THE PHILOSOPHY OF ANIMAL MINDS

EDITED BY

ROBERT W. LURZ

CAMBRIDGE
UNIVERSITY PRESS

CHAPTER 6

A language of baboon thought?

Elisabeth Camp

1. INTRODUCTION

Does thought precede language, or the other way around? How does having a language affect our thoughts? Who has a language, and who can think? These questions have traditionally been addressed by philosophers, especially by rationalists concerned to identify the essential difference between humans and other animals. More recently, theorists in cognitive science, evolutionary biology, and developmental psychology have been asking these questions in more empirically grounded ways. At its best, this confluence of philosophy and science promises to blend the respective strengths of each discipline, bringing abstract theory to bear on reality in a principled and focused way. At its worst, it risks degenerating into a war of words, with each side employing key expressions in its own idiosyncratic way – or worse, contaminating empirical research with a priori dogmas inherited from outmoded philosophical worldviews.

In *Baboon Metaphysics* (2007), Dorothy Cheney and Robert Seyfarth offer an analysis of baboon cognition that promises to exemplify the very best interaction of philosophical theory and empirical research. They argue that baboons have a language of thought: a language-like representational medium, which supports the sophisticated cognitive abilities required to negotiate their complex social environment. This claim is intended to be surprising in its own right, and also to shed light on the evolution of spoken language. Because our own ancestors likely lived in a similarly complex social environment, Cheney and Seyfarth propose that the earliest humans also developed language first as a cognitive medium, and that spoken language evolved as a means to express those thoughts.

There are two potential difficulties here. First, "Language of Thought" (LOT) is a term of art, with much associated theoretical baggage and often comparatively little careful exposition. Thus, evaluating the claim requires getting clearer about just what LOT implies in this context. Second, if

Cheney and Seyfarth are right, then we seem to be left with a rather surprising model, on which baboons think in a way very like we do but fail to talk. According to Cheney and Seyfarth, this is because they lack a theory of mind. But this still leaves a dramatic mismatch between the expressive potential of baboons' postulated representational system and the complexity of the behavior they manifest in other areas. Thus, it's worth examining whether an alternative form of thought might better explain the distinctive contours of their conceptual abilities and limitations.

2. FROM CONTENT TO FORM

The claim that thought is in language or is language-like is a claim about the form of thought. Information itself – how things are in the world – doesn't have any form: the same information, such as the shortest route between two locations, can be represented in multiple ways; that's why the question of form can even arise. However, while we have an intuitive understanding of the differences in form among ordinary external representations, it's harder to know what it means for thought to have one or another form: surely we don't mean that minds are ethereal sheets on which a homunculus writes or draws pictures, as we inscribe marks on paper. Rather, claims about thoughts' forms are usually intended functionally: as concerning a level of description at which distinct token thoughts share certain structural features despite lower-level differences. The features that constitute a given thought-type *X* are determined by the functional role *X* plays within the overall system: by what other thoughts, perceptual inputs, and behavioral outputs *X* is produced by and produces when the system is operating correctly. The formal features of *X* are then determined by abstracting away from, first, *X*'s attitude – e.g., belief or desire – and second, its informational content: by isolating the transitions and co-instantiation relations that obtain regardless of what particular state in the world *X* represents.[1]

Predicate logic provides the most familiar case in which we distinguish formal or syntactic features from contentful or semantic ones; and it serves as the model for thought on the "Language of Thought Hypothesis" (LOTH). However, we need to distinguish at least two operative interpretations of LOTH (for discussion, see Camp [2007]). On the first

[1] "Content" is a loaded term, and many theorists will object that contents do already have formal properties, given either by the structure of the objects and properties in the world (*à la* Russell) or by the structure of the mode of presentation (*à la* Frege). Because we need a term for what is being thought about that doesn't already fix form, I use the term in a more coarse-grained way (*à la* Stalnaker).

interpretation, LOTH starts with the observation that representational abilities appear to be complex, in the sense of employing multiple simpler, interacting abilities; it then argues that those complex abilities must be underwritten by analogously complex vehicles (ultimately, neural configurations), such that the constituents and structure of those vehicles mirror the constituents and structure of what is represented by the corresponding abilities. Thus on this interpretation, LOTH is primarily a way of spelling out computationalism, and it remains neutral about the specific formal features of the vehicles' parts and structure. On the second interpretation of LOTH, which is the version that is relevant for Cheney and Seyfarth's discussion, the claim is that mental representations, however these are implemented, have a specifically language-like structure.

These interpretations are not always clearly distinguished, for at least two reasons. First, most philosophers have been exclusively concerned with human thought, and it is plausible that the structure of human thought often shares a common structure with the sentences we use to express them. Second, the claim that thought is language-like is typically opposed to the claim that thought is imagistic (e.g., Pylyshyn [1973]), and most philosophers assume that images lack formal structure in the relevant sense. The governing assumption is that there are just two kinds of representation: imagistic or *iconic*, which is analog and employs direct resemblance between signifier and signified; and sentential or *propositional*, which is digital and employs a merely conventional or causal referential connection between signifier and signified. However, as we will see in sections 4 and 5, there are other representational formats – most notably, maps and diagrams – which mix iconic and non-iconic, analog and digital elements in different ways, and which are governed by their own distinctive formal principles (Camp [2007]). In many ways, these mixed systems, rather than imagistic ones, represent the most theoretically interesting alternatives to language.

Because form is an abstract property of thought, it cannot be accessed directly. However, we can make provisional inferences about a thought's form by examining the overall pattern of contents that can be represented within the cognitive system and how the system manipulates those contents (where these features are themselves inferred indirectly from behavior). Most representational formats exhibit distinctive limitations in what information they can represent. For instance, because imagistic systems represent by replicating the physical appearances of particular objects, they cannot directly represent disjunctive information.[2] Further, different sorts

[2] Although a larger cognitive system might be able to deduce from a picture that either P or Q must obtain.

of information are easier to access, compile, and update in different formats. For instance, in cartographic, but not linguistic, systems, altering the location of any one represented individual automatically alters the represented relations between it and every other represented individual. Thus, we can glimpse the distinctive contours of different forms of thought by attending to limitations in the contents that a system is capable of representing, and to how it clusters and distinguishes distinct contents in processing.

Finally, it's worth noting that insofar as we construe form functionally, a representational system need not be implemented in a way that looks like the formats we ordinarily employ. Instead, a physical substrate implements a system in virtue of there being some consistent, causally operative mapping from formal to physical properties. For instance, there must be some properties of the implementation of a functional map that correspond to "being east of" or "being to the left of," although they need not themselves be being to the east or left. Moreover, those properties must be appropriately related to each other, and to further properties like "being north of" and "being in front of," in such a way that one could construct a normal physical map just by substituting the functionally appropriate counterparts.

3. BABOON CONTENTS

Although Cheney and Seyfarth don't distinguish form and content as explicitly as I have, the argument of their book implicitly moves from a claim about the contents that baboons represent to a conclusion about how they represent them. By far the most notable feature of baboon cognition is their ability to track complex social dynamics. Each baboon troop consists of around eighty members, in which females form a fairly stable dominance structure while males migrate and change rank frequently. Females' social ranks are determined by two factors: first, by which (matrilineal) family they belong to, and second, by their birth order within that family.

Social rank has a pervasive effect on everyday interactions, determining who is allowed to eat what, sit where, groom or mate with whom, or handle whose baby. And baboons clearly behave in ways that indicate their awareness of others' and their own relative ranks: for instance, if a dominant baboon approaches two lower-ranking ones, the most subordinate typically retreats, as if she recognizes that she is subordinate to both of the others (p. 92).[3] Female baboons also regularly act to preserve and improve their rank, including in indirect ways. For instance, they join alliances with their

[3] All page numbers without author or date are from Cheney and Seyfarth (2007).

close relatives or higher-ranking females against opponents with whom they have no independent conflict, and they employ the close kin of baboons with whom they've recently fought as proxies for direct interaction – indeed, "kin-mediated reconciliation" is twice as common as direct reconciliation (p. 83).

Much of baboons' social interaction takes the form of vocalizations, especially threat grunts, fear barks, and reconciliation grunts. Although these vocalizations themselves lack any real internal syntactic structure, baboons who overhear unseen callers are able to recognize the caller and type of call being made. This information can in turn affect their own behavior: for instance, if two unrelated baboons are sitting near one another and overhear a call that represents a conflict between their relatives, the baboon from the higher-ranking family is then more likely to go assert her own dominance over the other than if they heard a call involving baboons they weren't related to (p. 97). They also demonstrate surprise when they hear (artificially constructed) sequences of calls that violate their expectations about social order. In particular, baboons respond with more surprise to sequences that represent a subordinate threatening a dominant from a different family than they do to calls that represent an analogous intra-familial conflict, even when the difference in overall rank order is the same (p. 108). This difference in evinced surprise provides compelling evidence for the hypothesis that baboons represent social rank in a two-tiered, hierarchical manner – a hypothesis we'll examine more closely in the next section. More direct evidence for this hypothesis comes from the fact that when a conflict between two baboons does lead to a change in dominance, the effects of that change depend on whether the baboons are related: intra-familial changes in rank affect only those two baboons, while a change in rank involving baboons from different families may cause many members of the losing baboon's family to lose rank, and even fall to the bottom of the overall order (p. 70).

4. LANGUAGE-LIKE FEATURES OF BABOON THOUGHT

Given this background, our question now becomes: what do the contents of baboons' thoughts reveal about their form? Why should we believe that they think in a language-like way? In keeping with many discussions of LOT, Cheney and Seyfarth merely claim that baboon cognition is importantly *like* natural language (p. 250). However, in contrast to many such discussions, which remain extremely vague, Cheney and Seyfarth identify six specific respects of similarity, which they summarize in the following passage:

> The vocalizations of monkeys clearly lack any properties that we would be tempted to call syntactic. Nevertheless, their social knowledge, their assessment of call meaning, and their parsing of call sequences display a number of syntactic properties. First, knowledge is *representational*, and highly specific. Second, social knowledge is based on properties that have *discrete values*... Third, animals combine these discrete-valued traits to create a representation of social relations that is *hierarchically structured*. Baboons, for example, create a nested hierarchy in which others are placed in a linear rank order and simultaneously grouped according to matrilineal kinship in a manner that preserves ranks both within and across families. Fourth, social knowledge is *rule-governed and open-ended*... Fifth, knowledge is *propositional*... That is, they represent in their minds (albeit in a limited way) the individuated concepts of "Sylvia," "Hannah," "threat-grunt," and "scream," and they combine these concepts to create a mental representation of one individual's intentions toward another... Sixth, knowledge is *independent of sensory modality*... These properties of non-human primates' social knowledge, while by no means fully human, bear important resemblances to the meanings we express in language, which are built up by combining discrete-valued entities in a structured, hierarchical, rule-governed, and open-ended manner. (pp. 268–269, italics in original)

I'll condense their six properties into three broad categories:

- Representational/propositional/independent of sensory modality
- Discrete-valued/rule-governed/open-ended
- Hierarchically structured

Taken together, these features appear to make a strong case for the claim that baboons think in something very like language. They are the features most commonly cited in discussions of LOT and of what constitutes a language; and I agree with Cheney and Seyfarth that baboon cognition is like language in all these respects. However, in this section I argue that none of these features are sufficient to make baboon cognition *distinctively* language-like, because they can also be exhibited by other representational systems, such as maps and diagrams. If this were just a disagreement about what constitutes a language or being language-like, it wouldn't be very substantive. In section 5, though, I argue that language differs from these other systems in a crucial further respect – its combinatorial principle lacks a substantive, dedicated semantic significance – and that this difference explains its comparatively greater expressive capacity. In section 6, I argue that this feature is likely absent from baboon cognition, even though at several points Cheney and Seyfarth talk as if it were present. Thus, the particular contours of baboon cognition are better explained on the hypothesis that baboons employ a representational format that differs importantly from language.

4.1. Representational thought: representational, propositional, and independent of sensory modality

The first property that Cheney and Seyfarth mention, of being *representational*, is fairly obviously a necessary condition on being a linguistic system; but just as obviously, it's also not sufficient: maps, pictures, and diagrams are all also representational. Likewise, die-hard behaviorists aside, nearly all parties will agree that baboon cognition is representational, simply in virtue of being a kind of thought.[4] However, this leaves open the very question we're asking: what form do those representations take?

The second property, *propositionality*, can be construed in two ways. Sometimes "propositional" is used in a way that's basically equivalent to "representational": it just means being truth-assessable, or representing a state of affairs, in which case it is again necessary but not sufficient for language.[5] In its other use, "propositional" is tied specifically to sentential structure: to say that thought is propositional in this sense implies that it has subject–predicate structure.[6] On this interpretation, though, the claim that baboon cognition is propositional is question-begging: it simply attributes linguistic structure to baboon cognition without further evidence.

The third property is *independence from sensory modality*; for example, baboons incorporate information derived from overheard vocalizations and from directly observed interactions into a single representation of social rank.[7] Again, though, this feature can be implemented in a variety of formats. For instance, the thought of *A* threatening *B* might be represented by an image of the two fighting, even if it is caused by overhearing a sequence of calls. And while images can only explicitly encode visual

[4] Some theorists (e.g., Gibson [1966]) do deny that thought is representational.

[5] This usage seems to be employed in the following passage from Cheney and Seyfarth:

> Applied to humans, the language of thought hypothesis assumes that people – even preverbal infants – have propositional attitudes: mental states with representational content. To have a certain propositional attitude is to be in a specified relation to an internal representation: to know, believe, or fear that something is the case. (p. 250)

[6] Cheney and Seyfarth appear to be employing this sense when they write that "the baboon's alarm wahoo seems best described as a proposition: a single utterance or thought that simultaneously incorporates a subject and a predicate" (p. 256). Likewise, immediately prior to the passage cited in footnote 5, Cheney and Seyfarth write: "Like natural languages, [the Language of Thought] includes some mental representations that correspond to objects in the world and others that specify the relation between these objects" (p. 250).

[7] I have argued (2009) that a related feature, of *stimulus-independence* – the ability to compile information from multiple sources and occasions, and to apply it in flexible ways in variable contexts – is a necessary condition on genuine conceptual thought.

properties, maps frequently represent objects and properties that lack any sensory appearance, such as an earthquake's epicenter.

4.2. Compositional thought: discrete-valued, rule-governed, and open-ended

The next class of properties that Cheney and Seyfarth discuss are more closely associated with LOTH. First, baboon cognition is *discrete-valued*: for example, baboons represent each baboon in their troop as a distinct entity. They also classify vocalizations into discrete categories with sharply different meanings, such as *Hannah's threat grunt*, *Shashe's threat grunt*, or *Shashe's fear bark*, even though in terms of their purely physical acoustic properties, those vocalizations form a continuum (p. 231). Second, baboon cognition is *rule-governed*: for instance, if a baboon observes a dominant lose a fight to a subordinate, this has systematic effects on the represented ranks of those two baboons and possibly of other troop members as well. Similarly, if a baboon hears an audio sequence that is ambiguous between two interpretations – say, it displays two baboons making threat-grunts and a third making a fear-bark, which could be a response to either threat – she will respond as if she "parses" the sequence according to a "rule" given by the established dominance order (p. 266). Third, baboon cognition appears to be *open-ended*: baboons are not just capable of representing specific dominance relations among specific baboons, but can represent nearly any configuration of relations among all the baboons they know.[8] Indeed, Cheney and Seyfarth's experimental methodology exploits just this feature: they present baboons with artificially constructed audio sequences that are unlikely to occur naturally. The baboons don't merely find such sequences anomalous, as they might if they heard an entirely new sound. Rather, they appear to find the sequences surprising precisely because of the information conveyed: for instance, they look at relatives of the individuals whose calls are being played, as if to check their reactions.[9]

Taken together, these three properties imply that baboon cognition is *compositional*: it combines discrete representational parts into larger wholes, such that (nearly) any combination of parts licensed by the system's

[8] Not absolutely all combinations need be possible: in particular, baboons are likely unable to represent alternative kinship relations, such as Sashe belonging to Hannah's family. For a discussion of what would underwrite such cognitive flexibility, see Camp (2009).

[9] This contrasts with a structurally analogous experiment involving honeybees who were confronted by bee dances representing the surprising information that nectar was located in the middle of a lake (Gould and Gould [1988]; Tautz *et al.* [2004]). The bees' "response" – inaction – doesn't give us reason to distinguish surprise or disbelief from incomprehension (Camp [2009]).

combinatorial rules is well-formed and significant.[10] The combinatorial rules of such a system are (relatively) *general* or *systematic*, in the sense that they apply to (nearly) all parts of the relevant formal types, regardless of their particular content. As a result, such systems also display (relative) *productivity*: the system can represent any of a large range of contents by combining a fairly small pool of basic elements in different ways.

Compositionality is traditionally taken to be an essential, even definitive, feature of language. There is significant debate about just how fully compositional natural languages really are, both because the intuitive meaning of a whole sentence often seems to depend on factors that are not explicitly coded by its constituent terms or mode of composition (e.g., Travis [1994]), and because the syntactic categories that govern the well-formedness of whole sentences are often quite narrow and overlapping (e.g., Johnson [2004]). But it is nearly undeniable that language is compositional to some large degree. Further, classic arguments for LOT rely exclusively on these features: because our own and other animals' cognitive abilities appear to be systematic and productive, the argument goes, the vehicle that underwrites those abilities must be language-like in its form.[11] The implicit assumption is that only language is compositional, and hence capable of producing systematic, productive representational outputs.

However, language isn't the only type of compositional system (Camp [2007]). Most obviously, diagrammatic systems, such as Venn diagrams, are compositional. But further, many maps are also constructed from discrete elements – a city map might employ black lines for streets, green squares for parks, and red dots for ATM locations, for instance – which may or may not bear any resemblance to what they represent, and which can be recombined in different ways according to a uniform principle to represent a wide range of different states of affairs.[12] Indeed, maps can be even more systematic than language, insofar as they don't impose any syntactic restrictions on which types of elements can be combined, as language does.

[10] Cheney and Seyfarth explicitly claim that baboons' vocalizations – or more properly, their interpretation of other baboons' vocalizations – are compositional (p. 266).

[11] Editor's note: See Tetzlaff and Rey (chapter 4) and Carruthers (chapter 5) for such an argument applied to honeybee and (other) invertebrate cognition.

[12] I use the term "map" narrowly, to refer to spatial depictions of isomorphic spatial situations. I call spatial depictions of, for example, abstract information "diagrams." One might object that maps are not discrete, because the elements are typically put together in, and represent locations in, a continuous space, and because they often employ icons that vary continuously, such as the shape of a blue line depicting the shape of a river. However, maps' signs need not be iconic and often carry discrete semantic import, and maps can employ discrete representational spaces, such as a seating chart.

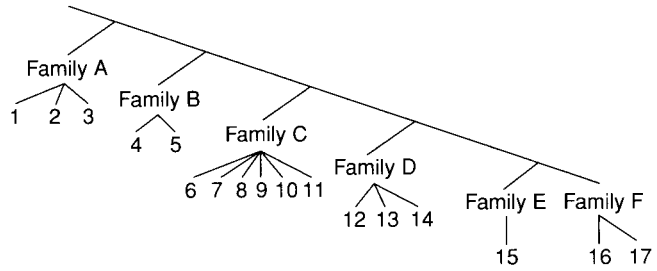

Figure 6.1 Hierarchical dominance relations in baboons (after Cheney and Seyfarth [2007], p. 107).

I discuss the implications of the combinatorial principles employed by these various systems in section 5.

4.3. *Hierarchically structured thought*

The final property that Cheney and Seyfarth identify in baboon cognition, of being *hierarchically structured,* is the most surprising. In section 3, I sketched the evidence that female baboons' social ranking is hierarchically structured: whole families are ranked relative to each other, and individuals are in turn ranked within each family, producing an overall absolute order. Further, baboons themselves seem to be aware of this hierarchical structure, insofar as they react differently to inter-familial and intra-familial conflicts involving baboons with the same absolute rank difference. From this evidence about the contents of baboons' representations of social order, Cheney and Seyfarth conclude that the form of their thought is also hierarchically structured: as they say, "the picture that emerges in a baboon's mind can be diagrammed as a branching, hierarchical tree" (p. 277), of the sort depicted in Figure 6.1.

It should be relatively uncontroversial that we *can* represent baboons' awareness of social rank as in Figure 6.1. The controversy comes, first, in moving to the claim that we need to represent their awareness in this way, and second, in drawing a conclusion from this about the form of baboon cognition across the board. In particular, if all we had to go on was static evidence about the social relations that baboons represent at a given moment, we wouldn't have much reason to prefer a two-layered structure over a linear ordering: it might just be coincidence that individuals from the same family all happen to cluster together, or it might be a by-product of some further factor, like size or aggressiveness. However, the postulation

of a two-layered structure gains significant support from the dynamics of baboon cognition, and specifically from the fact that baboons treat intra- and inter-familial rank reversals so differently. Thus, at least for current purposes, I want to grant that baboons do have some sort of two-tiered representation of social rank.

Next, we need to get clearer on the sense in which such a two-layered representation is hierarchical. Cheney and Seyfarth don't specify exactly what they mean. Baboons' thought about social structure shouldn't count as hierarchical simply in virtue of representing one individual or family as ranked higher than another – after all, Cheney and Seyfarth don't claim that the purely linear dominance structure among males requires a hierarchical representation. Rather, we are seeking to identify a formal, structural property of baboons' thought. More specifically, to support the hypothesis that baboons think in a language-like format, the hierarchical structure in question should be as closely analogous as possible to the hierarchy in language. I believe what Cheney and Seyfarth have in mind is this: baboons' representation of social relations is hierarchical in the sense that the same relation – *dominant to* – which is represented as holding between individuals is also represented as holding between units that are comprised of those individuals. This is a lot like the sort of hierarchical relationship we find among phrases in language. For example, in a sentence like "The girl who kissed the boy with the red shirt is wearing a blue dress," the noun phrase "the red shirt" is embedded within the larger noun phrase "the boy with the red shirt," which is itself embedded within the noun phrase "the girl who kissed the boy with the red shirt," so that the same operation which creates the second noun phrase from the first is applied again to form the entire noun phrase.[13]

Suppose we grant, then, that both baboon cognition and language exhibit the following sort of hierarchical structure: a relation which applies to individual elements can also be applied to relate units that are constructed from those individual elements. How language-like does this show baboon cognition to be?

One important difference between baboon cognition and language is that actual languages can generate structures of potentially infinite depth; but it seems quite unlikely that baboons' representations are capable of more than two layers of structure. For instance, although baboons rank all the families within their troop, it seems unlikely they are capable of

[13] More generally, modern grammars aim to employ a single construction rule – e.g., functional application, conjunction, or Merge – which operates on expressions of various types to produce larger phrases from smaller ones by a uniform operation.

ranking whole troops in relation to one another: there simply aren't the kinds of interactions among whole troops that would make this sort of ranking feasible. However, I don't think this difference matters much for current purposes.[14] A language-like system capable of just one layer of embedding, supplied with concepts for the contents that Cheney and Seyfarth argue baboons can think about – such as social relations, causal processes, patterns of food availability, and threats from predators – would still be a quite powerful representational medium.

Even putting these differences to the side, though, the mere presence of hierarchical structure is not enough to make baboon cognition distinctively language-like, because once again, other representational formats can also exhibit this feature. To take an example of particularly direct relevance, Linnean taxonomic trees employ branching lines to depict phenotypic differentiation. Here we have the same type of functional structure as in baboons' representations of social relations, simply with phenotypical similarity replacing dominance. Individual animals are classed together as a functional unit (a species) in virtue of sharing certain micro-features, and those functional units are in turn classed together with other such units to form a larger group (a genus), on the basis of sharing certain macro-features.[15]

Other non-linguistic forms of representation can also exhibit hierarchical structure. Thus, a Venn diagram might partition a collection of Venn diagrams, each of which represents some relation among the extensions of specific properties, according to whether those diagrams themselves possess certain features. It is even arguable that maps can exhibit hierarchical structure. For instance, a map collector might draw a map of the locations of the maps in his library. One likely difference between both of these cases and baboons' representations of dominance is that in the latter, the hierarchical relation is transitive: individuals from different families are represented as ranked in dominance in virtue of the represented ranking between their families. However, an especially eccentric map collector might arrange a library of maps of the United States according to the

[14] Cheney and Seyfarth (p. 252) explicitly include the closely related property of recursion among the features specific to modern natural languages that they're not claiming baboon cognition to possess. Another potentially significant difference is that in language, the application of the relevant operation (say, functional application) is what binds individual constituents together into a functional unit, whereas in baboon social cognition, a distinct operation – kinship – constitutes those units. Even so, I think the analogy between the two cases is tight enough, in the respect specified in the text, to warrant the claim that baboon cognition is like language in being hierarchical.

[15] Taxonomic trees are now often taken to represent phylogenetic history as well, though the phenotypic and phylogenetic interpretations do not always produce the same overall structure.

location in the country that the maps depict, so that maps of New Jersey are located on the east wall, maps of Texas on the south, and so on. In that case, the first-order functional unit would be an individual map, in which the spatial configuration of individual signs purports to replicate the spatial configuration of cities and other depicted entities. But in turn, those maps would themselves serve as signs representing the regions they depict, so that the overall arrangement of maps within the library replicates the spatial configuration of those regions.

Even putting these more recherché cases aside, the case of taxonomic trees shows clearly that hierarchical structure is necessary but not sufficient for language. Hence, it is invalid to infer from the presence of hierarchical structure to the conclusion that baboon cognition must be distinctively language-like. Further, because all of the representational formats we've been discussing can also possess all of the other features that Cheney and Seyfarth identify, the inference from the presence of all of these features to the conclusion that baboons have a language of thought is likewise invalid.

5. GENERALITY, FLEXIBILITY, AND THE CONTENTS OF THOUGHT

So far, we've established that appealing to structural features like compositionality or hierarchical structure doesn't suffice to establish the form of baboon cognition. To make further progress, we need to get clearer on how these various representational formats differ from language, and specifically on what implications this has for the contents that cognitive systems employing them are capable of representing.

I think the most important difference is this. In all of these formats, including language, we have a combinatorial principle – a rule for combining the system's basic parts into larger wholes – which is applied repeatedly, and potentially at distinct levels. In language, this principle is highly abstract, and has a very thin semantic significance. Predication, for instance, signifies instantiation or property-possession, so that the significance of combining "Socrates" with "is wise" is that the referent of the former instantiates the property expressed by the latter; and this relation is sufficiently abstract and general that it can relate nearly any property and object.[16] Further, in language the referential relation mapping basic expressions to objects and properties in the world is conventional or causal. Taken together, both the referential relation and the combinatorial principle are

[16] Other combinatorial principles for language, such as functional application, conjunction, and Merge, are even more general and abstract.

abstract enough that they don't impose substantive in-principle limitations on what can be assigned as referents to those basic expressions.[17]

Like languages, maps and diagrams can also employ arbitrary signs as their basic representational elements, and so they have a much wider expressive range than imagistic systems do. However, in maps and diagrams the basic combinatorial principle often has a quite robust semantic significance. And of course, a representational system can only represent the referents of its signs as being related by the relation that is represented by its combinatorial principle. When that combinatorial principle does have a robust significance, however, then this in turn constrains what objects and properties can be assigned as referents for the system's basic signs. So, for example, in maps (as opposed to diagrams) the combinatorial principle is one of (approximate) isomorphism, up to a scaling factor, between the signs' spatial configuration and the configuration of the objects and properties that they refer to. Because they are constructed this way, maps are only capable of depicting spatial relations among particular objects and properties. A map cannot, for instance, represent just that George is happy, or that someone somewhere is wearing a red shirt and carrying a gun, without placing George or some particular individual in some particular location; nor can a map assign a non-spatial relation, like dominance, as the referent of a basic expression (Camp [2007]). Likewise, the combinatorial principle that governs Linnean taxonomic trees is that entities on lower branches share a set of common features, which is represented by the node from which the branches radiate. As a result, Linnean trees cannot represent spatial relations among objects – nor can they even explicitly represent the shared and differentiating properties that determine each level of the tree.

The fact that a representational system's governing combinatorial principle can so heavily constrain the range of contents it is capable of representing provides, I think, a rather surprising justification for a very traditional intuition behind the Language of Thought Hypothesis. Theorists typically justify LOTH by appealing to the generality or systematicity of thought, by which they mean the general recombinability of one's basic concepts. As I said in section 4.2, an argument that moves straight from this sort of generality to the distinctively language-like form of thought is invalid, because maps and diagrams can be at least as recombinable as language. As we might put it, all of these formats are equally capable of representing a complete range of contents within a certain domain.

[17] However, because language is digital or discrete it cannot directly represent continuous values (Camp [2007]).

However, theorists have also ascribed another, arguably more profound sense of generality to thought, which is more distinctively connected to language. This is generality in the sense of the size of the domain the system can represent. Many theorists have held that it is an essential feature of thought or reason, as opposed to mere automatic response, that it be flexibly adaptable to a wide range of situations. Thus, Descartes (1637/1988, p. 140) concludes that men's souls must have an entirely different nature from those of animals, because even the stupidest man can invent and combine signs to make himself understood, whereas no animal or robot can do the same. Turing (1950) took a closely related criterion – the ability to converse in a way indistinguishable from a man, on any topic that might be presented – as an operational definition of intelligence. For both thinkers, it is not the ability to use language per se that is criterial of intelligence; rather, it is the ability to flexibly adapt to any situation. However, both also assume that only language provides an expressive medium that is flexible enough to manifest this cognitive adaptability.

We can now see why language should be so tightly connected to flexible, adaptable thought: it exploits referential and combinatorial principles that are abstract enough not to substantially constrain the range of contents the system can represent. By contrast, insofar as other systems employ combinatorial principles with more substantive semantic significance, the domain of contents that they can represent is more restricted. Neither Descartes nor Turing was interested in claims about thought's form; rather, they wanted to identify "sure signs" for discerning intelligence from a third-person perspective. But the same considerations apply if we turn to the medium of thought: maximal expressive generality is most easily achieved in a language-like format.

Unlike Descartes, I don't think we should restrict thought or reason to humans. Many other animals clearly do represent states of affairs in the world, and manipulate those representations in systematic ways to produce action. Further, some other animals clearly achieve a significant degree of both compositional and expressive generality (Camp [2009]). Rather than seeing thought or reason as an all-or-nothing affair, we should treat it as a matter of degree.[18] But if we do, then we should also acknowledge that thought can be implemented in many formats, with different degrees of expressive generality. In particular, while maps and taxonomic trees are expressively quite limited, Venn and Pierce diagrams are much more

[18] Editor's note: See Carruthers (chapter 5) for a comment on this point.

flexible – indeed, they can achieve expressive equivalence to first-order monadic predicate logic (Shin [1994]).[19]

6. BABOON THOUGHT

What are the implications of this theoretical discussion for baboon cognition? For one thing, we can clearly rule out the possibility that baboons represent dominance relations using maps, since dominance is a non-spatial relation, and maps can only explicitly represent spatial relations. More relevantly, it seems plausible that baboons might well represent female dominance relations in a form that can be drawn as the sort of branching tree depicted in Figure 6.1, as Cheney and Seyfarth suggest. However, this is fully compatible with the tree's branches and nodes having a dedicated semantic significance, as in a taxonomic tree. If that were the case, then baboons would have a form of thought with much less expressive generality than language, despite being discrete, rule-governed, open-ended, and hierarchical. And this would in turn explain the fact that they don't seem to exploit the full expressive generality that even a language capable of a single layer of embedding would provide.

At some points, it appears that Cheney and Seyfarth are only committed to the weaker claim that baboons' social knowledge alone forms a hierarchical, compositional, representational system: for instance, they write that "[Baboons'] social knowledge constitutes a discrete, combinatorial system of representations – a language of thought – that shares several features with human language" (p. 251). If baboons' hierarchical compositional system were indeed restricted to representations of dominance, then the claim that they think in a language-like format with a non-restrictive combinatorial principle that operates on a highly restricted conceptual domain would be empirically indistinguishable from the claim that they employ a system with a semantically restricted combinatorial principle. But in that case, the additional expressive generality of the language-like system would be doing

[19] Such systems are still limited by the fact that they exploit the spatial structure of the representational vehicle: for some contents, the appropriately intersecting figures cannot be drawn in a single plane (Lemon and Pratt [1998]). On the other hand, many theorists have held that language also falls short of full expressive generality. (We saw in section 4.2 that it falls short of full compositional generality.) That is, various philosophers and linguists have wanted to impose limitations on the range of terms that can be significantly combined, even when those combinations are syntactically well formed. This reflects, I think, an intuition about the semantic significance of predication: that, for example, you can't meaningfully predicate "is a prime number" of Julius Caesar because he's not the kind of object that could possibly possess that property. In Camp (2004), I argue against such restrictions on recombinability.

no explanatory work, and so it would be more parsimonious to explain baboons' cognitive abilities by hypothesizing that they employ a tree-like structure whose combinatorial principle has a dedicated significance, of dominance.

At other points, though, Cheney and Seyfarth commit themselves to the stronger claim that all of baboon cognition is robustly language-like, as in the following passage:

> The lack of syntax in non-human primate vocalizations cannot be traced to an inability to recognize argument structure – to understand that an event can be described as a sequence in which an agent performs some action on an object. Baboons clearly understand the difference between *Sylvia threatens Hannah* and *Hannah threatens Sylvia*. Nor does the lack of syntax arise because of an inability to mentally represent descriptive modifiers (a *big* leopard as opposed to a *small* one) and prepositions that specify locations (a leopard *in* a tree as opposed to one *under* a tree) . . . In their natural behavior, therefore, non-human primates and other animals certainly act as if they are capable of *thinking*, as it were, in sentences. But the ability to think in sentences does not lead them to *speak* in sentences – in our view, because their lack of a theory of mind causes them not to understand what others might need to know. (p. 264)

Here, Cheney and Seyfarth go beyond simply attributing some sort of hierarchical, compositional format for social knowledge: they speak as if all of baboon cognition occurs in a format whose structure fairly closely resembles that of natural language. And indeed, the considerations in section 5 might seem to provide an important motivation for this stronger view. Baboons don't just think about dominance: they also represent food, mates, predators, and other things, and their actions are often informed by information about multiple topics; as Cheney and Seyfarth say,

> A baboon's ability to acquire the most nutritious food depends, simultaneously, on both her ecological knowledge of plants and her skill in competition with others. Her ability to detect and evade predators depends, simultaneously, on both her knowledge of predator behavior and her ability to live cooperatively in a large group, where she benefits from predator detection and defense. (p. 122)

The fact that baboons can bring thoughts about various topics together to produce unified action seems to suggest that they have a general medium of thought. But since only language provides enough expressive generality to cover all of these domains, it appears that baboons must think in language after all.

Against this stands the fact that baboons don't manifest an ability to think hierarchically structured thoughts in any domain *except* dominance. For instance, they don't engage in any sort of tool use, as chimpanzees do, and as we might have expected if they could entertain thoughts with complex clausal structure, such as "Hitting the nut *with* a rock *causes* tasty nutmeats to come out." Cheney and Seyfarth explain a different lack of ability – the fact that baboons don't make syntactically complex, language-like vocalizations – by appealing to the hypothesis that they lack a theory of mind. The relationship between theory of mind and communication is complex enough to merit its own discussion.[20] But given how sophisticated baboons' social awareness already is, their very lack of theory of mind constitutes another domain where they don't manifest an ability to think hierarchically structured thoughts which would be quite useful for them.

Of course, there are many explanations for the absence of these abilities which are fully compatible with baboons' thinking in a language-like way: language is – at most – only a necessary condition on abilities like tool use and theory of mind. However, there is also no positive evidence in favor of their having a general hierarchically structured cognitive system. There is just evidence for hierarchically structured thoughts about dominance, on the one hand, and a general ability to represent various topics, on the other. But the latter ability is demonstrated by many, perhaps most, other animals; and Cheney and Seyfarth don't intend to claim, as for example Fodor and Pylyshyn (1988) do, that all thought is language-like. Rather, they intend to be advancing a bold claim specifically about baboon thought, with specific relevance for the evolution of spoken language in humans – where the most distinctive aspect of baboon thought is just the hierarchical structure of their representation of dominance.

The evidence presented by Cheney and Seyfarth is thus compatible with the hypothesis that baboons have multiple cognitive modules, perhaps encoded in different formats, with only the dominance module being hierarchically structured, as in a taxonomic tree. Proponents of modularity have offered various responses to the challenge to explain how thoughts from different modules might interact to produce action. It's worth noting, though, that if there were a dominance module, all the information about intra- and inter-matrilineal structure could be isolated within it. The only information that would need to be fed in and out would concern relations

[20] In humans, theory-of-mind deficits produce systematic impairments in linguistic production and comprehension, although the deficits are more pragmatic than semantic. Cheney and Seyfarth point out (p. 244) that a Gricean account of meaning presupposes a robust theory of mind; but Grice intended his account as a rational reconstruction of speaker meaning, rather than a psychological analysis (see Camp [2006]). In practice, something less than fully explicit mental-state attribution and closer to mere coordination might suffice for the evolution of a syntactically complex communicative system (Skyrms [2004]), so long as the number of messages that needed to be communicated was sufficiently high (Nowak *et al.* [2000]).

between particular individuals: who is related to whom, and who has, or is likely to, defer to or supplant whom.[21] Thus, there would be no need for other modules, or a general faculty of cognition, to process hierarchically structured representations.

7. CONCLUSION

Thus, we have seen that all six of the features that Cheney and Seyfarth identify are indeed respects in which baboon cognition is plausibly like language. Nonetheless, we shouldn't conclude that baboons think in language, because we get a simpler and more efficient explanation of the distinctive contours of their cognitive abilities and limitations if we hypothesize that their representational system also differs from language in at least one crucial respect: that the component of their thoughts which is hierarchically structured employs a combinatorial principle with a robust function, that of representing dominance. If we adopt this weaker view, then Cheney and Seyfarth's "social origins hypothesis" could be reformulated as the still-substantive claim that the hierarchical structure in the analogous module employed by our ancestors was eventually appropriated to represent other relations, and ultimately to represent something as abstract as predication.[22]

Beyond baboon cognition and the origins of language, we can also draw two more secure, if also more purely methodological, morals. The first is that speculating about the form of thought is a dangerous business. It can seem almost irresistible to assume that thought has a form; but because this claim is usually understood functionally, it cannot be tested directly. Nor is it generally possible to determine the form of thought just by examining structural features, like being discrete, rule-governed, or even hierarchical. Instead, we need to appeal to the particular contents that an agent is and is not capable of representing, and to how they manipulate those contents. Even then, though, inferences from content and processing to form will always be inferences to the best explanation; and such inferences are always vulnerable to the charge of imaginative failure. In this sense, Cheney and Seyfarth's discussion provides a valuable case study.

Second, in general, theorists have neglected alternative representational possibilities because they have too unreflectively subscribed to a dichotomy between imagistic and linguistic formats. One way in which philosophers might usefully contribute to debates about the form of thought is by developing a more general understanding of the combinatorial and referential principles that govern various representational systems and of the relations among those systems. This would enable us to systematically delineate the space of representational possibilities, with special attention to the functional strengths and weaknesses of each system.

[21] Other modules or general cognition might have representations of dominance and kinship; the claim is just that they wouldn't need to represent hierarchical structure.

[22] Though, as several authors (e.g., Pinker and Bloom [1990]) have pointed out, hierarchical organization may be a feature of many complex cognitive systems.