Chrysippus's Dog as a Case Study in Non-Linguistic Cognition

Michael Rescorla

Abstract: I critique an ancient argument for the possibility of non-linguistic deductive inference. The argument, attributed to Chrysippus, describes a dog whose behavior supposedly reflects disjunctive syllogistic reasoning. Drawing on contemporary robotics, I urge that we can equally well explain the dog's behavior by citing probabilistic reasoning over cognitive maps. I then critique various experimentally-based arguments from scientific psychology that echo Chrysippus's anecdotal presentation.

§1. Language and thought

Do non-linguistic creatures think? Debate over this question tends to calcify into two extreme doctrines. The first, espoused by Descartes, regards language as necessary for cognition. Modern proponents include Brandom (1994, pp. 145-157), Davidson (1984, pp. 155-170), McDowell (1996), and Sellars (1963, pp. 177-189). Cartesians may grant that ascribing cognitive activity to non-linguistic creatures is instrumentally useful, but they regard such ascriptions as strictly speaking false. The second extreme doctrine, espoused by Gassendi, Hume, and Locke, maintains that linguistic and non-linguistic cognition are fundamentally the same. Modern proponents include Fodor (2003), Peacocke (1997), Stalnaker (1984), and many others. Proponents may grant that non-linguistic creatures entertain a narrower range of thoughts than us, but they deny any principled difference in kind.[1]

An intermediate position holds that non-linguistic creatures display cognitive activity of a fundamentally different kind than human thought. Hobbes and Leibniz favored this intermediate position. Modern advocates include Bermudez (2003), Carruthers (2002, 2004), Dummett (1993, pp. 147-149), Malcolm (1972), and Putnam (1992, pp. 28-30). Proponents may grant that our lower-level cognition resembles the mental activity of languageless creatures, but they insist that we *also* manifest higher-level cognition unavailable to such creatures. The main challenge facing such a view is to describe non-linguistic cognitive processes that differ in a principled way from higher-level human thought.

I will try to meet this challenge by exploring a putative example of non-linguistic cognition. Tolman (1948) introduced the notion of *cognitive map* to explain how rats in a laboratory maze take detours and shortcuts to reach destinations. Although Tolman's analysis proved controversial, many psychologists have followed him in proposing that human and animal navigation exploits cognitive maps. These have representational content: they represent the world as being a certain way, so we can evaluate them as veridical or non-veridical. Moreover, they participate in rational mental processes that update them based on perception and exploit them during locomotion. Thus, cognitive maps are genuinely cognitive. Yet they differ from higher-level human thought in two crucial ways: they do not have logical form, and they do not figure in deductive inference.

To illustrate the explanatory potential of cognitive maps, I will deploy them against a venerable philosophical argument for languageless thought and reasoning. Sextus Empiricus, who credits the argument to Chrysippus, presents it as follows:

> [Chrysippus] declares that the dog makes use of the fifth complex indemonstrable syllogism when, on arriving at a spot where three ways meet…, after smelling at the two roads by which the quarry did not pass, he rushes off at once by the third without stopping to smell. For, says the old writer, the dog implicitly reasons thus: "The animal went either by this road, or by that, or by the other: but it did not go by this or that, therefore he went the other way."[2]

Many commentators, including Aquinas, Gassendi, Jevons, Montaigne, and even King James I of England, have argued on this basis that non-linguistic creatures execute logical reasoning. More recently, Glock (2000) and Horgan and Tienson (1996, p. 93) concur.

Despite millennia of discussion, opponents of non-linguistic deduction have not yet answered this argument satisfactorily. They have not shown how to explain the described phenomena without imputing logical reasoning to Chrysippus's dog. I seek to fill this gap. My proposed explanation, which draws heavily upon research from contemporary robotics, cites rational, non-deductive mental processes defined over cognitive maps. The explanation shows that we can accommodate Chrysippus's dog without assimilating animal minds to human minds. Yet it also militates against Cartesianism, because it invokes representational mental states and rational processes defined over those states. Thus, my treatment illustrates the explanatory resources of an intermediate position that countenances non-linguistic cognition while sharply distinguishing it from linguistic cognition.

§2. Logical form and cognitive maps

I focus on two crucial features of human propositional attitudes: they have logical form, and they participate in deductive reasoning sensitive to that form. Both features have been recognized since Aristotle, but Frege profoundly enhanced our understanding of them. The Fregean tradition, culminating in Tarski's work, analyzes how truth-conditions of logically complex thoughts depend on semantic values of their parts. It thereby illuminates the semantically and inferentially relevant structure of human thought (Burge 2005, pp. 12-26). The paradigmatic structural elements are compositional mechanisms of the predicate calculus: truth-functional connectives, universal and existential quantifiers, and predication. Additional compositional devices, such as modal operators and generalized quantifiers, have also received compelling semantic analysis over the past century.

My question is whether non-linguistic animals enter into mental states with logical form, where "logical form" minimally includes the familiar compositional mechanisms of the predicate calculus. As already indicated, I will study a single putative case of non-logical cognition: spatial representation. The phrase "cognitive map" appears frequently in contemporary psychology and philosophy. As Bermudez (1998, pp. 203-207) and Kitchin (1994) document, researchers associate it with diverse meanings. Gallistel (1990, p. 103) defines a cognitive map as "a record in the central nervous system of macroscopic geometric relations among surfaces in the environment used to plan movements through the environment." This definition remains neutral about the extent to which cognitive maps resemble ordinary concrete maps. I will employ a more literal usage: a cognitive map is a mental representation that represents geometric features of the physical environment and that employs the same basic representational mechanisms as a concrete map. On this usage, which seems close to that of O'Keefe and Nadel (1978, pp. 80-96, pp. 389-391), cognitive and concrete maps share a common representational format. Specifically, they have comparable compositional structures. A fuller elucidation of my usage would require systematic discussion of the compositional mechanisms underlying concrete cartographic representation. Unfortunately, those mechanisms are not completely understood. Despite the efforts of such authors as Casati and Varzi (1999), Pratt (1993), and Sloman (1978), we have no canonical cartographic semantics analogous to Tarski's semantics for the predicate calculus. Nevertheless, my definition of "cognitive map" seems clear enough for present purposes.

I will assume that cognitive maps do not have logical form. This assumption follows from two premises: first, ordinary concrete maps do not have logical form; second, cognitive maps and concrete maps employ the same basic representational mechanisms. Regarding the first premise, most philosophers who address the matter agree that maps do not express negation, disjunction, the conditional, or the quantifiers (Fodor 1991, p. 295), (Millikan 1993, p. 302), (Pylyshyn 2003, p. 424-5).[3] This verdict merits more extended defense than it typically receives, but it seems plausible enough for us to assume it here. Whether concrete maps can express conjunction strikes me as more debatable, but I will again assume a negative verdict. More debatable still is the thesis that concrete maps do not feature predication. One might propose that attaching a map symbol (such as a symbol denoting mountains) to a map coordinate amounts to predicating the corresponding property of the corresponding spatial location. Casati and Varzi (1999) develop a formal cartographic semantics designed to incorporate this proposal. If the proposal is correct, then maps feature rudimentary logical form akin to atomic sentences. I attack this proposal in (Rescorla, in press b), arguing that maps do not feature predication, as construed within Fregean or Tarskian semantics.

However, suppose we were to concede that maps have predicative structure. A system of cartographic mental representations would still not support deductive reasoning. It would not allow familiar inference patterns like modus ponens, argument by cases, or universal instantiation. Not even the laws of identity would apply, since concrete maps do not feature an identity sign. Thus, even if we were to concede that cognitive maps have predicative structure, a principled distinction would separate mental processes defined over them from higher-level human cognition.

Given that cognitive maps do not participate in logical inference, can they figure in *any* rational mental processes? Philosophers often suggest that they cannot, on the grounds that rationality requires logically structured mental states (Pylyshyn 1984, pp. 195-196; Rey 1995, p. 203; Devitt 2006, pp. 146-147). If so, then mental activity defined solely over cartographic mental representations is not rational. Accordingly, many philosophers will resist calling the activity "cognitive."

To address these worries, I want to discuss an instructive case study: Chrysippus's dog. I will analyze this case study through a detailed description of rational mental activity defined over cognitive maps.

§3. Explaining the phenomena

Discussions of Chrysippus's dog typically choose among four strategies:

(1) Treat the dog as executing a deductive inference.

This strategy is probably the most popular. Fry (2002), who works within a cybernetic framework, develops it in fairly rigorous detail. Since logical reasoning presupposes logical structure, (1) requires us to ascribe logically structured mental states to the dog.

(2) Attribute logical reasoning to the dog, but construe the attribution instrumentally.

According to (2), the dog does not "really" execute deductive inference. When we impute logical reasoning to the dog, we are *interpreting* its behavior, not describing observer-independent states and processes. Ironically, Sorabji (1993, p. 26) suggests that Chrysippus himself advocated (2). Similarly, Dennett (1996, p. 115) suggests that our attribution of logical reasoning to Chrysippus's dog is just a matter of adopting the "intentional stance."

(3) Do not attribute logical reasoning to the dog. Instead, maintain that the dog records relevant sensory observations regarding the third path.

Plutarch favored (3): "it is perception itself, by means of track and spoor, which indicates the way the creature fled; [the dog] does not bother with disjunctive and copulative propositions" (1957, 969.a-b). Samuel Coleridge adopted a similar analysis. More recently, Gärdenfors maintains that the dog "could have smelled the scent so clearly along the third path that it did not need to do any sniffing" (2003, p. 71).

(4) Grant that the dog records no additional relevant observations beyond those mentioned by Chrysippus. Explain the dog's behavior by citing observer-independent mental processes distinct from logical reasoning.

Proponents of (4) include Philo of Alexandria, Basil of Caesarea, and Ambrose. For instance, Ambrose cites not logical inference but rather "the training given by nature" (1961, vi.4.23).[4]

Strategy (3) is inadequate. It denies what Chrysippus takes as a datum: that the dog forms no relevant observations of the third path. Setting aside whether Chrysippus himself saw a dog behave as he describes, many philosophers seem convinced that non-linguistic animals routinely exhibit similar behavior. To satisfy these philosophers, we must show that the behavior could result from a mental process other than logical reasoning.

Strategy (2) is appealing only if we accept an instrumentalist or "interpretivist" approach to intentionality, like that espoused by Brandom, Davidson, or Dennett. I reject any such approach. From the realist perspective that I favor, creatures enter into representational mental states that depict the world as being a certain way, states whose representational contents do not result from interpretation by an observer. Cognitive psychology should isolate laws governing how representational mental states interact with one another, with perceptual inputs, and with behavioral outputs. As Fodor (1981, pp. 100-123) argues, there is no more reason to adopt an instrumentalist stance towards the theoretical posits of cognitive psychology than towards those of any other science. Thus, we should reject an instrumentalist treatment of Chrysippus's dog.

Strategy (4) is more promising than (2) and (3). The problem is that no one has developed it satisfactorily. Vague appeals to "nature" do not suffice. We must describe a psychological mechanism that differs in a principled way from logical reasoning, and we must show that the proposed mechanism generates the desired behavior.[5]

In my view, an adequate development of strategy (4) should satisfy three constraints. First, it should predict the desired behavior, rather than dismissing Chrysippus's description in the style of strategy (3). Second, it should support appropriate counterfactuals about how the dog would have behaved had circumstances been different. For instance, it should support the counterfactual: if the prey had traversed the second path, then the dog would have chosen that path rather than the third. Similarly, it should support the counterfactual: if the dog had sniffed the second and third paths without detecting the prey, then it would have chosen the first path without bothering to sniff. Finally, a good account should isolate a general cognitive mechanism that produces the dog's behavior and that is deployable in diverse circumstances. To take an absurd example, we should not posit an innate "hunting at a three-way fork in the road" cognitive module with the following property: if the dog pursues its prey to a three-way fork in the road, and if the dog detects no signs that the prey traversed two of the three forks, then the dog immediately proceeds down the third fork. This putative explanation is unsatisfying, because it cites an ad hoc mental module rather than a general mental capacity applicable in various environmental contexts.

Strategy (1) satisfies all three constraints. I will develop an approach that satisfies the constraints while eschewing logical form and logical inference. My proposal is that Chrysippus's dog performs a probabilistic inference over the space of cognitive maps.

§4. Bayesian reasoning over cognitive maps

Bayesian decision theory is a formal framework for modeling probabilistic reasoning and decision-making. It represents a subject's "degree of belief" in various hypotheses through a subjective probability distribution p. Mathematically precise rules dictate how to update p in light of new evidence and how to act based upon p and one's utilities. Given that p measures degree of belief, the question naturally arises: belief in what? It might seem that any adequate answer will cite sentences, propositions, or their ilk.[6] Accordingly, philosophers often present subjective probability distributions as defined over logically structured entities. But Bayesianism is more general than this. It presupposes only an hypothesis space satisfying certain closure constraints.[7] Elements of the hypothesis space must represent possible states of the world, but they need not have logical structure. In particular, they might be cognitive maps. Thus, we may posit a probability distribution defined not over logically structured representations but over cartographic mental representations.

To illustrate, I will present a Bayesian-cum-cartographic model of Chrysippus's dog. My treatment deploys ideas from probabilistic robotics. Indeed, one of my unofficial goals is to publicize this important field, which philosophers have largely ignored.

I assume that Chrysippus's dog (henceforth *D*) hunts its prey (henceforth *X*) by updating and consulting a probability distribution defined over the space of possible cognitive maps. When *D* reaches the crossroads, it recognizes three relevant possible states of the world, represented by three cognitive maps, M_1, M_2, and M_3, that correspond respectively to the following three concrete maps:

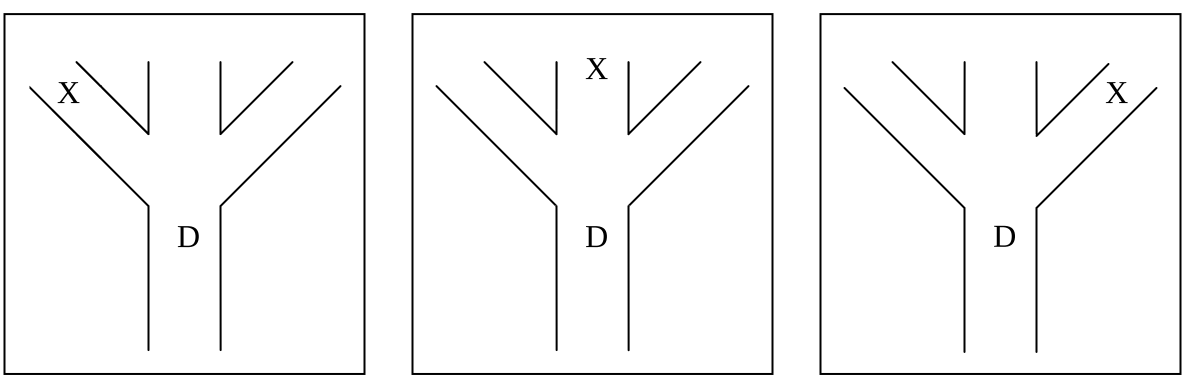

M_1 M_2 M_3

M_1, M_2, and M_3 do not purport to represent distances accurately, but they purport to capture relations of location and connectedness. More technically: they purport to capture topological properties but not metric properties. In this regard, they resemble subway maps.[8] Since D initially lacks evidence regarding which path X chose, D's initial probability distribution treats all three maps on a par:

$$p(M_1) = p(M_2) = p(M_3) = 1/3.$$

The probabilities sum to 1, because M_1, M_2, and M_3 exhaust the space of possibilities.

From left to right, label the three branches 1, 2, 3. D can sniff each branch i, an observation with two possible outcomes: y_i, signifying that D detects some olfactory trace of X on branch i; and n_i, signifying that D detects no such olfactory trace. Since these are the only two possible options,

$$p(y_i) + p(n_i) = 1.$$

I assume that D assigns *conditional probabilities* $p(y_i \mid M_j)$: the probability of measurement y_i when D sniffs branch i, assuming that M_j is veridical. Since D's perceptual systems are fallible, $p(y_i \mid M_i)$ is less than 1: even if X chose branch i, D may not smell it. To be conservative, I assume the "prior likelihood"

$$p(y_i \mid M_i) = 2/3.$$

Since $p(y_i) + p(n_i) = 1$, it follows that the chance of false negatives is

$$p(n_i \mid M_i) = 1/3.$$

There is also a slight chance of false positives, presumably less than that of false negatives. For $i \neq j$, I assume that

$$p(y_i \mid M_j) = 1/6,$$

and hence that

$p(n_i \mid M_j) = 5/6.$

As I explain below, one could vary these numbers considerably without altering the thrust of my analysis.

I divide *D*'s activity into three stages. **Stage One**: *D* sniffs branch 1. **Stage Two**: *D* sniffs branch 2. **Stage Three**: *D* chooses branch 3 without sniffing. I now describe each stage in more detail.

Stage One

D sniffs branch 1, obtaining result n_1. How does this observation lead *D* to update *p*? Bayes's law, a fundamental result of probability theory, asserts that:

$$p(a \mid b) = \frac{p(b \mid a)\, p(a)}{p(b)}.$$

It is convenient to rewrite this formula as:

$$p(a \mid b) = \eta\, p(b \mid a)\, p(a),$$

where we regard η as a normalization constant ensuring that relevant probabilities sum to 1. In the cases that interest us, Bayes's law entails

$$p(M_i \mid n_1) = \eta\, p(n_1 \mid M_i)\, p(M_i).$$

Intuitively: the probability of a given hypothesis, conditional on our evidence, is proportional to the prior probability of the hypothesis times the prior likelihood of our evidence conditional on the hypothesis. Substituting our assumed values for relevant probabilities, it is easy to show that

$$p(M_1 \mid n_1) = 1/6$$

$$p(M_2 \mid n_1) = 5/12$$

$$p(M_3 \mid n_1) = 5/12.$$

Following standard Bayesian procedure, we assume that D "conditionalizes": as a result of observation n_1, D updates its probabilities so that the new probability assigned to M_i is $p(M_i \mid n_1)$. Thus, D redistributes probabilities over M_1, M_2, and M_3 to 1/6, 5/12, and 5/12, respectively. This is intuitively plausible. Since D did not detect any sign of X down path 1, D lowers the probability it assigns to M_1. No evidence yet differentiates between M_1 and M_2, so D assigns them equal probability.

Stage Two

D sniffs branch 2, obtaining result n_2. How does this observation lead D to update p? We employ a generalized form of Bayes's law:

$$p(a \mid b, c) = \beta\, p(b \mid a, c)\, p(a \mid c),$$

where $p(x \mid y, z)$ is the probability of x given that y and z obtain, and where β is a normalization constant ensuring that relevant probabilities sum to 1. Thus,

$$p(M_i \mid n_1, n_2) = \beta\, p(n_2 \mid M_i, n_1)\, p(M_i \mid n_1).$$

Following typical practice in probabilistic robotics (Thrun, et. al., 2005, p. 33), we deploy the *Markov assumption*: given the current state of the world, past observations are irrelevant to predictions about future observations. More formally,

$$p(n_2 \mid M_i, n_1) = p(n_2 \mid M_i).$$

Under this assumption, it is easy to show that

$$p(M_1 \mid n_1, n_2) = 2/9$$

$$p(M_2 \mid n_1, n_2) = 2/9$$

$$p(M_3 \mid n_1, n_2) = 5/9.$$

Assume that D conditionalizes once again. Then D assigns probabilities 2/9, 2/9, and 5/9 to M_1, M_2, and M_3, respectively. This is intuitively plausible. D's two observations render M_3 most probable, and they do not differentiate between M_1 and M_2.

Stage Three

So far, we have considered how D updates probabilities. We must now examine D's utilities. I assume that D has four available actions: remain at the fork of the road (whether performing further observations or merely abandoning the chase); or else traverse branches 1, 2, or 3, respectively. Call these actions H, x_1, x_2, and x_3. Action u and map M_i jointly determine a new map, depicting what the world would be like if it begins in the state depicted by M_i and then changes only in that D performs u. I denote this new map with "(M_i, u)." For instance, (M_1, H) is just M_1, while (M_1, x_2) is

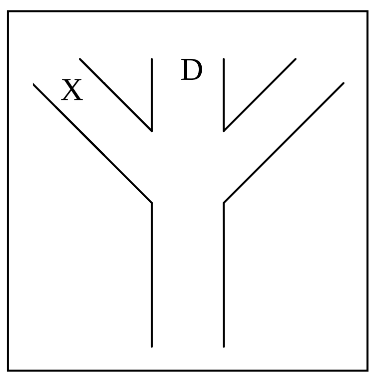

Note that, although I defined the meta-linguistic term "(M_i, u)" by using the material conditional, the cognitive map denoted by this meta-linguistic term does not itself have conditional structure, or logical form more generally. There are twelve maps (M_i, u), each associated with a *payoff* C(M_i, u). The following chart summarizes one possible set of payoffs:

	M_1	M_2	M_3
H	0	0	0
x_1	1	-1	-1
x_2	-1	1	-1
x_3	-1	-1	1

Columns correspond to possible initial states of the world. Rows correspond to *D*'s four possible actions. The entry in column M_i and row u is the payoff $C(M_i, u)$ if map (M_i, u) is veridical. Our chart reflects two assumptions: *D* prefers catching *X* to not catching *X*; and *D* prefers remaining immobile to wasting resources in failed pursuit of *X*.

The *expected payoff E(u)* of action *u* is a weighted average of the possible payoffs resulting from *u*:

$$E(u) = \sum_{i=1,2,3} p((M_i, u) \mid u) \times C(M_i, u),$$

where $p((M_i, u) \mid u)$ is the conditional probability that map (M_i, u) is veridical given that *D* performs *u*. I assume that *D*'s actions generate no additional uncertainty, so that $p((M_i, u) \mid u) = p(M_i)$. Thus,

$$E(u) = \sum_{i=1,2,3} p(M_i) \times C(M_i, u).$$

At Stage One,

$$E(H) = 0,\ E(x_1) = E(x_2) = E(x_3) = -1/3$$

At Stage Two,

$$E(H) = 0,\ E(x_1) = -2/3,\ E(x_2) = E(x_3) = -1/6$$

At Stage Three,

$$E(H) = 0,\ E(x_1) = E(x_2) = -5/9,\ E(x_3) = 1/9.$$

Under the standard assumption that *D* performs the action with highest expected payoff, our analysis predicts that *D* performs *H* at Stage One, *H* at Stage Two, and x_3 at Stage Three. This is the phenomenon we wished to explain.

§5. A general theoretical framework

Our Bayesian model of Chrysippus's dog satisfies the three criteria from §3. First, it explains the desired phenomena. Second, it supports appropriate counterfactuals. For instance, if *X* had chosen path 2 rather than path 3, then at Stage Two *D* would have recorded observation y_2. It is easy to show that *D* would then have immediately chosen path 2. Third, our model depicts *D*'s behavior as reflecting a more general capacity to perform Bayesian reasoning over cognitive maps.

That capacity extends far beyond the scenario described in §4. Probabilistic robotics studies how robots equipped with the capacity can navigate through diverse physical environments (Thrun, et al. 2005). It offers computational models describing how perceptual input induces a rational robot to update its probability distribution over a space of possible maps. Recently, robots along these lines performed impressively in the DARPA Grand Challenge Race, sponsored by the United States Department of Defense. Of course, current robotics algorithms are far more sophisticated than the tinker-toy model from §4. But the basic ingredients are the same: probability distributions over cognitive maps, Bayes's law, conditionalization, the Markov assumption, expected utility maximization, and so on. Thus, our simplistic model of Chrysippus's dog illustrates a flexible theoretical framework that has already enjoyed great practical success.

We could refine our model in various ways. For instance, we might consider maps that represent *metric* features of the environment, such as distances and angles. Or we might describe how D alters its heading as it sniffs different branches. Or we might describe D's actions not in terms of their environment-involving consequences (e.g. *D travels down the third path*) but in terms of more specific environment-independent motor commands. Or we might describe D as updating its probability distribution and its motor commands even after it begins moving down path 3. Or we might treat D's actions as introducing additional uncertainty, in which case we would abandon the assumption that $p((M_i, u) \mid u) = p(M_i)$. Contemporary robotics offers many ideas about how to effect these and other refinements. Although refinement might improve our analysis, it would not change the moral: we can build a robot that employs Bayesian reasoning over cognitive maps and that behaves as Chrysippus describes. The robot would also behave appropriately in counterfactual variants upon Chrysippus's scenario.

§6. Comparing the two explanatory strategies

How might proponents of strategy (1) react to my discussion so far? I consider three possible objections.

Objection: The model from §4 is an instance of strategy (1), not an alternative to it. Even if we grant that the model assigns probabilities and utilities to non-logical representations, the assignments employ identities, such as "$p(M_1) = 1/3$" and "$C(M_1, H) = 0$," that have logical structure. Moreover, the proposed model treats D as performing mathematical calculations, which requires logical inference from mathematical axioms. Thus, §4 implicitly posits deductive reasoning over logically structured mental states.

Reply: Identities such as "$p(M_3) = 1/3$" and "$C(M_1, H) = 0$" describe which probabilities and utilities D assigns to which maps. But D itself can assign probabilities and utilities without employing these identities. The "assignment" consists in suitable functional relations between cognitive maps and mental representations denoting numbers.[9] What is it for D to assign probability 1/3 to M_1? It is for D to enter into a mental state bearing appropriate functional relations to other mental states: the functional relations described in §4. An assignment of probabilities or utilities to cognitive maps is a mental state that occupies a suitable role in probabilistic calculation. Nor does probabilistic calculation presuppose logical structure or logical inference. The literature offers numerous models of mathematical computation, such as Turing machines and register machines, that do not employ logical inference. These computational models demonstrate that mathematical calculation need not involve deducing theorems from axioms.

Objection: The map-like character of cognitive maps plays no essential role in §4. The same probabilistic calculations would apply if p were defined over unstructured representations, rather than representations with internal structure akin to concrete maps. Since Bayesianism rather than cartographic structure does all the work, §4 illustrates nothing about the explanatory potential of cognitive maps.

Reply: A model defined over unstructured representations fails to explain a crucial phenomenon: *systematicity*. As Fodor (1987, pp. 147-153) emphasizes, a basic empirical fact about cognition is that creatures able to instantiate certain contentful mental states are necessarily able to instantiate certain other contentful mental states. For instance, it seems plausible that any terrestrial animal able to entertain the possibility represented by M_1 could also entertain the possibilities represented by M_2 or M_3. Fodor argues that a cognitive model based on unstructured mental representations cannot explain these systematic interrelations among possible mental states. He argues that a satisfying explanation should treat contentful mental states as relations to structured mental representations. We explain systematicity by noting that a representation's parts can be recombined to form new representations. The model from §4 implements this proposal, treating relevant mental states as relations to mental representations whose semantically relevant structure resembles that of ordinary concrete maps. As Braddon-Mitchell and Jackson (2007, p. 182) note, ordinary cartographic representation is systematic, so a system of cartographic mental representations would likewise be systematic. I conclude that §4 is far more satisfying than a model that replaces cognitive maps with unstructured representations.

Objection: §4 relies on *post hoc* assumptions about *D*'s mental states. It assumes that *D* has winnowed the space of possibilities to M_1, M_2, and M_3. It assumes that *D* has selected appropriate *prior probabilities* --- $p(M_1)$, $p(M_2)$, and $p(M_3)$ --- and *prior likelihoods* --- $p(y_i \mid M_j)$ and $p(n_i \mid M_j)$. It assumes that *D* associates various payoffs with various outcomes. Altering these assumptions would block the proposed derivation of *D*'s behavior or, even worse, generate an incorrect prediction. Since §4 retroactively tailors its assumptions to the desired outcome, it yields a thoroughly vacuous account.

Reply: The key question is how my account compares to strategy (1). Presumably, proponents of (1) envisage something like the following explanation:

(*) *D* tracks *X* to the fork in the road, so *D* believes that *X* chose one of the three paths. *D* observes no olfactory signs that *X* chose either path 1 or path 2, so it concludes that *X* chose neither path. By logical reasoning, *D* concludes that *X* chose path 3. Since *D* prefers catching *X* to not catching *X*, *D* chooses path 3.

(*) assumes that *D* believes itself to be facing a three-pronged fork in the road. It assumes that *D* believes *X* could have traveled down any of the three paths. It assumes that *D* can record relevant olfactory observations, which *D* then takes at face value as veridical. Finally, it assumes that *D* prefers catching *X* to not catching *X*. Altering any of these assumptions would block the proposed explanation of *D*'s behavior, or, even worse, yield an incorrect prediction. Our question is whether (*)'s assumptions are any less "post hoc" than those required by my treatment.

In many respects, the assumptions seem comparable. The main apparent difference is that my assumptions are more specific. I assign precise numerical probabilities and payoffs, and I exploit my numerical assumptions in an essential way. Let us consider the various assumptions in turn.

Prior probabilities: It is highly plausible that *D* initially assigns approximately equal probabilities to M_1, M_2, and M_3, since by stipulation *D* lacks any differentiating evidence. If desired, we could extend our model backwards in time, discussing how *D* generates an initial probability distribution over the space of possible maps. Mapping is a central research topic of probabilistic robotics, which offers a wealth of theories. Any such extended model should depict *D* as tending to assign roughly equal probabilities to M_1, M_2, and M_3 until it collects relevant differentiating evidence.

Prior likelihoods: My assumptions here are quite conservative. Specifically, the assumed high probability of false negatives --- $p(n_i \mid M_i) = 1/3$ --- shows that *D* can act appropriately even if it treats its own perceptual faculties as unreliable. Furthermore, we could vary my specific numerical assumptions considerably while generating the same predictions. Holding fixed the utilities and prior probabilities, and assuming that *D* treats the three paths symmetrically, we can easily show that any assignment of prior likelihoods yields the desired behavior if it satisfied these constraints:

$p(n_i \mid M_i) > 0$

$p(n_i \mid M_i) < p(n_i \mid M_j) / 2$,

where $i \neq j$. In other words, we achieve our desired result if *D* allows some probability of false negatives and if that probability is not too outrageously large.

Payoffs: Once again, we can vary the specific numbers considerably while generating the same predicted behavior. Let $a = C(M_i, u_i)$. If $i \neq j$, let $b = -C(M_i, u_j)$. Let $C(M_i, H) = 0$. Holding fixed the prior probabilities and likelihoods, we can easily show that any values of *a* and *b* yield the desired behavior if

$4b/5 < a < 7b/5$.

This inequality imposes a non-trivial constraint upon possible models of *D*. The constraint is a virtue, not a defect. It yields a quantitative analysis of how *D*'s payoffs and probabilities jointly determine action. For instance, if $a < 4b/5$, then the prospect of wasted resources outweighs the prospect of catching *X*, so *D* remains immobile even while realizing that *X* probably chose path 3. In contrast, (*) appeals vaguely to a "preference" for catching *X*. It offers no guidance in comparing that preference with *D*'s preference for conserving resources. Bayesian decision theory provides an appealing framework for conducting such comparisons, as illustrated by the above analysis.

Apparently, then, §4 is no less *a priori* plausible than (*). If so, then a satisfying treatment of Chrysippus's dog need not cite logical reasoning over logically structured mental states. We can instead cite Bayesian reasoning over cognitive maps.

§7. Animal cognition

Chrysippus's dog is scarcely more than a thought experiment. To what extent does my discussion bear on actual non-linguistic creatures? Do such creatures perform Bayesian reasoning over cognitive maps? Do they execute deductive inferences? These questions, which transgress the limits of *a priori* analysis, impel us to leave the armchair and engage with scientific psychology. Thought experiments are useful. They facilitate comparison of alternative theories in a relatively clean setting. But they must eventually yield to actual experiments featuring quantitative measures and proper controls.

In this spirit, let us first ask whether the approach sketched in §§4-5 illuminates navigation. As already noted, many psychologists enthusiastically posit cognitive maps. Bayesian models are also popular within psychology, as applied to perception, word learning, and many other phenomena. A few researchers combine these two strands into cognitive models that posit probabilistic reasoning over cognitive maps. For instance, Balakrishnan, Bousquet, and Honavar (1999) argue that such models help explain various experimental results. This approach seems likely to receive further empirical application as its success within robotics becomes better known among psychologists. Of course, the particular model from §4 is highly simplified. An empirically credible model would feature many refinements, including those mentioned in §5. But the basic strategy looks promising. For further discussion, see (Rescorla, in press a).

What about logical reasoning? Many psychologists claim that non-linguistic creatures perform deductive inferences. Building on work of Premack and Premack (1994), Call and Carpinter (2001) presented a chimpanzee with two hollow tubes. As the chimpanzee watched, experimenters placed food inside one tube. A screen obscured which tube they selected. They then allowed the chimpanzee to search inside the tubes for the food. In 20-30% of the trials, a chimpanzee who discovered that one tube was empty immediately selected the second tube without searching inside it. Call (2004) concludes that the chimpanzee employed disjunctive syllogistic reasoning to determine which cup contained food. Erdőhegyi, et al. (2007) report a similar but somewhat weaker result for dogs. In another series of experiments, Call (2004) presented apes with two opaque cups, only one of which contained food. He then shook the empty cup so that the ape could observe that no noise was produced. Three out of nine test subjects performed above chance in selecting the unshaken cup. Call (2004; 2006) argues that these subjects executed a disjunctive syllogism.

Each of the foregoing studies sought to disconfirm rival hypotheses through appropriate controls. For instance, Call (2004) employed controls to show that his apes were not just smelling the hidden food or detecting inadvertent cues about its location from experimenters. He performed additional experiments designed to show that the apes' behavior did not simply reflect a reinforcement history resulting in learned associations, such as an aversion to a noiseless shaken cup (Call 2006; Call 2007).

Nevertheless, many psychologists remain unconvinced. According to Penn and Povinelli (2007), proponents of non-linguistic logical inference mistakenly assume a rigid dichotomy between associationist and deductivist theories of mental activity: either learned associations exhaust a creature's mental activity, or else the creature performs logical inferences. This rigid dichotomy neglects the possibility of non-deductive mental processes vastly more sophisticated than associative learning. Penn and Povinelli caution that, even if we disconfirm an associationist explanation of some observed phenomena, we should not immediately embrace a deductivist conclusion.

My approach, articulated in §§4-5, occupies Penn and Povinelli's desired middle ground. It is neither deductivist nor associationist. Rather than positing logical inferences or learned associations, it posits rational probabilistic inferences over contentful mental representations. This approach can explain many behavioral phenomena supposedly indicative of non-linguistic syllogistic reasoning. For instance, the results from (Call and Carpinter 2001; Erdőhegyi, et al. 2007) recall Chrysippus's dog. We can readily explain those results through suitably altered versions of §4's simplistic model. We treat the animal as updating a probability distribution over possible maps of its surroundings.

The results reported in (Call 2004) are harder to accommodate, because they introduce a novel element: the relation between the cup's contents, the shaking of the cup, and the noise thereby produced. We might accommodate this novel element through prior likelihoods $p(n \mid M, s)$, where M is a map, s represents that a shaking of the cup occurs, and n represents that the ape's sensors detect noise. But this maneuver substantially alters the model from §4, in which prior likelihoods are conditional only upon cognitive maps, not upon spatiotemporal events (such as a shaking of a cup). An adequate theory of how the ape arrives at a suitable prior likelihood $p(n \mid M, s)$ will go substantially beyond the models of map-learning and navigation currently employed within probabilistic robotics.

What additional representational and inferential resources beyond those posited in §4 would an adequate theory require? According to Call, his results show that the ape represents and reasons about causal relations. Penn and Povinelli (2007, pp. 109-110) disagree. But suppose we concede that Call is correct. In particular, suppose we concede that the prior likelihood $p(n \mid M, s)$ reflects the ape's grasp of causal relations among physical events. We do not thereby concede that the ape engages in anything resembling *logical* reasoning. There is no obvious reason why a systematic theory of causal representation and reasoning must invoke logically structured mental states.

For instance, consider the theory of *causal Bayes nets* (Pearle 2000). A causal Bayes net is a directed acyclic graph. Each node is a variable whose values represent possible events or states of the world. A directed edge from one node to another represents direct causal influence of the former upon the latter. Every node is associated with a subjective probability distribution, conditional on the values of its parents. Causal influence and probabilistic dependence relate through the *causal Markov Condition*: a node is probabilistically independent of its nondescendants, conditional on its parents. Under the Markov assumption, conditional probability distributions for individual nodes determine a unique joint probability distribution defined over all the nodes.

This formalism yields an elegant framework for representing and reasoning about causal relations among events. Working within the framework, researchers have proposed various algorithms for inferring causal relations from observed correlations. Although the framework originated in statistics, computer science, and philosophy, several cognitive scientists have recently deployed it within empirical theories of causal reasoning in humans (Gopnik and Schulz 2007) and non-humans (Blaisdell, et al. 2006). To my knowledge, no one has yet integrated cognitive maps and causal Bayes nets into a synthesized theory of map-learning, navigation, and causal reasoning. But I see no obvious bar to an integrated theory. Since neither element of the proposed synthesis requires logically structured representations or deductive inferences, I see no reason why the proposed synthesis would require those resources. Clearly, the topic deserves further investigation.

Although my discussion of causal reasoning may seem intolerably vague, I have offered more detail than psychologists who attribute logical reasoning to non-linguistic animals. As Penn and Povinelli note, such psychologists never provide or even gesture towards formal models. For instance, Call offers only vague folk psychological talk about "causal-logical reasoning," without hinting how to convert such talk into actual psychological models. Thus, my proposal is no vaguer than Call's.

I submit that many experimental results supposedly indicative of non-linguistic deduction can be explained without citing logical reasoning or logically structured mental states. Admittedly, I have addressed only one strand in the relevant psychological literature. A more thorough discussion would survey the many other experimental results that purportedly reflect non-linguistic deduction.[10] Nevertheless, we can draw some preliminary conclusions. First, rational psychological processes defined over cartographic mental representations are possible. The relevant processes, grounded in Bayesian decision theory, differ markedly from deduction. Second, behavioral phenomena that superficially appear to involve logical reasoning may instead reflect non-deductive mental processes defined over non-logical mental representations. Hence, we must exercise caution when arguing through anecdotes, thought experiments, or scientific experiments that non-linguistic creatures' mental states have logical form. Whether non-linguistic cognition features logically structured mental states is an open question, to be settled through sustained engagement between philosophy and scientific psychology.

Acknowledgments

I am indebted to José Luis Bermudez, Robert Lurz, and Jason Newman for helpful comments.

Works Cited

Allen, C. (2006). Transitive inference in animals: reasoning or conditioned associations? In S. Hurley and M. Nudds (Eds.), *Rational Animals?*. Oxford: Oxford University Press.

Ambrose. (1961). *Hexameron*. In J. Savage (Trans.), *The Fathers of the Church*: *A New Translation*, vol. 42. New York: Fathers of the Church, Inc.

Balakrishnan, K., Bousquet, O., and Honavar, V. (1999). Spatial learning and localization in rodents: a computational model of the hippocampus and its implications for mobile robots. *Adaptive Behavior*, 7, 173-216.

Bermudez, J. L. (1998). *The Paradox of Self-Consciousness*. Cambridge, MA: MIT Press.

---. (2003). *Thinking without Words*. Oxford: Oxford University Press.

Blaisdell, A., Sawa, K., Leising, K., and Walmann, M. (2006). Causal reasoning in rats. *Science*, 311, 1020-1022.

Braddon-Mitchell, D. and Jackson, F. (2007). *Philosophy of Mind and Cognition*, 2nd ed. Cambridge, MA: Blackwell.

Brandom, Robert. *Making it Explicit*. Cambridge: Harvard University Press, 1994.

Burge, T. (2005). *Truth, Thought, and Reason,* Oxford: Oxford University Press.

Call, J. (2004). Inferences about the location of food in the great apes. *Journal of Comparative Psychology*, 118, 232-241.

---. (2006). Descartes's two errors: reason and reflection in the great apes. In S. Hurley and M. Nudds (Eds.), *Rational Animals?*. Oxford: Oxford University Press.

---. (2007). Apes know that hidden objects can affect the orientation of other objects. *Cognition*, 105, 1-25.

Call, J. and Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition*, 4, 207-220.

Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, 21, 145-182.

Carruthers, P. (2002). The cognitive functions of language. *Behavioral and Brain Sciences*, 25, 657-726.

---. (2004). On being simple-minded. *American Philosophical Quarterly*, 41, 205-220.

Casati, R. and Varzi, A. (1999). *Parts and Places*. Cambridge, MA: MIT Press.

Davidson, D. (1984). *Inquiries into Truth and Interpretation*. Oxford: Clarendon Press.

Dennett, D. (1996). *Kinds of Minds*. New York: Basic Books.

Devitt, M. (2006). *Ignorance of Language*. Oxford: Clarendon Press.

Dummett, M. (1993). *The Seas of Language*. Oxford: Clarendon Press.

Erdőhegyi, Á., Topál, J., Vrányi, Z., and Miklósi, Á. (2007). Dog-logic: inferential reasoning in a two-way choice task and its restricted use. *Animal Behavior*, 74, 725-737.

Floridi, L. (1997). Skepticism and animal rationality: the fortune of Chrysippus's dog in the history of Western thought. *Archiv für Geschichte der Philosophie*, 79, 27-57.

Fodor, J. (1981). *Representations*. Cambridge, MA: MIT Press.

---. (1987). *Psychosemantics*. Cambridge, MA: MIT Press.

---. (1991). Replies. In Barry Loewer (Ed.), *Meaning in Mind.* Cambridge, MA: Blackwell.

---. (2003). Review of Bermudez's *Thinking Without Words*. *London Review of Books*, 25 (19), pp. 16-17.

Fry, R. (2002). The engineering of cybernetic systems. *Bayesian Inference and Maximum Entropy Systems in Science and Engineering*: *American Institute of Physics Conference Proceedings*, 617, 497-528.

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

Gärdenfors, P. (2003). *How Homo Became Sapiens: On the Evolution of Thinking*. Oxford: Oxford University Press.

Glock, H.-J. (2000). Animals, thoughts, and concepts. *Synthese*, 123, 35-64.

Gopnik, A. and Schulz, L. (Eds.). (2007). *Causal Learning*. Oxford: Oxford University Press.

Horgan, T. and Tienson, J. (1996). *Connectionism and the Philosophy of Psychology*. Cambridge, MA: MIT Press.

Kitchin, R. (1994). Cognitive maps: what are they and why study them? *Journal of Environmental Psychology*, 14, 1-19.

Knill, D., and Richards, W. (Eds.). (1996). *Perception as Bayesian Inference*. Cambridge: Cambridge University Press.

Lurz, Robert. (2007). In defense of wordless thoughts about thoughts. *Mind and Language*, 22, 270-296.

McDowell, J. (1996). *Mind and World*, 2nd ed. Cambridge, MA: Harvard University Press.

Malcolm, N. (1972). Thoughtless brutes. *Proceedings and Addresses of the American Philosophical Association*, 46, 5-20.

Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press.

O'Keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford: Clarendon University Press.

Peacocke, C. 1997. Concepts without words. In R. Heck (Ed.), *Language, Thought, and Logic*. Oxford: Oxford University Press,.

Pearle, J. (2000). *Causality*. Cambridge: Cambridge University Press.

Penn, D. and Povinelli, D. (2007). Causal cognition in humans and non-human animals: a comparative, critical review. *Annual Review of Psychology*, 58, 97-118.

Penn, D., Holyoak, K., and Povinelli, D. (In press). Darwin's mistake: explaining the discontinuity between human and non-human minds. *Behavioral and Brain Sciences*.

Plutarch. (1957). *Moralia: De Solertia Animalium*, vol. xii. H. Cherniss and W.

Helmbold (Trans.). Cambridge: Harvard University Press.

Pratt, Ian. (1993). Map semantics. In A. Frank and I. Campari (Eds.), *Spatial Information Theory: A Theoretical Basis for GIS*. Berlin: Springer-Verlag.

Premack, D. and Premack, A. J. (1994). Levels of causal understanding in chimpanzees and children. *Cognition*, 50, 347-362.

Putnam, H. (1992). *Renewing Philosophy*. Cambridge, MA: Harvard University Press.

Pylyshyn, Z. (1986). *Computation and Cognition.* Cambridge, MA: MIT Press.

---. (2003). *Seeing and Visualizing*. Cambridge, MA: MIT Press.

Rescorla, M. (In press a). Cognitive maps and the language of thought. *British Journal for the Philosophy of Science*.

---. (In press b). Predication and cartographic representation. *Synthese*.

Rey, G. (1995). A not ‘merely empirical’ argument for a language of thought. *Philosophical Perspectives*, 9, 201-222.

Sloman, A. (1978). *The Computer Revolution in Philosophy*. Hassocks: The Harvester Press.

Sorabji, R. (1993). *Animal Minds and Human Morals*. Cornell: Cornell University Press.

Stalnaker, R. (1999). *Context and Content*. Oxford: Oxford University Press.

Thrun, S., Burgard, W., and Fox, D. (2005). *Probabilistic Robotics*. Cambridge, MA: MIT Press.

Tolman, E. (1948). Cognitive maps in rats and men. *Psychological Review*, 55, 189-208.

Wilson, M. (1995). Animal ideas. *Proceedings and Addresses of the American Philosophical Association*, 69, 7-25.

Notes

[1] Wilson (1995) offers an historical overview.
[2] From *Outlines of Pyrrhonism*, I.69, as translated in (Floridi, 1997).
[3] Camp (2007) considers how one might enrich maps with semantic devices akin to negation, disjunction, the conditional, and the quantifiers. However, she seems to grant that these semantic devices go beyond ordinary cartographic representational mechanisms.
[4] Floridi (1997) offers a useful overview of the extensive historical literature.
[5] Bermudez (2003, pp. 140-149) denies that non-linguistic creatures execute logical reasoning, but he ascribes to them a more primitive mode of reasoning ("proto-logic") not involving standard logical connectives. Specifically, he isolates a proto-logic analogue to the inference "*p* or *q*; not-*p*; therefore *q*." He does not apply his discussion of proto-logic to the case of Chrysippus's dog. However, one might try to elaborate (4) by citing Bermudez's proto-logical analogue to the disjunctive syllogism. My alternative approach should be seen as complementary to Bermudez's. For criticism of Bermudez, see (Lurz, 2007).
[6] I am indebted to José Luis Bermudez for pressing this question.
[7] More precisely, the standard Kolmogorov axiomatization of probability defines a probability space as (Ω, A, p), where Ω is a non-empty set, A is a σ-algebra over Ω (i.e. a set of subsets of Ω that contains Ω and is closed under countable union and complementation in Ω), and p is a probability measure. It is consistent with this axiomatization to construe Ω as containing cognitive maps.
[8] Some writers restrict the phrase "cognitive map" to representations that represent metric structure. However, this usage is hardly universal. For instance, Gallistel's official definition, quoted in §2, mentions "geometric relations" without privileging metric over topological properties.
[9] For empirical evidence that even fairly primitive animals such as rats represent and perform computations involving numbers, including non-integral rationals, see (Gallistel, 1990, pp. 317-383). Note also that much of the scientific literature on perception treats low-level visual processes as performing sophisticated Bayesian computations (Knill and Richards, 1996). Thus, the fact that my Bayesian-cum-cartographic model attributes numerical computation to Chrysippus's dog should not seem at all problematic.
[10] For critical discussion of relevant psychological literature, see (Allen, 2006), (Bermudez, 2003, pp. 112-114), and (Penn, Holyoak, and Povinelli, forthcoming).