a monadic concept Φ(_) with a dyadic concept Δ(_, _) is ∃[Δ(_, _)^Φ(_)], which applies to an entity if and only if Δ(_, _) applies to the pair consisting of that entity and something to which Φ(_) applies.

One can describe Θ-junction as an operation that yields a concept of type <e, **t**> given inputs of type <e, <e, **t**>> and <e, **t**>. But this way of encoding semantic typology is misleading if all nonmonadic inputs to composition operations are dyadic and atomic. Even if one takes Θ-junction to be a special case of saturation, since a dyadic concept combines with something to yield a monadic concept, this operation does not support "adicity reduction" for concepts of any other type. Again, Θ-junction does not permit supradyadic conjuncts like BETWEEN(_, _, _) or conjunctions of dyadic concepts. Indeed, Θ-junction does not even permit junction of PATIENT(_, _)^COW(_) with a monadic concept: VIOLENT(_)^[PATIENT(_, _)^COW(_)] is not generable. The envisioned system cannot create a complex dyadic concept by linking the unsaturated slot of VIOLENT(_) to one of the other slots; nor can it create a triadic concept by leaving the new slot unlinked. Only a little dyadicity is allowed, and none that percolates.

One can view '^' and '∃' as together indicating a conjunctive operation that can be described as bipartite, leaving it open whether the implemented operation is divisible into junction and closure components. Humans may enjoy a slightly flexible mode of junction that always yields a monadic concept, but one whose *form* depends on the inputs: given Φ(_) and Ψ(_), the result is Φ^Ψ(_), in which the slot of Φ(_) is identified with the slot of Ψ(_); given Δ(_, _) and Φ(_), the result is ∃[Δ(_, _)^Φ(_)]. Alternatively, we may enjoy two related forms of junction, each type-inflexible. (In which case, using '^' for both M-junction and Θ-junction could be misleading.) It is hard to tell the difference, in part because one might analyze concepts like BROWN-COW as Θ-junctions like ∃[BROWN(_, _)^COW(_)], with BROWN(_, _) applying to an ordered pair ⟨α, β⟩ if and only if α is: *brown and identical to* β; or *the predominantly brown surface of* β; or *identical to β and brown for a thing of the specified sort*; etc.

With those caveats entered, I'll use '^' to encode both M-junction and Θ-junction. Thus, ∃[PATIENT(_, _)^[BROWN^COW(_)]] applies to an entity if and only if that entity has a brown cow as its patient. Likewise, CHASE(_)^∃[PATIENT(_, _)^COW(_)] applies to an entity if and only if it is a chase that has a cow as its patient. And to flag a more interesting case, discussed later,

∃[PATIENT(_, _)^[ARRIVAL(_)^∃[PATIENT(_, _)^[BROWN^COW(_)]]]]

applies to an entity if and only if it has *an arrival whose patient is a brown cow* as its patient.[22] If one wants to make the contrast between M-junction and Θ-junction more vivid, one can

22 For simplicity of *initial* exposition, pretend that all "internal" arguments of verbs correspond to PATIENT(_, _). Chapters five and six address the complications concerning which thematic concepts correspond to which verbs. For example, things that arrived may be *themes* of arrivals, even if things kicked/chased are patients of kicks/chases. And arrivals that are seen may be themes of perceptual events that have *experiencers* rather than agents.

recode M-junctions as follows: BROWN^{+}COW(_). But I'll want to highlight a potential parallel with predicate-adjunct and predicate-argument combination in human language syntax: while adjunction differs from complementation, expressions are combined in both cases, suggesting some common mode of syntactic combination that can be supplemented; see Chametsky (19xx), Hornstein and Pietroski (2009), Hunter (2011). Let me defer such details, though, and finally get back to clarifying the idea of concept assembly instructions.

4.2 Fetch, Join, Repeat

If biology somehow implements M-junction and Θ-junction, one can envision a mind with further capacities to (i) use lexical items to *access* concepts that can be inputs to these operations, and (ii) combine lexical items in ways that trigger *assembly* of concepts. Suppose that a "merging" two expressions (lexical or phrasal) is an instruction to send a pair of corresponding concepts to a "joiner," whose output can be an input to further operations of joining, perhaps via storage in some form of short term memory. And suppose that a second way of combining two expressions, "thematizing," sends a corresponding pair of concepts—one monadic but perhaps complex, the other dyadic and atomic—to the Θ-joiner, whose monadic output can also be an input to further operations. This would yield a kind of recursion that goes beyond M-junction.

As a concrete illustration, imagine a mind—call it Joyce—that has some lexical items, each with a long-term address that may be shared by two or more polysemously related concepts; see chapters five and six. In addition to having lexical items, Joyce can produce and execute instructions like *fetch*@'cow'; where for any lexical item L, an instruction of the form *fetch*@L is executed by generating a copy of a concept that resides at the long-term address of L. Joyce has a few memory registers that can be used to briefly store (copies of) concepts. Joyce can also produce and execute instructions of the forms **M-join**[**I**, **I′**] and **Θ-join**[**I**, **I′**]; where **I** and **I′** are themselves generable instructions. An instruction of the form **M/Θ-join**[**I**, **I′**] is executed by executing **I** and **I′**, in some suitable workspace, and M/Θ-joining the results. For example, Joyce form the concept CHASE(_)^∃[PATIENT(_, _)^[BROWN^COW(_)]] by executing the instruction **M-join**[*fetch*@'chase', **Θ-join**[*fetch*@'Θ:pat', **M-join**[*fetch*@'brown', *fetch*@'cow']]].

If it helps, imagine the following sequence of operations.

1. *fetch*@'cow'
2. *fetch*@'brown'
3. **M-join**[1, 2]
4. *fetch*@Θ:PAT
5. **Θ-join**[3, 4]
6. *fetch*@'chase'
7. **M-join**[5, 6]

Let the numbers correspond to reusable memory registers, where results of executing instructions can be stored, and let 'Θ:pat' be a functional item whose address is that of PATIENT(_, _). If SEE(_) and ARRIVAL(_) are stored at the lexical addresses for 'see' and 'arrive', then the monadic concept SEE(_)^∃[PATIENT(_, _)^[ARRIVAL(_)^∃[PATIENT(_, _)^[BROWN^COW(_)]]]] can be similarly assembled by executing the following instruction.

M-join[*fetch*@'see',
  Θ-join[*fetch*@Θ:PAT ,
    M-join[*fetch*@'arrive',
      Θ-join[*fetch*@Θ:PAT, **M-join**[*fetch*@'brown', *fetch*@'cow']]]]]

Executing this instruction yields a concept akin to but distinct from the more sophisticated

SEE(E) & ∃E′[PATIENT(E, E′) & [ARRIVAL(E′)^∃X[PATIENT(E′, X) & BROWN(X) & COW(X)]]]].

For these purposes, instructions can take many forms, perhaps even strings of '1's and '0's that get used—as in a von Neumann machine—to access other such strings and perform certain operations on them. The instruction '1101011010' might be executed by performing *operation six* (110) on whatever *number is stored in* (1) register *twenty-six* (011010), while '0100110101' calls for performing *operation two* (010) on the *number* (0) *fifty-three* (110101). And instead of arithmetic operations that are performed on accessible/generable numbers, one can imagine conjunctive operations that are performed on accessible/generable concepts.

I won't appeal to instructions of the form **Saturate**[**I**, **I′**], which would call for saturating the result of executing one instruction with the result of executing another. But others can hypothesize that CHASE(_, _) and BESSIE are lexically fetchable, and that CHASE(_, BESSIE) is assembled by executing **Saturate**[*fetch*@'chase', *fetch*@'Bessie']. Likewise, EVERY[BROWN^COW(_)]—a concept of type <<e, **t**>, **t**>—might be assembled by executing **Saturate**[*fetch*@'every', **M-join**[*fetch*@'brown', *fetch*@'cow']].[23] The idea of executing Begriffsplans, thereby assembling concepts, is neutral with regard to the atomic concepts and composition operations. But the availability of certain concepts/operations in cognition does not ensure the constructability or executability of corresponding instructions to fetch/combine.

Imagine a mind that can saturate some concepts with others, conjoin certain concepts, etc. If BROWN(_) and COW(_) are the only concepts stored at the corresponding lexical addresses, then even if the instruction **Saturate**[*fetch*@'brown', *fetch*@'cow'] is generable, it will not be executable. Given the lexical limitation, 'brown cow' would not be an executable instruction for a mind that could only generate Begriffsplans of the form **Saturate**[**I**, **I′**]. Though of course, type shifting can extend the utility of a constrained instruction generator. Suppose that BROWN(_) can be used to introduce BROWN[_, Φ(_)]: a concept of type <<e, **t**>, <e, **t**>> that when saturated by a monadic concept Φ(_) yields a concept that applies to an entity if and only if that entity falls under both BROWN(_) and Φ(_). If BROWN[_, Φ(_)] is assigned to the lexical address for 'brown', along with BROWN(_), then one way to execute *fetch*@'brown' is to fetch the higher-order concept BROWN[_, Φ(_)]. In which case, executing **Saturate**[*fetch*@'brown', *fetch*@'cow'] could result in the assembly of BROWN[_, COW(_)]—a concept equivalent to BROWN^COW(_).[24]

This kind of type shifting may or may not be natural for humans. So let's also imagine a mind that can only generate assembly instructions of the form **M-join**[**I**, **I′**] or **Θ-join**[**I**, **I′**]. If CHASE(_, _) and BESSIE are the only concepts stored at the corresponding lexical addresses, then the instruction **M-join**[*fetch*@'chase', *fetch*@'Bessie'] will not be executable, and likewise for **Θ-join**[*fetch*@'chase', *fetch*@'Bessie']. The concepts CHASE(_, _) and BESSIE cannot be joined. But suppose that CHASE(_, _) and BESSIE are used to introduce CHASE(_) and BESSIE(_); where to a first approximation, BESSIE(_) applies to an entity if and only if it is the entity thought about with BESSIE. If CHASE(_) and BESSIE(_) are assigned to the relevant lexical addresses, then executing **M-join**[*fetch*@'chase', **Θ-join**[*fetch*@'Θ:pat', *fetch*@'Bessie']] can be a way to assemble CHASE(_)^∃[PATIENT(_, _)^BESSIE(_)], a concept of chases of a/the Bessie. Similarly,

---

23 Higginbotham (1985) appeals to both theta-linking and theta-binding; see also Larson and Segal (1995). Heim and Kratzer (1998) introduce predicate-modification and predicate-saturation as composition principles.

24 Cp. Parsons (1970), Kamp (1970), Montague (1974). This raises the question of whether *disjunctive* concepts could be introduced, so that 'brown cow' might be roughly synonymous with 'is brown or is a cow'. Schein (2002, forthcoming) discusses related questions—illustrated with '*John and Mary* formed a happy union/make a good couple/are happy together'—in the context of an approach to plurality adopted in chapters three and four.

∃[AGENT(_, _)^DOG(_)]^[CHASE(_)^∃[PATIENT(_, _)^COW(_)]] can be assembled by executing the instruction below, letting the address of 'Θ:ag' be that of the concept AGENT(_, _).[25]

M-join[**Θ-join**[*fetch*@'Θ:ag' ,*fetch*@'dog'],
M-join[*fetch*@'chase', **Θ-join**[*fetch*@'Θ:pat',*fetch*@'cow']]]

There are many details to work out. But my suggestion is that human linguistic expressions pair instructions of this sort with pronunciations. Or put another way, human languages are (biologically implementable) procedures that generate expressions whose meanings *are* concept assembly instructions of a certain sort. Each generable expression has a phonological side (PHON) and a semantic side (SEM), by virtue of which the expression can interface with the articulatory/perceptual systems and intentional/conceptual systems that are employed in the production and comprehension of speech. The idea is that meanings/SEMs can be described as generable instructions for how to assemble concepts. Each meaningful and articulable expression can also be described syntactically—e.g., as a noun phrase headed by 'cow'. Indeed, such description is crucial for theories of the generative procedure.[26] So we want to know how syntactic descriptions are related to SEMs. But a noun phrase can be a concept assembly instruction, as well as an instruction for how to activate articulators. (Noun phrases can be analytically related; they can also rhyme.)

Meaningful expressions can *be* articulable Begriffsplans that have syntactic properties. But such expressions can also *be* syntactic structures that have meanings and articulations, since such structures can pair PHONs with SEMs. On this view, meanings can be as composite as the corresponding concepts. Meanings—like phrases—are not merely specifiable in terms of lexical items and syntax. Instructions like **M-join**[*fetch*@'brown',*fetch*@'cow'] have simpler instructions as *parts*. But individuals can share Begriffsplans without sharing concepts. Speakers might fetch different concepts with 'cow', or articulate this word in slightly different ways, yet still count (at least for many purposes) as pairing the same phonological instruction with the same semantic instruction. We may not have clear criteria for what counts as the same lexical address across speakers; and the vague facts may be determined in part by what fetched concepts are concepts of. But as suggested by polysemy, a single instruction may correspond to multiple concepts.

I hope the analogy to elementary computer programs, which can be compiled and implemented, makes the operative notion of instruction tolerably clear and unobjectionable in the present context. One can raise many questions about the sense(s) in which Begriffsplans are intentional, and which philosophical projects would be forfeited by appealing to unreduced capacities to generate and execute the instructions posited here. But my task is not to reduce linguistic meaning to some nonintentional basis. I'll say more about Begriffsplans in later

---

25 This may require three memory registers, and likewise for **Saturate**[**Saturate**[*fetch*@'every',*fetch*@'dog'], **Saturate**[*fetch*@'chase',*fetch*@'Bessie']]. However semantic composition is described, concept assembly can reflect complexity that is due to more than one slot of a polyadic concept. One expects limits on how much such complexity a natural system can tolerate within one cycle of computation; cp. Chomsky (2005), Boeckx (200x). But 'Every dog that saw a cat chased a cow that saw a rat' may well call for storage of at least one concept constructed by executing a "phase" of the expression; cp. Uriagereka (2002). So it seems safe to assume repeatable use of at least three memory sites in which concepts can be stored as inputs to a later stage of combination.

26 Following Chomsky (1964, 1965) and many others, I think the I-languages that *generate* complex expressions (from atomic expressions) can and should be described in purely formal/syntactic terms, without reference to the interfacing systems. So I suspect that human languages can be described without talk of instructions, phonological or semantic. But in describing a generative procedure as a *language*, one suggests—and I grant—that the procedure can be used in a way that supports communication.

chapters. But in offering an account of how human linguistic meanings combine, I assume concepts and capacities to generate and execute instances of certain instruction types: **M-join**[**I**, **I′**]; **Θ-join**[**I**, **I′**]; and a couple of others to be discussed.[27] For even if meanings are Begriffsplans, semanticists face an amply hard two-part question: what types of concepts can lexical items call for, and what modes of concept composition do phrasal meanings call for? Working out and defending a relatively spare answer is task enough.

For example, since neither M-junction nor Θ-junction provides a way of forming complete thoughts, I need to say something about sentential conjunction, sentential negation, and the very idea of a matrix sentence. Relatedly, one might think that any plausible analysis of quantificational constructions (e.g., 'Fido chased every cow') will appeal to Fregean instructions to saturate concepts. I'll try to show otherwise, drawing on Pietroski (2005a). More generally, I think that we can eliminate appeal to **Saturate**[**I**, **I′**] in favor of appeal to **Θ-join**[**I**, **I′**], and that we cannot avoid appeal to the latter. But in any case, accommodating relative clauses—as in 'dog that chased a cow that chased a dog', etc.—requires a further "abstractive" operation; see chapters two and six. This highlights a methodological point, with which I'll end this section.

Relative clauses reflect a cognitively productive form of recursive combination. They also suggest that sentential expressions can be converted into phrases that correspond to monadic concepts that can be conjoined with others. In light of such clauses, one wants to know—even for sentences with no additional adjectives, adverbs, or prepositional phrases—(i) what forms of conjunction and abstraction will be required by any decent theory that *also* appeals to a form of saturation, and (ii) what forms of conjunction and abstraction are required by a decent theory that *does not* appeal to a form of saturation. Given any posited operation of conceptual composition, one wants to know if it is needed to achieve the kind of productivity illustrated with 'dog that chased a cow that chased a dog'. For if not, then one needs other reasons for supposing that the operation is invoked by a composition instruction. I propose a bundle of operations/instructions that does not include **Saturate**[**I**, **I′**]. On this view, lexical items can only call for monadic or dyadic concepts, because semantic composition is a matter of building monadic concepts from monadic or (atomic) dyadic inputs. We'll see if this turns out to be plausible.

5. Flavors of Conjunction

At this point, let me enter a confessional note. For many years, I used '&' to encode claims about logical forms without asking (much less knowing) what I intended to represent with that squiggle. I knew the ampersand of an invented predicate calculus, as in 'Brown(x) & Cow(x)', was not the one defined in terms of truth tables. And I suspected that understanding 'brown cow' did not require tacit grasp of Tarski's (1933) notion of conjunction—characterized in terms of satisfaction of open sentences, which can have unboundedly many variables, by unboundedly long sequences of entities. But I had no clear alternative in mind when inscribing formulae like '∃e[Stab(e, Brutus, x) & ∃z[With(e, z) & Knife(z)] & Today(e)]]' or 'Stab(e) & Patient(e, x)', in the course of talking about human languages. Mea culpa. My intention is to do better.

The posited concept ∃[PATIENT(_, _)^[BROWN(_)^COW(_)]]—formed by M-joining BROWN(_) with COW(_), and Θ-joining the result with PATIENT(_, _)—may be a figment of theoretical imagination. But it is not ∃X[PATIENT(E, X) & [BROWN(X) & COW(X)]] in disguise. My proposed modes of junction may turn out to be unnatural, in that ordinary thinkers do not employ them. Though as we'll see, Tarskian conjunction is unlikely to be more natural, despite its

27 I allow for further reduction without providing any. Perhaps **Θ-join** can be reduced to **M-join** and an instruction to perform a restricted kind of closure. And perhaps the basic instructions are intentional only in a limited sense.

distinguished history. And if only for these reasons, I want to locate M/Θ-junction within a range of possible conjunctive operations. Repeating a point from §2.1, encoding possible concept types helps in specifying which types correspond to human linguistic meanings. A typology of conjunctions will also help clarify my own proposal, indicate some other options, and further illustrate the point that Θ-junction allows for only a smidgeon of polyadicity.

5.1 Propositions and Predicates

For simplicity, let's encode the options in terms of dyadic concepts of conjunction, and say that a conjunctive thought/concept is the product of saturating a concept of conjunction; see note 20 above. We can worry later, if we're not too tired, about converting talk of conjunction-*concepts* into other terms. We have enough to think about now, since in principle, there are endlessly many conjunction-concepts. But let's start with some obvious examples that allow for recursion.

Let **&**(_, _) be a concept of type <**t**, <**t**, **t**>> that can be saturated twice, by a thought each time, to yield a complex thought that is true iff both saturating thoughts are true. For these purposes, ignore any differences between **&**(_, _) and two formally distinct concepts: **&**(⟨_, _⟩), the corresponding concept of type <**t**^{2}, **t**> that can be saturated by an ordered pair of thoughts; and **&**({_, _}), the corresponding concept that can be saturated by an unordered pair of thoughts. Then we can use '**&**' to indicate conjunction of *thoughts*—representations of type <**t**>—abstracting from further formal details. Restricting the ampersand in this way leaves room for the hypothesis that thoughts like VENUS-ROSE-AND-BESSIE-CHASED-VENUS are formed via some other (perhaps more permissive) operation. But one can also hypothesize, not implausibly, that humans and other animals enjoy a conjunction-concept of type <**t**, <**t**, **t**>>.

Let ∩[Φ(_), Ψ(_)] be of the type <<e, **t**>, <<e, **t**>, <e, **t**>>>—or abbreviating '<e, **t**>', the type <e**t**, <e**t**, e**t**>>. When saturated by monadic concepts, ∩[Φ(_), Ψ(_)] yields their M-junction. So the instruction below has the same effect as **M-join**[*fetch*@'cow', *fetch*@'brown'].

Saturate[**Saturate**[*fetch*@'∩', *fetch*@'cow'], *fetch*@'brown']

The notation reflects the old idea that such junctions represent intersections of properties. Again, ∩[COW(_), BROWN(_)] differs formally from ∩[⟨COW(_), BROWN(_)⟩] and ∩[{COW(_), BROWN(_)}]. But let's use the intersection sign to indicate any such conjunction of monadic concepts, without concern about order, leaving room for various hypotheses about if and when humans deploy concepts of type <e**t**, <e**t**, e**t**>>; again, see note 20.

One might balk at this "higher" conceptual type. But **&**(_, _) and ∩[Φ(_), Ψ(_)] are both concepts of types that exemplify the pattern <F, <F, F>>; where F is a Fregean type. Given joinable thoughts, positing a junction-concept of type <**t**, <**t**, **t**>> does not seem like a big leap, and likewise for a junction-concept of type <e**t**, <e**t**, e**t**>> given joinable monadic concepts. Appeal to higher-order relational concepts seems unavoidable, given thoughts like MOST-OF-THE-COWS-ARE-BROWN and FIDO-CHASED-MOST-OF-THE-COWS, often analyzed as follows: MOST[BROWN(_), COW(_)], and MOST[CHASED(FIDO, _), COW(_)]; where the quantificational concept MOST[Φ(_), Ψ(_)] is of type <e**t**, <e**t**, **t**>>.[28]

---

28 In chapter six, I'll say how a dyadic concept like MOST[Φ(_), Ψ(_)] can be used to introduce a monadic analog that can be conjoined with others. For now, just note that unlike COW(_), MOST[Φ(_), Ψ(_)] may well be analyzable. Indeed, there is more than candidate analysis. A concept of one-to-one correspondence—saturatable with COW(_) and BROWN(_) to form ONE-TO-ONE[BROWN(_), COW(_)], according to which the cows correspond one-to-one with the brown things—might be used to introduce ONE-TO-ONE-PLUS[Φ(_), Ψ(_)], permitting construction of a thought according to which: some (but not all) of the cows correspond one-to-one with (all of) the brown things. One could then identify MOST[Φ(_), Ψ(_)] with ONE-TO-ONE-PLUS[Φ(_), Ψ(_)]. Or suppose we have a concept *NUMBER*[Φ(_)] of type <e**t**, *e*>; where here, italics just suggest concepts that let one think about numbers. Given GREATER(*X*, *Y*)—a

That said, ∩[Φ(_), Ψ(_)] differs from **&**(_, _) in a way that merits attention: COW(_) is a monadic concept with a "slot" of its own, and likewise for BROWN(_). Using a link as below,

```
        ________
       |        |
BROWN(_)^COW(_)]
```

or a variable letter as in '∩[COW(X), BROWN(X)]', makes it explicit that the "two" slots count as occurrences of the same unsaturated position in the joined concept. Perhaps such linking is trivial for concepts of type <e, **t**>. But if only because biological AND-gates for representations of type <**t**> may be easier to build than monadic analogs, the distinction between propositional conjunction and (any variant of) monadic concept conjunction is worth noting.

More generally, having concepts of type F does not ensure the availability of a conjunction-concept of type <F, <F, F>>. Perhaps CHASE(_, _) and SAW(_, _) can be joined to form a third concept of type <e, e**t**> that can be indicated with variable letters as follows: ^{2}∩[CHASE(X, Y), SAW(X, Y)]. But there is no guarantee that such junction/linking is possible. We may not have a conjunction-concept of type <<e, e**t**>, <<e, e**t**>, <e, e**t**>>>; cp. Montague (197x). And even if we have a dyadic concept-conjoiner, which might be useful for various cognitive purposes, it may not be accessible for purposes of linguistic comprehension. Similar points apply, with more force, to supradyadic concepts like GIVE(_, _, _) and SELL(_, _, _, _).

4.2 Syncategorematic Conjunction

With this in mind, consider a third notion of conjunction, reflected in formal languages that allow for open sentences like 'Rxy *&* Syzw' and closed sentences like 'Rab *&* Sbcd'. This italicized ampersand, unlike '**&**', can connect sentences that contain arbitrarily many variables (treating constants as special cases); see chapter two. Thanks to Tarski (193x), this kind of conjunction has become part of the logician's standard toolkit. But it has some striking features that may not be exhibited by any natural mode of concept junction.

Note that 'Rxy *&* Syzw' corresponds to a tetradic concept, even though each conjunct corresponds to a dyadic or triadic concept. Because this ampersand connects sentences that are spelled with variable letters, which determine how the predicative slots are linked, 'Rxy *&* Syzw' has a fixed adicity.[29] But since there are no constraints on how variables can reoccur across sentences, the adicity of a Tarskian conjunction is not determined by the adicity of the conjuncts, except in special cases where at least one conjunct has zero variables. While each of the conjuncts in 'Bx *&* Cy' is monadic, the conjunction is as dyadic as 'Rxy'.

One can hypothesize that humans enjoy a concept of conjunction that takes a pair of "pre-linked thought-frames" as inputs; where a *thought-frame* is any (complete thought or)

---

concept of type <*e*, *e***t**> with which one can think about one number being greater than another—a thinker might then introduce a complex concept of type <e**t**, <e**t**, **t**>>: GREATER(*NUMBER*[Φ(_)^ Ψ(_)], *NUMBER*[Φ(_)^~(Ψ(_))]). Actually, this requires a concept of *concept*-negation, while '~' was defined for thoughts of type <**t**>. But instead of appealing to negation of any kind, one might analyze MOST[BROWN(_), COW(_)] as GREATER[*NUMBER*[Φ(_)^Ψ(_)], *MINUS*(*NUMBER*[Φ(_)], *NUMBER*[Φ(_)^Ψ(_)])]; where *MINUS*(*X*, *Y*) is a concept of type <*e*, <*e*, *e*>> with which one can think about a number as the result of subtracting one number from another. While this last option may seem recherché, Pietroski et. al. (2009) and Lidz et. al. (2011) provide psycholinguistic evidence against the first option, and in favor of third over the second. Still, one would like to know why "logical" vocabulary items permit analyses, and why we have no determiner for ONE-TO-ONE[Φ(_), Ψ(_)]; cp. note 24.

29 By contrast, one could stipulate that 'R_ _ & S_ _ _' is well-formed but of no fixed adicity. Depending on how the slots get linked, this conjunction might be triadic, tetradic, or pentadic—or monadic or dyadic if there can be linking within a sentence that has no sentential consitutents, as in 'Rxx'.

predicative concept, and thought-frames are *pre-linked* if their slots are marked—as with variables—in a way that determines which are linked to which, and which are left unlinked. Saturating such a concept (of conjunction) twice would yield another pre-linked thought-frame, though one whose adicity might exceed that of either saturater. We can say that such a concept would be of the type *<PLTF*, *<PLTF*, *PLTF>>*. But this is not a Fregean type: the first conjunct might be of any type; likewise for the second; and the input types need not determine the ouput type. One can say that the inputs are Fregean concepts like ABOVE(_, _) and GIVE(_, _, _) along with a choice of how to link the slots. But this third input, the choice, would not be a saturater in any intuitive sense. Correlatively, "saturating" a concept of type *<PLTF*, *<PLTF*, *PLTF>>* with ABOVE(_, _) and GIVE(_, _, _) would not yield a monadic concept, saturatable by BESSIE; it would yield a representation with no fixed adicity (see note 29).

If we want to talk about a *concept* of conjunction that has these features, I think our notation should highlight its distinctive character. It is potentially misleading to use the same shape used for the relatively boring concept of propositional conjunction, which is of the fixed type <**t**, <**t**, **t**>>. So let's replace '*&*' with '!' and stipulate that ABOVE(X, Y) ! GIVE(Y, Z, W) is a concept with which one can, if one has the concept, think about one thing being above a second that a third gave to a fourth. Given special training and a blackboard, one can invent and perhaps think with such symbolism. But inventing a procedure that can generate 'Ba ! Cb', 'Bx ! Cx', and 'Rxy ! Gyzw' hardly shows that a natural procedure generates analogs of *all* these invented conjunctions if it generates analogs of one.[30] Assuming that human concepts are conjoinable in this Tarskian way is to assume, rather than explain, a very strong form of cognitive productivity.

Questions about the available kind(s) of junction raise corresponding questions about existential closure. The concept ∃[Φ(_)]—of type <e**t**, **t**>—can be saturated by ABOVE(_, BESSIE) or ABOVE(VENUS, _). But ∃[Φ(_)] cannot combine with ABOVE(_, _) or any other concept of type <e, e**t**>; see chapter two for further discussion. We can introduce a more permissive concept of existential closure that can combine with any predicative concept, and bind any variable, to allow for triadic concepts like ∃z[ABOVE(X, Z) ! GIVE(Y, Z, W)]. Perhaps humans can use this concept to think about a first thing being above something which a second thing gave to a third. But perhaps in so far as we naturally do such thinking, we use concepts that exhibit a different kind of architecture. Similar points apply to negation, since by hypothesis, ~[_] is of type <**t**, **t**>.

Put another way, one can learn to see '∃y(Axy ! Gzyw)' as an expression that is satisfied by certain sequences of things—in accord with Tarski's algorithm, discussed in chapter two. And one can hypothesize that humans naturally form such sentences, using a mental analog of '!'. But perhaps theorists should posit sparer conceptual capacities in accounts of meaning for human languages that permit expressions like 'is above something that was given'.

In terms of explanatory order, one can speculate that humans form and use concepts like ∃z[ABOVE(X, Y) ! GIVE(Y, Z, W)] independent of capacities to generate and understand linguistic expressions. Then one might appeal to such concepts in theories of linguistic meaning. Alternatively, one might try to explain linguistic understanding in the sparest possible way—

---

30 Using '*&*' fosters the illusion that the posited species of conjunction is mundane. I assume that human inquirers can invent and come to think with technical concepts whose properties make them unnatural for purposes of lexical acquisition and ordinary human thought. Consider **twix**(**p**, **q**, **r**): **p** and **r** are equivalent with respect to truth—both are true, or both are false—while **p** and **q** are *not* equivalent in this respect. Or **maj**(**p**, **q**, **r**): exactly two of the three thoughts are true. Note that **twix**(**p**, **q**, **r**) **&** ~(**q**) implies **maj**(**p**, **q**, **r**). Compare 'grue' (Goodman [1954]), 'urp' (Quine [1960] on undetached rabbit parts), and nonconservative determiners (Barwise and Cooper [1981], Higginbotham and May [1981]).

positing only the concepts and combination operations required in order to account for the relevant natural human phenomena—and argue that this minimalist approach yields a *better* account of meaning in the following sense: it explains, without sacrificing descriptive adequacy, many facts that are puzzling given richer assumptions about human concepts; cp. Chomsky (1995). Then theorists are free to use the posited concepts and spare combination operations in explaining the unbounded cognitive productivity of human concepts. The latter approach is more ambitious. But I'll argue that it is also more empirically plausible.

5.3 Minimal Polyadicity (Again)

Ugly facts will, no doubt, force us to a view more complicated than the one urged here. But positing an ideal Fregean/Tarskian mind, to account for human linguistic competence, seems like bad methodology. So I approach the issues from a different direction. Of course, nobody thinks that human heads house ideal minds. But many semantic theories invoke powerful machinery. This makes it tempting for authors to hedge about whether, or in what sense, they are offering empirical claims about human psychology; see, e.g., Davidson (1967b). My preference is to start with an overtly psychological hypothesis that may be too simplistic, and ask what else we need to posit. Instead of assuming that we can form concepts like ∃Z[ABOVE(X, Y) ! GIVE(Y, Z, W)], and asking how we deploy this capacity in comprehending speech, I assume that we can use linguistic expressions as instructions for how to build concepts like BROWN(_)^COW(_) and ∃[ABOVE(_, _)^COW(_)]; where this latter concept applies to things above a cow.

Recalling the discussion above, one can characterize M-junction in terms of the concept ∩[Φ(_), Ψ(_)]. Though instead of describing 'cow' as an expression of type <e, **t**>—taking the types <e> and <**t**> as basic for human language semantics—one might describe 'cow' as an expression/Begriffsplan of type <M> for monadic. And one might describe M-junction as an operation of type <M, <M, M>>. One can likewise describe Θ-junction as an operation of type <D, <M, M>>; where 'D' indicates a secondary, very limited semantic type. While Θ-junction permits construction of ∃[ABOVE(_, _)^COW(_)], the idea is not to generate concepts like ∃X[GIVE(Z, Y, X) & COW(X)] or ∃X[ABOVE(Y, Z) & GIVE(Z, Y, X)]. Even the monadic concept ∃X[COW(Y) & BROWN(X)] fails to be generable via M-junction and/or Θ-junction. And as discussed in chapter five, human languages may not permit instructions of type <e>. In which case, the basic instruction types may be <D> and <M>, with no non-basic types, since all complex expressions are of type <M>; cp. Partee (2006).

Distinguishing semantic instruction types from Fregean conceptual types, and not appealing to <**t**> as a basic instruction type—and certainly not invoking an existential quantifier of type <e**t**, **t**>, much less a *more* flexible quantifier that can combine with *any* predicative concept—highlights the question of what a sentence is. I take this to be a good thing, since as noted above, our best current conceptions of generative grammar make it unclear which if any generable expressions can plausibly be identified with sentences. This unclarity is unsurprising if the notion of sentence has its home in *subject*-predicate conceptions of *thought*, while the notion of subject has no stable place in generative *grammar*; cp. Chomsky (1965). It may be that the only natural sentences of type <**t**> belong to languages of thought with roots that are phylogenetically older than the spoken/signed languages that human children acquire. And perhaps expressions of these newer languages are used to build complex monadic concepts, including some special cases that are importantly equivalent to natural thoughts of type <**t**>, without ever generating expressions of type <**t**>. In which case, the very idea of a *truth* conditional semantics for a human language may be fundamentally misguided.

Such speculations aside, Tarski (1933) himself showed how closed sentences could be viewed as special cases of monadic predicates that have satisfaction conditions, and why this does not require appeal to truth values; see chapter two. Drawing on Tarski, I'll introduce a pair of operations/instructions that correspond to existential closure and negated existential closure of a monadic concept. But these will be of type <M, M>; and <M, M> is not <e**t**, **t**> in disguise. The idea, spelled out in chapter six, is that COW(_) can be used to build either of two monadic concepts—↑COW(_) or ↓COW(_)—that can be used to think about everything or nothing, depending on whether or not there is a cow: if COW(_) applies to something, then ↑COW(_) applies to each thing, and ↓COW(_) applies to nothing; if COW(_) applies to nothing, then ↑COW(_) applies to nothing, and ↓COW(_) applies to each thing. The arrows—suggesting upward and downward entailment, modern versions of inference patterns discussed in section two—can only attach to monadic concepts: ↑COW(_) is equivalent to though not identical to ∃X[COW(X)]; by contrast, ↑ABOVE(_, _) is neither generable nor coherent.

Propositional conjunction is easily reconstructed in these terms. This in turn invites a proposal about the kind of abstraction required for relative clauses and quantificational expressions like 'chased every cow that most dogs chased'. One might think it obvious that such expressions are not instructions for how to assemble monadic concepts, even if quantifiers raise and leave traces of displacement, as in the following (partial) syntactic analysis of 'Fido chased every cow': [[every cow]$_1$ [Fido [chased t$_1$]]]. But chapter six shows how to treat traces as instructions of type <M>, akin to the proper noun in 'the John we know', and how to characterize the meaning of 'every cow' in terms of Θ-junction.

I'll argue that if quantificational expressions are instructions for how to build concepts of the form ∃[Θ(_, _)^Φ(_)], then these expressions are no harder to accommodate than on views that appeal to saturation and concepts of type <e**t**, <e**t**, **t**>>. In fact, the constraints make it easier to isolate what is really required to accommodate quantificational constructions. Similar points apply, I claim, to plural constructions like 'chased three dogs'. Specialists will want to see the details, including how appeal to truth values can be avoided when it comes to introducing analogs of (nonfirstorderizable) concepts like MOST[Φ(_), Ψ(_)], which cannot be captured in terms of Tarski's (1933) apparatus. But these are matters for chapters five and six.

5.4 Following Instructions

Let me instead end this first chapter by returning to the relevant notion of instruction and an important sense in which it is intentional. My suggestion is not that 'gave a cow a dog' is an instruction to build a concept with which one can think about things that gave a cow a dog. This instruction might be executed by building the concept ∃Y∃Z[GAVE(X, Y, Z) ! COW(Z) ! DOG(Y)], which has a triadic constituent. My claim is that 'gave a cow a dog' is an instruction for how to build a concept like [[GIVE(_)^PAST(_)]^∃[RECIPIENT(_, _)^COW(_)]]^∃[PATIENT(_, _), DOG(_)], which has only an occasional dyadic concept that is "sealed in." This M-junction has constituents that include two Θ-junctions. For some purposes, the differences between the concepts may not matter. But formal distinctions can be crucial when the issues concern the *natural generative procedures* that underlie human linguistic competence and cognitive productivity.

Put another way, to satisfy semantic instructions of the sort envisioned here—in the sense in which an instruction can be satisfied by suitable *activity*, as opposed to an invented sentence being satisfied certain *entities*—monadic concepts must be assembled in certain ways from building blocks of certain sorts. Assembling monadic concepts in *other* ways does not satisfy the instructions. (Compare: I told you to make ice cubes from the tap water, not the expensive

bottled water.) So executing assembly instructions *may but need not* yield concepts that have satisfaction conditions in some other sense.

If meanings are human Begriffsplans (concept assembly instructions) that have mechanical execution conditions, executing these instructions need not lead to ideal concepts. Though in some cases, the assembled concepts may approximate Fregean models. When conditions are optimal, following a recipe may lead to an especially good instance of a product type, even if the recipe is not an instruction for how to create especially good instances of that type. In a kitchen stocked with excellent ingredients, a skilled baker might create a sublime pie by following the procedure specified in a certain book, while another baker in another kitchen makes a merely adequate pie by following that same procedure. Often, an instruction can be carried out in many ways, yielding results that reflect many factors. To take a nonculinary example, suppose we are each told to find a red box and a blue box, and then put the former on the latter. The stability of our stacks will depend on the size and rigidity of the available boxes. But the *instruction* can be satisfied, even if the result is unstable or otherwise less than ideal.

As we'll see, if meanings satisfy demanding compositionality constraints, then the corresponding concepts may "fit each other" better than they "fit the world." For many purposes, it may do no harm to suppose that ∃Y∃Z[GIVE(X, Y, Z) ! COW(Z) ! DOG(Y)] is "extensionally" equivalent to its natural analog. But this idealization may be limited in ways that matter for both meaning/psychology and truth/logic. My suspicion is that most natural concepts do not have extensions; cp. Travis (1985, 200x), Sainsbury (1996). But in any case, I deny that meanings are instructions for how to build concepts that exhibit classical semantic properties, even if executing such instructions sometimes leads to assembly of concepts that have classical semantic properties.[31] Procedure matters. In my view, the posited instructions/Begriffsplans have mechanical execution conditions that make no reference to the things we usually think and talk about. Indeed, at least with regard to the core recursive cases, the steps are rather simple-minded: *fetch*, **JOIN**, repeat.

From this perspective, theories of human linguistic meaning are theories of how human linguistic expressions are related to human concepts, whose relation to truth may turn out to be quite complicated and orthogonal to the central issues concerning natural semantic composition. This internalistic conception of meaning can be viewed as a variant on the old idea that meanings *are* concepts—mental representations with which thinkers can think about things—as opposed to things about which speakers can speak. But this runs counter to a newer tradition, according to which meanings (if such there be) are more like referents, and meaning is a lot like referring.

This sets the agenda for chapter two, which describes a possible language of thought, PL, that has a truth conditional semantics: sentences have satisfaction conditions that are recursively specified in terms of a given domain, and truth is characterized in terms of satisfaction; see Tarski (1933). In chapters five and six, I describe another possible language of thought, SMPL. This *second-order* but massively *monadic* predicate calculus permits *plural* predicates, as discussed by Boolos (1998). In some respects, SMPL is more restricted than PL; in other respects, it is richer and more Fregean. That is, SMPL is an extension of a fragment of PL. Eventually, my proposal will be that meanings are instructions for how to build (concepts that

31 In earlier work (e.g., Pietroski 2005b), I simply described meanings as instructions to build concepts, without worrying about the distinction between commands and blueprints: 'instruction' is polysemous in a way that lets us talk about an instruction *to* Φ or an instruction *for how to* Φ. I had mistakenly thought that analogies to executable programs would make the idea clear enough, and that talk of instructions could be accepted as relatively innocuous without further terminological fuss.

are) predicates of SMPL by fetching atomic predicates and combining them in the simple ways outlined in chapter one. As we'll see, this account of how meaning is related to psychology has implications for how meaning is (not) related to truth, while the Davidson-Montague conjecture has implications for how meaning is (not) related to psychology. But first, we need to know what it is for any systematically productive language—invented or natural—to have a truth conditional semantics.

Chapters five and six also return to details concerning lexicalization, concept introduction, and semantic composition in human languages. The idea is that concepts like GIVE(_), introduced so that they can be M/Θ-joined with others, tend to fit the world *less* well than prior concepts like GIVE(_, _, _). Condensing a long line of thought, and leaving out the qualifications: GIVE(_) does not have mind-independent satisfiers; though unlike UNICORN(_), GIVE(_) is lexically related to a concept that does apply to real things; but unlike GIVE(_), GIVE(_, _, _) cannot be systematically combined with other concepts, at least not via the operations invoked by phrasal syntax. Condensing further, neither GIVE(_) nor GIVE(_, _, _) has classical semantic properties: the introduced monadic concept is too far removed from the world; the prior relational concept does not have the right form. As Frege and Tarski knew, *inventing* a language with a truth conditional semantics is hard enough. It takes work to sync form and content in the requisite ways.