4

DO NONHUMAN ANIMALS HAVE A LANGUAGE OF THOUGHT?

Jacob Beck

1 Introduction

In the second half of the 20th century, behaviorism slowly gave way to the computational and representational paradigm of cognitive science. Human language may have been the first beneficiary, but it wasn't the last. Even much animal cognition came to be routinely explained in computational and representational terms (Gallistel 1990).

One influential, if controversial, idea that accompanied the ascendency of cognitive science is the language of thought hypothesis (LOTH), which maintains that mental representations are formatted like sentences (Fodor 1975). Because we human animals speak a public language, there has always been a special reason to accept LOTH as true of us. Our linguistic utterances are naturally construed as direct translations of our internal thoughts, which suggests that our internal thoughts mirror the structure of their public-language expressions.

When it comes to nonhuman animals (hereafter: animals), this special reason is missing. Insofar as animals communicate, they do so without employing the richly structured public languages that humans employ. One might therefore be tempted to infer that animals' mental representations have a nonlinguistic format – for example, an imagistic, map-like, or analog format. But this conclusion does not follow of necessity. The language of thought hypothesis for animals (LOTHA) could be true even if animals lack a public language in which to express their thoughts.

This chapter has two aims. The first is to review evidence that animals have at least some representations with a nonlinguistic format. The second is to argue that although we don't know enough as of yet to determine whether LOTHA is true, there is a clearly defined research program into the logical abilities of animals that can help to deliver an answer.

2 LOTHA

Sometimes LOTH is interpreted to mean only that cognizers have mental representations that are *compositional*. The representations consist of atomic parts that compose into complexes such that the contents of the atomic parts determine the contents of the complexes in a rule-governed way. But the common refrain that mental states are relations to LOT *sentences* suggests that proponents of LOTH often have something stronger in mind. According to this stronger conception, LOT representations exhibit the same basic representational and compositional properties as paradigmatic sentences (cf. Camp 2007 on "Weak-LOT" vs. "Strong-LOT").

We can make this stronger conception more precise by noting two properties of paradigmatic sentences. First, the constituents of sentences bear an *arbitrary* relation to their referents. There is nothing intrinsic to the English word "dog" that makes it especially well suited to represent dogs as opposed to cats or anything else. By contrast, a picture of a dog is especially well suited to represent dogs because it resembles dogs. Second, sentences have *logical form*. Their basic compositional mechanisms include predication and logical constants, such as negation, disjunction, implication, identity, universal quantification, and existential generalization (cf. Burge 2010a: 542–5 and Burge 2010b on "propositional thought"). Proponents of LOTH often emphasize this feature of sentences. For example, Margolis and Laurence (2007: 562) write that LOTH is committed to a "language-like syntax" that "incorporates, at the very least, a distinction between predicates and subjects, and that includes logical devices, such as quantifiers and variables."

In evaluating LOTHA, I will interpret it in this stronger sense. So interpreted, LOTHA contrasts with other accounts of animal cognition that are compatible with the representational and computational paradigm that dominates cognitive science (and is assumed here). For example, Rescorla (Chapter 3 in this volume) reviews evidence that animals navigate using cognitive maps, which, if interpreted in what he calls the "strict sense," have geometric rather than logical form. Similarly, Gauker (Chapter 2 in this volume) argues that animals' tool use and physical reasoning can be explained by imagistic representations, which are bereft of general concepts of the sort associated with predicates. If these hypotheses are correct, animals have at least some cognitive representations that defy LOTHA.

LOTHA is of interest, in part, because it provides a way to understand, from within the representational and computational paradigm of contemporary cognitive science, the question whether human and animal cognition differ in kind or only degree. If humans have a LOT but animals do not, then there is a clear sense in which human cognition has a fundamentally different representational format from animal cognition. By contrast, if humans and animals both have a LOT, then it remains an open possibility that, at least from the perspective of contemporary cognitive science, human and animal cognition differ only in degree.

3 Analog magnitude representations

Insofar as we can fully explain animal cognition by appeal to representations with a nonlinguistic format, such as cognitive maps and imagistic representations, we have reason to be skeptical of LOTHA. In this section, I want to briefly review evidence for one additional, but oft-overlooked type of nonlinguistic representation: *analog magnitude representations*. Hundreds of studies indicate that a wide range of animals, including mammals, birds, and fish, can represent numerosities, durations, rates, distances, sizes, and other worldly magnitudes (Gallistel 1990; Beck 2015). As an illustration, I'll review a now-classic set of experiments on rats.

After training rats to press the left lever in response to a two-second sequence of two tones and the right lever in response to an eight-second sequence of eight tones, Meck and Church (1983) tested the rats on intermediate stimuli, either holding duration constant at four seconds while varying the number of tones or holding number constant at four tones while varying the duration of tones (Figure 4.1). When duration was held constant, the rats were most likely to press the left lever in response to two or three tones and most likely to press the right lever in response to five, six, or eight tones, suggesting that they represented the numerosity of the tones. When number was held constant, the rats were most likely to press the left lever in response to a two- or three-second tone and most likely to press the right lever in response to a five-, six-, or eight-second

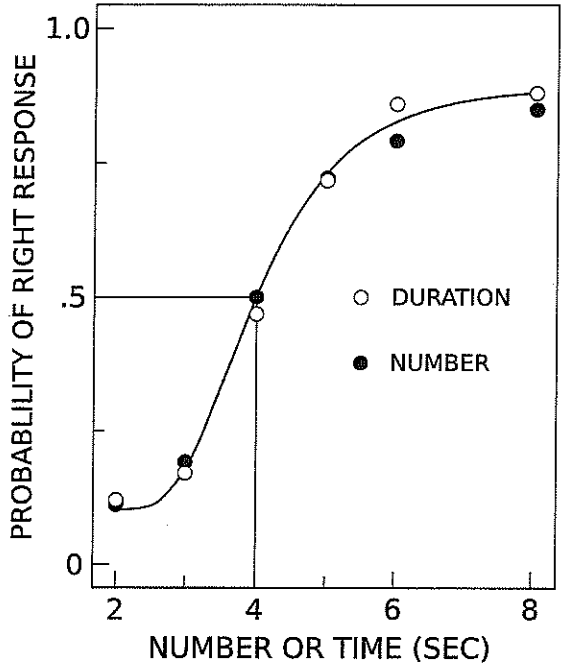

Figure 4.1 The probability that the rat will press the right lever as a function of the duration or number of tones. Redrawn from Meck and Church (1983).

tone, suggesting that they also represented the duration of the tones. By contrast, the rats were equally likely to press the right and left levers when presented with a four-second sequence of four tones. At first blush, it may seem surprising that rats treat four (the geometric mean), rather than five (the arithmetic mean), as the point of subjective equality between two and eight. This result makes perfect sense, however, if the rats represent magnitudes in terms of ratios, since 2:4 = 4:8.

Do the rats in these experiments really represent duration and numerosity? Or can their behavior be explained more simply in terms of low-level acoustic properties that correlate with duration and numerosity? One reason to think rats represent duration and numerosity themselves is that they transfer their training across modalities. For example, when rats are trained on auditory stimuli as summarized above and then presented with flashes of light (visual stimuli), they'll press the left lever when presented with a two-second light and press the right lever when presented with an eight-second light (Meck and Church 1982). Since vision and audition operate over disparate low-level sensory stimuli, these results support the hypothesis that rats really are glomming onto the abstract properties of duration and numerosity.

There are at least three related reasons to think that animals' magnitude representations are nonlinguistic. First, they have a nonarbitrary, analog format (Beck 2015). Animals' magnitude discriminations are ratio sensitive: as the ratio of two magnitudes approaches one, the ability to discriminate them deteriorates. Rats thus find it easier to discriminate three tones from four tones than four tones from five tones. (This is why the rats in Meck and Church's study treat five as more similar to eight than to two.) On the assumption that magnitude representations involve some internal magnitude (say, neural firing rate) that increases or decreases in proportion to the magnitude represented (say, number of tones), and are thus a direct analog of the magnitudes they represent, this ratio sensitivity is exactly what one would predict. As the ratio of two external magnitudes approaches one, the ratio of the internal magnitudes will follow suit. Thus, assuming some noise in the system, the internal magnitudes will themselves become increasingly difficult to discern, leading to discrimination errors of the magnitudes they represent. Because of their arbitrary referential relation, paradigmatic linguistic representations are not analog in this sense. For example, there is no property of Arabic numerals or English number words that increases as the number represented increases: "9" is not intrinsically more similar to "7" than to "5," and "nine" is not intrinsically more similar to "seven" than to "five."

Second, because they exhibit logical form, LOT representations are systematically recombinable like the words in a sentence (Fodor 1987). Thus, given LOTH, if you can think that Amy likes Ben and that Ben likes Cam, then you can also think that Amy likes Cam. It is doubtful that analog magnitude representations are systematically recombinable in this way. For example, although a rat can form representations with something like the content that 9 tones are fewer than 18 tones and that 10 tones are fewer than 20 tones, it is questionable whether it can form representations with anything like the content that 9 tones are fewer than 10 tones or that 18 tones are fewer than 20 tones (Beck 2012a). The reason, once again, is that magnitude discriminations are ratio sensitive, and as the ratio of two magnitudes approaches one, the ability to discriminate them deteriorates. When the ratio of two numbers is close enough to one, rats cannot reliably represent the numbers as distinct.

A final reason to doubt that analog magnitude representations have a linguistic format is that the computations they enter into can be fully described without any appeal to logical constants such as negation, disjunction, or identity. Rather, the computations that analog magnitude representations enter into are arithmetic. They include addition, subtraction, multiplication, division, and comparison (or primitive analogs thereof). For example, animals might divide a numerical representation by a duration representation to yield a representation of the rate of return of a given feeding source, and then compare the result to its representation of the rate of return of a second feeding source in order to help it decide which source to visit. Logical constants play no role in this explanation (Beck 2015).

4 Logical inference as a test for LOTHA

The existence of various types of nonlinguistic representations – cognitive maps, mental images, analog magnitude representations – places pressure on LOTHA. The more that intelligent animal behavior can be explained through various types of nonlinguistic representation, the less theoretical work there is left over for LOTHA to do.

Still, we are a long way from being certain that all of animal cognition can be explained by appeal to nonlinguistic representations, and the mere fact that some aspects of animal cognition have been so explained is hardly reason to conclude that animal minds are bereft of any sentence-like representations. It is thus worth considering whether there aren't more direct ways to test LOTHA.

I want to suggest that we can gain some traction on this issue from the idea that the inferences a thinker is capable of undertaking form a window into the structure of the thinker's thoughts (Evans 1985: 337; Burge 2010a: 542–7; Beck 2012b: 225–6). If a thinker's cognition is supported by a LOT, and thus by representations with logical form, we should expect that thinker to be capable of engaging in logical inferences. This suggests that we can test LOTHA by testing an animal's facility with logic.

In the next section, I will review one experimental paradigm that has been put to use to test the logical abilities of animals. First, however, I want to note, if only to set aside, five complications that attend to the strategy of using logical inference as a test for LOTHA.

The first complication derives from alternatives to LOTH for which talk about the "format" of a thinker's mental representations seems to have no place. Chief among these are connectionist networks. It is trivial to get connectionist networks to compute logical functions. But because the representations in a connectionist network can be distributed across the network, questions about their format are arguably misplaced.

This complication deserves more attention than I have space for here, so I'll have to settle for three brief comments. First, while it is uncontroversial that connectionist networks can compute logical functions, it is far less clear whether they can do so without *implementing* a LOT.[1] Second, my main concern here is not to resolve the question of whether LOTHA is true, but to show how it can be approached empirically, and connectionist models can surely be empirically tested. Third, whether or not the ability to draw logical inferences is evidence *for* LOTHA, the inability to draw logical inferences is surely evidence *against* LOTHA. For those interested in the status of LOTHA, there is thus good reason to examine animals' logical abilities regardless of what one thinks of the relation between LOTH and connectionism.

A second complication derives from *hybrid* formats that combine linguistic and nonlinguistic elements, such as Venn diagrams that are outfitted with special symbols (Peirce 1933; Shin 1994) or maps that have markers for negation or disjunction (Camp 2007). Such hybrid formats have been pursued in psychology under the guise of mental-models theory, which posits imagistic representations that are supplemented with arbitrary symbols (Johnson-Laird 2006). In some sense, hybrid representations such as mental models are sentence-like since they have arbitrary symbols that represent logical constants. But they also have components (e.g., imagistic elements) that are nonlinguistic. As with connectionist networks, it thus isn't always clear whether hybrid mental representations ought to count as implementing LOTH or competing with it. But one thing that *is* fairly clear is that, as with connectionist models, mental models can be empirically evaluated. It should thus be possible to empirically distinguish mental models from a purer form of LOTHA. Furthermore, whatever one thinks of the relation between LOTHA and mental models, the inability to draw logical inferences would surely count as evidence *against* LOTHA.

Third, performance on logical reasoning tasks is a function not only of a thinker's logical concepts, but also of various additional factors such as attention, working memory, and background beliefs. Thus, while possessing logical concepts plausibly entails the *capacity* to draw certain inferences, it does not entail the *successful exercise* of that capacity. A thinker could possess a logical concept and yet fail this or that logical reasoning task. But by testing a thinker in a wide variety of inference tasks that require the same logical concept but place varying demands on performance factors, it should be possible to tease conceptual competence apart from performance.

A fourth complication is that the kingdom of animals is diverse, and there is every reason to think that reasoning abilities will vary across species. This means that in the long run, empirical research into the logical abilities of animals should sample from a diversity of species. Here, however, I focus on the question of how the logical abilities of *any* nonlinguistic species can be empirically tested.

A final complication is that logic is not monolithic. There are many forms of logical inference. In the long run, there is no reason to limit an inquiry to this or that form of inference. If we want a full picture of the contours of the format of animal cognition, we should test as large a variety of inferences as is feasible. In the more medium term, however, we would do best to focus our efforts on nontrivial but fairly basic forms of inference. One such form that has the advantage of having been extensively studied is the disjunctive syllogism, *modus tollendo ponens*, or reasoning by exclusion.[2]

P or Q
Not P
Q

In the following section, I focus on how we can tell whether animals can execute inferences of this form.

5 One cup, two cups, three cups, four

A now-standard tool researchers have used to probe animals' capacity to reason by exclusion is the "two-cups task." An animal is presented with two opaque cups, A and B, and shown that both are initially empty. The animal sees the experimenter bait one of the cups with food, but a barrier prevents it from seeing which cup was baited. In the crucial condition, the experimenter reveals that cup A is empty and then allows the animal to choose a cup for inspection. If the animal is capable of executing a disjunctive syllogism, it should choose cup B. It should reason: *The food is either in cup A or in cup B; it's not in cup A; so it's in cup B.*

Several species have succeeded at this task, including great apes (Call and Carpenter 2001; Call 2004), monkeys (Grether and Maslow 1937; Petit et al. 2015; Marsh et al. 2015), ravens (Schloegl et al. 2009), and dogs (Erdőhegyi et al. 2007). Does that mean that these species can execute disjunctive syllogisms?

Surely not. No single task is ever sufficient to establish a conceptual ability. There are, as always, competing interpretations that need to be evaluated. In the remainder of this section, I discuss three such interpretations and indicate how further experiments might address them. My aim is to show how empirical evidence can, in principle, be used to decide among competing interpretations, not to defend any one interpretation in particular.

5.1 Avoid the empty cup

First, animals could succeed on the two-cup task by following the simple heuristic: *avoid the empty cup*. According to this interpretation, it's not that animals search in cup B because they infer that cup B has food. Rather, they search in cup B because they see that cup A is empty, want to avoid empty cups, and cup B is the only other hiding place in view.

In order to evaluate this interpretation, Call (2016) has developed a three-cup task. Subjects see the experimenter place food in either cup A or cup B, but not in cup C. The experimenter then reveals that cup A is empty. If the subjects only avoid the empty cup, they should have no preference as between cups B and C. But if they execute a disjunctive syllogism, they should choose cup B, reasoning that it was in A or B, not in A, and thus must be in B. As of this writing, Call was still in the process of collecting data on great apes, and so the results are not yet known. But the design of the study clearly illustrates how the avoid-the-empty-cup heuristic can be empirically distinguished from exclusionary reasoning.

5.2 Maybe A, maybe B

A second competing explanation of the two-cup task, suggested by Mody and Carey (2016), is that subjects represent each cup as a possible location of food and then eliminate cup A when it is shown to be empty. Unlike in a disjunctive syllogism, this does not lead subjects to conclude that the food is definitely in cup B. They don't update their representations of cup B at all. It's just that once cup A is eliminated, cup B is the only location remaining in which food is represented as possibly being present, and so subjects choose it. Mody and Carey call this the "maybe-A-maybe-B" interpretation.

Notice that the three-cup task doesn't distinguish the maybe-A-maybe-B interpretation from the disjunctive-syllogism interpretation. When subjects see food hidden in cup A or cup B, they represent the food as maybe in A and maybe in B. By contrast, they do not represent food as maybe in C. So when A is shown to be empty, they eliminate that cup, leaving cup B as the only cup represented as possibly containing the food. So both the maybe-A-maybe-B interpretation and the disjunctive-syllogism interpretation predict that subjects should choose cup B in the three-cup task.

To make headway, Mody and Carey add a fourth cup. The child (Mody and Carey ran their study on human children) sees the experimenter place one sticker in either cup A or cup B and another sticker in either cup C or cup D. The experimenter then reveals that cup A is empty and allows the child to select a cup for search. If the maybe-A-maybe-B interpretation were correct and the child never updated her representation of B upon learning that A is empty, she should be equally likely to choose B, C, or D. But if the disjunctive-syllogism interpretation were correct, then when the child was shown that A is empty, she should update her representation of B as definitely containing a sticker, and should thus prefer to search in that cup. Mody and Carey found that 2.5-year-old children failed at this task even though they succeeded at the two-cup task. By contrast, three- to five-year-old children succeeded at both tasks. The four-cup task has yet to be run on nonhuman animals, but it surely could be in principle.

5.3 Probabilistic reasoning

A third alternative to the disjunctive-syllogism interpretation, articulated by Rescorla (2009), maintains that animals use cognitive maps to represent the possible locations of objects, tag the maps with subjective probabilities of how likely they are to be accurate, and then update those probabilities in accordance with Bayes' Law. In the two-cup task, when animals initially see the food hidden, they represent cups A and B as equally likely to contain food. When cup A is then shown to be empty, they raise the probability that food is in cup B and lower the probability that it's in A. But contrary to the disjunctive-syllogism interpretation, there is no deterministic inference that cup B definitely contains food. Yet because Rescorla further assumes that animals will conform to expected utility theory and thus search in the location that is most likely to contain food, their behavior will be indistinguishable from that of a subject who reasons by way of the disjunctive syllogism. When cup A is shown to be empty, they'll always choose cup B.

Of all the alternative interpretations we have thus far considered, this one seems to be the most difficult to test. Mody and Carey (2016: 46) claim that their results from the four-cup task tell against it since "3- to 5-year-old children chose the target cup [cup B] just as often in test trials as they did in training trials, in which they could directly observe that a sticker was being hidden there." But if the children are conforming to expected utility theory, that's exactly what one would predict. In the four-cup study, cup B is the most likely to contain a sticker. So if children approach the task using probabilistic representations and a decision procedure that has them select the greatest expected payoff, they'll choose cup B as surely as if they saw the sticker there directly.

How, then, can we test between the probabilistic and deductive alternatives? As Rescorla knows, the proposal that thinkers use probabilistic representations only generates predictions about behavior when tied to an assumption about how those representations figure into practical reasoning. Rescorla, as we've seen, assumes a version of expected utility theory, whereby thinkers try to maximize their expected payoffs. Yet this, too, is an empirical assumption – one that can be, and has been, subjected to its own tests. As it happens, one of the more interesting findings to emerge from the animal cognition literature over the past few decades is that there are circumstances in which animals systematically *violate* expected utility theory.

A fish tank contains two feeding tubes: one that releases a food morsel every second and another that releases a food morsel every two seconds. Fish that obeyed expected utility theory would spend all of their time in front of the first tube. But that's not what fish actually do. Instead, they spend two-thirds of their time in front of the first tube and one-third of their time in front of the second tube. In other words, they adopt a probability matching decision procedure (Godin and Keenleyside 1984). While initially puzzling, this procedure has clear selection benefits for group foragers. A lone fish that adopted a probability matching strategy while its peers all conformed to expected utility theory would reap a superior harvest.

The tendency to follow a probability matching strategy in certain circumstances is widespread throughout the animal kingdom. (Even humans sometimes display this tendency, which is one reason casinos are so profitable.) This gives us a wedge to distinguish between the probabilistic and deductive interpretations. All we need to do is run a version of the four-cup task for animals when we have independent evidence that they are disposed to adopt a probability matching decision procedure over their probabilistic representations.

For example, we might begin by putting up a barrier that covers all four cups and hiding food in a way that prevents the animal from telling which of cups A, B, C, or D it is hidden in. We could then privilege cup A by hiding the food in it on 40 percent of trials (and hiding the food in cups B, C, and D on 20 percent of trials each). On each trial, the animal chooses a cup and keeps the food if it guesses correctly. Over time, we could see how the animal responds. If, as seems likely, it adopts a probability matching strategy, it should eventually learn to choose cup A 40 percent of the time and each of the other cups 20 percent of the time. We can then continue to a version of the four-cup task in which the animal sees one piece of food hidden in A or B, a second piece hidden in C or D, and is then shown that A is empty. Given Bayes' Law and reasonable background assumptions, the animal should assign a very high probability to food being in cup B, and lower probabilities of roughly .5 each to food being in cup C or in cup D. Given a probability matching strategy, it should thus choose cup B more often than it chooses C and more often than it chooses D, but not more often than it chooses C or D. If the animal is executing a disjunctive syllogism, however, it should choose cup B almost all of the time (or, allowing for performance errors, about as often as when it directly sees the food hidden in B).[3]

6 Conclusion

We began with the question: Is LOTHA true? That is: do animals have sentence-like mental representations – mental representations with an arbitrary referential relation and logical form? As we saw in our discussion of analog magnitude representations, there is evidence that at least some animal mental representations are *not* sentence-like. But of course that doesn't settle the question of LOTHA since animals could have sentence-like representations in addition. We thus sought out a more direct way of evaluating LOTHA.

This led us to examine experiments that test whether animals are capable of logical inferences such as the disjunctive syllogism. Such experiments are always open to alternative interpretations. Thus, no single experiment can hope to settle the matter on its own. But by developing a series of experiments that are designed to pit two or more competing interpretations against one another, we saw how we can acquire empirical evidence that enables us to rationally decide among competing interpretations. As a result, we can begin to see – modulo the challenges identified in Section 4 – how empirical methods can help to reveal whether LOTHA is true.[4]

Notes

1 This is a delicate issue. For an overview and references, see Aydede (2015: §8).

2 One indication of the nontriviality of this form of inference is that it presupposes the conceptual resources to express all possible truth-functions in propositional logic.

3 Bermúdez (2003) suggests a fourth alternative interpretation that I lack the space to properly discuss. Animals could solve the various cup tasks without a genuine negation operator by employing contrary representations, such as *present* and *absent*. For a response, see Burge (2010b: 62–3).

4 Thanks to Kristin Andrews, Roman Feiman, and Christopher Peacocke for helpful comments, and to Matthew Cutone for assistance with Figure 4.1.

Further reading

S. Shettleworth, *Cognition, Evolution, and Behavior* (New York: Oxford University Press, 2010) is a comprehensive textbook on animal cognition that covers both behaviorist and computational/representational approaches. C. J. Völter and J. Call, "Causal and Inferential Reasoning in Animals," in G. M. Burghardt, I. M. Pepperberg, C. T. Snowdon, and T. Zentall (Eds.), *APA Handbook of Comparative Psychology Vol. 2: Perception, Learning, and Cognition*, pp. 643–671 (Washington, DC: American Psychological Association, 2017), reviews empirical research on exclusionary reasoning. Chapter 4 of K. Andrews, *The Animal Mind: An Introduction to the Philosophy of Animal Cognition* (New York: Routledge, 2015) contains an accessible introduction to many of the issues discussed here, including logical reasoning in animals.

References

Aydede, M. (2015) "The Language of Thought Hypothesis," in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition), http://plato.stanford.edu/archives/fall2015/entries/language-thought/

Beck, J. (2012a) "The Generality Constraint and the Structure of Thought," *Mind* 121: 563–600.

——— (2012b) "Do Animals Engage in Conceptual Thought?" *Philosophy Compass* 7: 218–229.

——— (2015) "Analogue Magnitude Representations: A Philosophical Introduction," *British Journal for the Philosophy of Science* 66: 829–855.

Bermúdez, J. L. (2003) *Thinking Without Words*, New York: Oxford University Press.

Burge, T. (2010a) *Origins of Objectivity*, New York: Oxford University Press.

——— (2010b) "Steps Towards Origins of Propositional Thought," *Disputatio* 4: 39–67.

Call, J. (2004) "Inferences About the Location of Food in the Great Apes (Pan Paniscus, Pan Troglodytes, Gorilla Gorilla, and Pongo Pygmaeus)," *Journal of Comparative Psychology* 118: 232–241.

——— (2016) "Avoiding the Empty Cup: Is There Anything More to Apes' Inferential Reasoning?" Talk delivered at the Origins of Logical Reasoning Workshop, York University, Toronto.

Call, J., and Carpenter, M. (2001) "Do Apes and Children Know What They Have Seen?" *Animal Cognition* 3: 207–220.

Camp, E. (2007) "Thinking With Maps," *Philosophical Perspectives* 21: 145–182.

Erdőhegyi, Á., Topál, J., Virányi, Z., and Miklósi, Á. (2007) "Dog-Logic: Inferential Reasoning in a Two-Way Choice Task and Its Restricted Use," *Animal Behaviour* 74: 725–737.

Evans, G. (1985) "Semantic Theory and Tacit Knowledge," in G. Evans (ed.) *Collected Papers*, New York: Oxford University Press, 322–342.

Fodor, J. (1975) *The Language of Thought*, New York: Crowell.

——— (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, Cambridge, MA: MIT Press.

Gallistel, C. R. (1990) *The Organization of Learning*, Cambridge, MA: MIT Press.

Godin, J. J., and Keenleyside, M. H. A. (1984) "Foraging on Patchily Distributed Prey By a Cichlid Fish (Teleosti, Cichlidae): A Test of the Ideal Free Distribution Theory," *Animal Behaviour* 32: 120–131.

Grether, W. F., and Maslow, A. H. (1937) "An Experimental Study of Insight in Monkeys," *Journal of Comparative Psychology* 24: 127–134.

Johnson-Laird, P. (2006) *How We Reason*, New York: Oxford University Press.

Margolis, E., and Laurence, S. (2007) "The Ontology of Concepts – Abstract Objects or Mental Representations?" *Noûs* 41: 561–593.

Marsh, H. L., Vining, A. Q., Levendoski, E. K., and Judge, P. G. (2015) "Inference by Exclusion in Lion-Tailed Macaques (Macaca Silenus), a Hamadryas Baboon (Papio Hamadryas), Capuchins (Sapajus Apella), and Squirrel Monkeys (Saimiri Sciureus)," *Journal of Comparative Psychology* 129: 256–267.

Meck, W. H., and Church, R. M. (1982) "Abstraction of Temporal Attributes," *Journal of Experimental Psychology: Animal Behavior Processes* 8: 226–243.

——— (1983) "A Mode Control Model of Counting and Timing Processes," *Journal of Experimental Psychology: Animal Behavior Processes* 9: 320–334.

Mody, S., and Carey, S. (2016) "The Emergence of Reasoning by the Disjunctive Syllogism in Early Childhood," *Cognition* 154: 40–48.

Peirce, C. S. (1933) *Collected Papers*, Cambridge, MA: Harvard University Press.

Petit, O., Dufour, V., Herrenschmidt, M., De Marco, A., Sterck, E. H., and Call, J. (2015) "Inferences About Food Location in Three Cercopithecine Species: An Insight Into the Socioecological Cognition of Primates," *Animal Cognition* 18: 821–830.

Rescorla, M. (2009) "Chrysippus' Dog as a Case Study in Non-Linguistic Cognition," in R. Lurz (ed.) *The Philosophy of Animal Minds*, New York: Cambridge University Press, 52–71.

Schloegl, C., Dierks, A., Gajdon, G. K., Huber, L., Kotrschal, K., and Bugnyar, T. (2009) "What You See Is What You Get? Exclusion Performances in Ravens and Keas," *PLoS ONE* 4: e6368.

Shin, S. (1994) *The Logical Status of Diagrams*, Cambridge: Cambridge University Press.