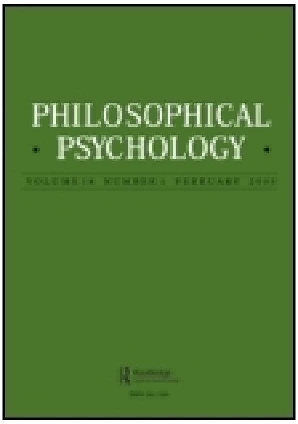

Philosophical Psychology
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/cphp20

Why we can't say what animals think

Jacob Beck

Published online: 29 Mar 2012.

Philosophical Psychology, 2013
Vol. 26, No. 4, 520–546, http://dx.doi.org/10.1080/09515089.2012.670922

Routledge
Taylor & Francis Group

Why we can't say what animals think

Jacob Beck

Realists about animal cognitive representations confront a puzzle. If animals have real, contentful cognitive representations, why can't anyone say precisely what the contents of those representations are? I argue that realists can best resolve this puzzle by appealing to differences in the format of animal cognition and human language.

Keywords: Animal Cognition; Animal Thought; Content; Format; Indeterminacy; Representation

> "If a lion could talk, we could not understand him."
> Ludwig Wittgenstein (1973) *Philosophical Investigations*, p. IIxi

1. The Problem

The past few decades have witnessed an explosion of research in cognitive ethology and animal psychology, much of which explicitly appeals to cognitive representations to explain the sophisticated behavior of non-human animals. As a result, the skepticism about such representations that guided much philosophy and psychology in the 20th century has largely receded. Most animal researchers now accept that animal cognition involves operations over causally efficacious representations with intentional content—representations that characterize the world as being a certain way. More colloquially, it is now widely accepted that animals think.

But our attributions of representations to animals face an embarrassing difficulty: when we try to put words to those representations, our articulations always seem to mischaracterize their contents. Even the staunchest realists about animal cognition acknowledge this difficulty. For example, in a paper arguing that animals as simple as *bees* think, Tetzlaff and Rey write: "of course, a problem of enormous significance and difficulty for psychology in general is how to characterize *precisely* the intentional content of [animals'] representations .... We do not pretend to have any adequate answer to this question" (2009, pp. 74–75). The current state of theorizing about

Jacob Beck is an Assistant Professor in the Department of Philosophy at Texas Tech University.
Correspondence to: Jacob Beck, Box 43092, Texas Tech University, Lubbock, TX 79409, USA. Email: jacob.beck@ttu.edu

animal cognition is thus awkwardly committed to both of the following theses (see also Jamieson, 2009, p. 17):

> *Realism*: animals have causally efficacious cognitive representations with determinate contents.
>
> *Indeterminacy*: we are currently unable to provide precise linguistic characterizations of the contents of animals' cognitive representations.

This dual commitment cries out for explanation. If animals really have cognitive representations with determinate contents, why can't anybody say what those contents are?

Some philosophers have responded to this tension by rejecting Realism. (I will capitalize the initial letters in 'Realism' and 'Indeterminacy' whenever I intend to refer to the specific theses defined above.) For example, Davidson (1975/1984) takes our inability to precisely specify the contents of animals' cognitive representations as evidence that they don't have any. Similarly, Dennett (1987, p. 108) and Jamieson (2009) argue that Indeterminacy is best explained by *interpretivism*, according to which the contents of animals' cognitive representations are necessarily co-extensive with the contents that would be attributed by an ideal interpreter (roughly: an interpreter that knows all the behavioral facts and is rational and sober).[1] Since interpretivists expect ideal interpreters to reach divergent conclusions about the contents of any animal's cognitive representations, they conclude that there is simply no fact of the matter about which of a range of contents an animal "really" has. Yet while these views certainly help to explain Indeterminacy, they are difficult to reconcile with the many flourishing research programs in cognitive ethology and animal psychology whose explanations of animal behavior are structured around an appeal to real cognitive representations with determinate contents. Those impressed by these research programs are thus likely to maintain that animals have cognitive representations whose contents are not necessarily co-extensive with the judgments of an interpreter. But how can this realist orientation towards animal cognition be reconciled with our inability to precisely characterize animals' cognitive representations? How can the Realist explain away Indeterminacy?

This paper has two aims. The first is to show that Realists are not without resources to accommodate Indeterminacy by distinguishing four explanations of Indeterminacy that are open to Realists. These explanations differ not only in where they place the principal blame for Indeterminacy, but also in how deeply entrenched they take Indeterminacy to be. According to the *explanation from insufficient time*, there is no principled obstacle to the linguistic expression of animals' cognitive contents. Animal researchers are on the right track, and are making considerable progress. We just need to give them more time to carry out further observations, run additional experiments, and continue to hone the methods that they have already begun to develop. According to the *explanations from muteness and unfamiliar contents*, by contrast, Indeterminacy runs deeper. It is caused by fundamental limitations in our epistemic position that our current methods are ill suited to overcome. Because animals don't speak and thus cannot tell us what they are thinking (the explanation from muteness), or because the contents of their cognitive representations are so different from our own that we lack the capacity to grasp or express them (the explanation from unfamiliar contents), we don't have any real idea how to access the contents of their cognitive representations, real and determinate though they are. Finally, according to the *explanation from nonlinguistic format*, the roots of Indeterminacy run deeper still. The problem is not one of impoverished methods or epistemic limitations. Rather, the format of animal cognition makes its contents impossible to translate into natural language. Even God couldn't tell us precisely what a lion is thinking.

While my first aim is simply to delineate these four explanations of Indeterminacy, my second aim involves playing favorites. I will argue that the explanation from nonlinguistic format is most likely to account for the bulk of Indeterminacy. I say "most likely" because, for reasons that will emerge, I doubt that we can say for certain at present what the primary source of Indeterminacy is. Nevertheless, I will argue that there is much more to be said in favor of the explanation from nonlinguistic format—and against the other three explanations—than many have supposed. As a result, even the Realist has reason to think that Indeterminacy is an inevitable feature of animal cognition.

Three brief prefatory comments are in order. First, I will assume not only that animals have real cognitive representations with determinate contents, but also that those representations are compositional, such that the contents of complex representations are determined by the contents of the primitive representations from which they are composed. These assumptions are, perhaps, equivalent to a weak form of the language of thought hypothesis (Fodor, 1975). But the language of thought hypothesis is also sometimes given a stronger interpretation, according to which the representations it posits involve compositional mechanisms familiar from natural language, such as predication and quantification, and I make no assumptions about animal representations being governed by such mechanisms. In fact, I will question such assumptions in section 6.

Second, I will assume that animals are capable of taking a variety of functionally individuated attitudes, including belief and desire, towards complexes of their cognitive representations. As a result, I will sometimes discuss Indeterminacy as a problem that concerns not just our ability to precisely characterize the contents of animals' cognitive representations, but also our ability to precisely characterize the contents of animals' beliefs and other attitudes.

Finally, I want to guard against a potential misinterpretation. One might be tempted to assimilate the problem of reconciling Realism and Indeterminacy to the familiar worry that we are ignorant of the peculiar phenomenal states of animals. Just as we cannot know "what it's like" to experience the sonar echolocation of a bat (Nagel, 1974), nor can we know what animals are thinking. But as we will see, the case for Indeterminacy does not derive from the assumption that animals are conscious. Whether or not there is something it's like to be a bird, bat, or bee, the problem of reconciling Indeterminacy and Realism thus remains.

2. Indeterminacy

To linguistically characterize a cognitive representation is to find a linguistic expression—a word, phrase, sentence, or set of sentences—that expresses its content. The thesis of Indeterminacy thus claims that we are currently unable to find linguistic expressions that precisely express the contents of animals' cognitive representations. In this section I clarify and motivate this thesis. In sections 3–6, I then consider the explanations from insufficient time, muteness, unfamiliar contents, and nonlinguistic format, respectively.

I will assume that sentences have determinate contents, and that those contents are *at least* as fine grained as the sets of possible worlds circumscribed by their truth conditions. For example, I will assume that the sentence 'The 44th president of the United States is in the White House' expresses a determinate content that differs from the content expressed by the sentence 'Barack Obama is in the White House', since there are possible worlds where the truth values of these sentences come apart. But I will not take a stand on whether sentential contents should be individuated by Fregean senses, Russellian propositions, sets of possible worlds, or something else besides. I will further suppose that natural language sentences have a compositional semantics, and thus that the contents of words and phrases are sufficiently fine grained to explain differences in the contents of the sentences to which they contribute. For example, I will assume that the content of the phrase 'The 44th president of the United States' is distinct from the content of 'Barack Obama'.

It is common to distinguish two types of content ascription. On a de re ascription, co-extensive terms can always be freely substituted without changing the accuracy of the ascription. By contrast, on a de dicto ascription the substitution of co-extensive terms can change the accuracy of the ascription. For example, suppose that Tim believes that Barack Obama is in the White House, but isn't aware that Obama is the 44th president. The ascription "Tim believes that the 44th president of the United States is in the White House" would then be false on a de dicto construal, but true on a de re construal.

Let us say that a linguistic expression, E, *precisely* expresses the content of a cognitive representation, R, if and only if: (i) E is used to render a de dicto ascription of content to R; and (ii) the content of E is the same as the content of R. Indeterminacy is thus the claim that we are currently unable to find linguistic expressions that, when used to render de dicto ascriptions of cognitive representations to animals, have the same contents as those representations.

I now turn to some considerations that are intended to illustrate the plausibility of Indeterminacy. I want to emphasize, however, that I do not take these considerations to prove that Indeterminacy is true. My aim in this paper is to show how Indeterminacy and Realism might be reconciled, not to establish the veracity of either thesis on its own.

To borrow an example from Stich (1983), suppose we watch a dog, Fido, chase a squirrel up an oak tree, and then stand at the base of the tree barking.[2] We might be tempted to assert, "Fido believes that a squirrel ran up the oak tree." But as Stich argues, a skeptic might reasonably question whether this ascription gets things quite right.

> "Does Fido really believe it is a squirrel up in the oak tree? Are there not indefinitely many logically possible creatures which are not squirrels but which Fido would treat indistinguishably from the way he treats real squirrels? Indeed, does he believe, or even care, that the thing up the tree is an *animal*? Would it not be quite the same to Fido if he had been chasing some bit of squirrel-shaped and squirrel-smelling machinery, like the mechanical rabbits used at dog-racing tracks? The concept of animal is tied to the distinction between living and nonliving, as well as to the distinction between animals and plants. But Fido has little grasp of these distinctions. How can you say that he believes it is a squirrel if he doesn't even know that squirrels are animals?" Confronted with this challenge . . . it no longer sounds quite right to say that Fido believes there is a *squirrel* up the oak tree. (Stich, 1983, pp. 104–105)

Stich's worry is compelling. It is doubtful that the sentence 'A squirrel ran up the oak tree' precisely expresses the content of Fido's belief. Moreover, the worry does not seem to rely on any idiosyncratic features of this particular example. Other ascriptions of cognitive contents to animals raise analogous worries. For example, does Fido really believe that the object he buried in the backyard is a *bone*? Would he not need to display a better understanding of anatomy, and a greater sensitivity to the differences between real bones and ersatz bones (Stich, 1979, p. 18)? Similarly, does he really believe that his bowl contains *meat*? Would he not need to be able to distinguish meat from synthetic protein, or at least know that meat comes from an *animal* (Dennett, 1969, p. 84; Putnam, 1992, pp. 28–31)? Could he even believe that the stuff in his bowl is *food* when he shows no interest in many types of food and has no conception of nutritional value (Dennett, 1969, p. 85)?[3]

While these examples support Indeterminacy, it is important to be clear about why. It is not because they rely on the implausibly stringent assumption that representing Xs requires the capacity to distinguish Xs from all conceivable non-Xs—for example, that representing squirrels requires the ability to distinguish squirrels from all conceivable non-squirrels. After all, a really good replica of a squirrel created by an evil demon might fool even the most knowledgeable zoologist, yet surely *humans* can represent squirrels. Nor is it because of a highly controversial (Fodor & Lepore, 1992) commitment to conceptual holism, according to which having any one belief (e.g., that a squirrel is in the tree) presupposes having a host of others (e.g., that squirrels are animals, that squirrels are alive, etc.). It seems possible that a person could believe that a squirrel is in the tree even if he were under the mistaken impression that squirrels are robots controlled from Mars (Putnam, 1962).[4] Rather, the crux of the worry raised by these examples is that the behavior of animals provides little evidence that they have representations with the same contents as our words. For example, Fido's squirrel-directed behavior (including Fido's capacity to distinguish squirrels from non-squirrels) is so impoverished that there is little reason to think that he has a representation with the very same content as our word 'squirrel'. This is not to deny that Fido represents some property or other when he sees a squirrel. But given the nature of Fido's behavior towards squirrels, it *is* doubtful that he represents the property of being a squirrel.[5]

Of course, one can respond to this and other similar examples by attempting to adjust the ascription (Allen, 1992). For example, perhaps what Fido really believes is that an object with properties $P_1 \ldots P_n$ ascended an object with properties $Q_1 \ldots Q_n$. However, this strategy faces a considerable hurdle. Unless we spell out exactly what the predicates $P_1 \ldots P_n$ and $Q_1 \ldots Q_n$ are, we won't have succeeded in precisely expressing the content of the dog's belief, and at least as things presently stand, no one seems to know how to do this, even for the simple canine belief we have been considering. Compounding this problem is the fact that precise expression requires an accurate de dicto ascription. Thus, it would not suffice if the predicates $P_1 \ldots P_n$ and $Q_1 \ldots Q_n$ were extensionally equivalent to the dog's representations (though even the difficulty of this task should not be underestimated). They would need to represent the relevant properties in exactly the same way as the dog does. But it's hard to imagine how we could find such predicates. Suppose, for example, that we became convinced that the dog represents the squirrel by means of its shape. We might then say that the dog believes that a squirrel-shaped object ascended the tree. However, the representation SQUIRREL-SHAPED is partially composed from the representation SQUIRREL, which the dog, we assumed, does not possess. We would thus need to find some other predicate to characterize the shape that the dog represents. But again, it's hard to imagine what predicate we might use.

Even if one is prepared to grant that *many* of our attributions of cognitive contents to animals are not precise, one might nevertheless insist that they are *sometimes* precise. There are two types of attributions, in particular, that seem like good prima facie candidates for precision. The first consists of basic sensory attributions, such as "ouch!" (attributed to an animal in pain) or "red!" (attributed to an animal with color vision that is similar to our own). The second consists of purely demonstrative attributions, such as "that is there," or partially demonstrative and partially sensory attributions, such as "that is red" or "pain is there." A full discussion of whether these attributions really are precise, however, would take us far afield, requiring, among other things, forays into the nature of animal consciousness and the semantics of demonstratives. To avoid such detours, I will stipulate that the thesis of Indeterminacy should be understood as applying only to the attribution of non-sensory and non-demonstrative cognitive representations. Our question is thus how to reconcile Realism with our inability to precisely characterize the large array of non-sensory and non-demonstrative cognitive representations that animals have.

Notice that I have not (yet) argued that our inability to precisely express the contents of animals' cognitive representations will extend indefinitely into the future, nor even that this inability is based on any principled obstacles. I have merely provided some reasons to accept Indeterminacy, which says that we are *currently* unable to precisely express the contents of animals' cognitive representations.

Granting that Indeterminacy characterizes our ascriptions of cognitive contents to animals, one might worry that the same can be said of our ascriptions of cognitive contents to our fellow humans, and thus that there is no special problem for animals. I submit, however, that there are many *best cases* of (non-sensory and non-demonstrative) de dicto ascriptions to humans in which the contents of the expressions we use do match the contents of the cognitive representations they are intended to characterize. For example, if you and I both watch Fido chase a squirrel up the tree, and I then say, "you believe that a squirrel ran up the tree," I take it that I will have succeeded in precisely characterizing a belief of yours. Similarly, if you and I are both standing on the front porch watching the snow fall in the light of day, and then I say, "you believe that it's snowing," I take it that the content of the sentence 'It's snowing' will have the same content as one of your beliefs.[6]

In claiming that there exist many best cases of content ascriptions to humans, I do not mean to deny that there will be many other non-best cases, or that our judgments of when a case is a best case are highly fallible. As anyone who has been in a relationship can attest, getting the contents of others' mental representations precisely right is a challenge at least as often as it is straightforward. In fact, it will be a feature of the view I eventually defend that we are sometimes incapable of precisely expressing even humans' cognitive contents. But as I understand Indeterminacy, it claims that there are close to zero best cases of (non-sensory and non-demonstrative) de dicto ascriptions to animals. Thus, while the contents we attribute to our fellow humans are often precise, if Indeterminacy is correct the contents we attribute to animals are almost never precise.

Some philosophers may still balk at my insistence that there exist many best cases in which the linguistic expressions we use to express the contents of our fellow humans' cognitive representations have exactly the *same* content as those representations themselves. It thus bears notice that the present issue could be recast in terms of degrees of content similarity. When we attribute cognitive representations to humans, we can often find expressions that are *highly similar* in content to the representations they are intended to characterize. But when we attribute cognitive representations to animals, we can almost never find expressions that have such a high degree of similarity in content to the representations they are intended to characterize. For example, even if the sentence, 'A squirrel ran up the oak tree', does not *exactly* match the content of your belief, it surely better approximates the content of your belief than the content of Fido's. Thus, the differences between human and animal ascriptions could be couched in the graded terms of content similarity rather than in the binary terms of content identity and difference. Although I will not pursue this alternative framing of the issue here, I want to emphasize that it is available, and that little would be lost from the discussion to follow if it were adopted.

Henceforth I will take for granted that Indeterminacy exists and is a special problem for animal cognition. The question I now wish to consider is how Realists can account for it. If animals really have cognitive representations with determinate contents, why are we currently unable to say what contents those representations have?

3. The Explanation from Insufficient Time

Perhaps the most straightforward Realist explanation of Indeterminacy—and certainly the most optimistic—is that it is merely the product of a young science. The methods of animal psychology and cognitive ethology have only broken free from the shackles of behaviorism over the past few decades, and so the disciplines have not had sufficient time to make the kind of progress required to overcome Indeterminacy. Nevertheless, researchers are on the right track. The methods they apply are sound, and they are making genuine progress. We need only to continue, with patience, in the same spirit.

Applied to Stich's Fido, this optimistic assessment would maintain that we could find a sentence that precisely expresses the content of Fido's belief about the squirrel if we only performed the appropriate experiments on Fido and thought hard about how to describe his cognitive representations, perhaps while consulting a suitably expansive thesaurus. As Allen (1992) observes, Stich's argument does not rule this possibility out. We cannot conclude from the fact that we are currently unable to find a sentence that precisely specifies the content of Fido's belief that we will forever be unable to find such a sentence. More research may point us in the right direction.

While Allen is surely right that the sorts of casual observations to which Stich appeals do not refute this possibility, the past forty years of research into animal cognition tells against it. Even when a wide variety of experiments are carried out, and researchers devote themselves to the task of specifying the contents of animals' mental representations, precise characterizations are still elusive. Thus, as I read the empirical literature on animal cognition, while we have been learning a lot about animal cognition, no one is making serious progress in finding sentences that *precisely* express the contents of animals' cognitive representations.

The only way to really convince oneself that animal researchers are not converging on precise characterizations of animal cognitive representations is to dig into the primary literature and consider a wide range of examples. Unfortunately, such a review would take more pages than I have here, and likely test the reader's patience. So rather than undertaking an exhaustive review, I will briefly discuss two examples. Readers who are independently familiar with the literature on animal cognition will recognize these examples as representative.

First, consider the question whether chimpanzees have a "theory of mind"—i.e., whether they represent the beliefs, desires, and other mental states of agents. On the one hand, there have been a number of experiments in which chimps manifest impressive forms of behavior. For example, Premack and Woodruff (1978) found that they can choose the appropriate photograph to complete an action sequence, suggesting that they have some understanding of what an actor is trying to do—what the actor's goal or intention is. More recently, Tomasello and colleagues (Tomasello, Call, & Hare, 2003) found that chimps can determine what their conspecifics have and have not seen. A subordinate will only approach a piece of hidden food if it knows that the dominant chimp didn't see it hidden, suggesting that the subordinate appreciates the dominant's perspective. On the other hand, however, there are contexts in which chimps are remarkably oblivious to others' perspectives. Povinelli and Eddy (1996) found that chimps will beg for food from a person with a bucket over his head as much as from a person who is gazing straight at them. But they also found that chimps will beg less from a person who is facing away. Finally, Call and Tomasello (1999) found that chimps have a particularly difficult time discerning an agent's *false* beliefs, once again calling into question the claim that chimps have a theory of mind. It's no easy task trying to say what all these results come to. Do chimpanzees represent the mental states of others? If so, which ones? Primatologists are themselves deeply divided over the answers to these questions.[7] Not that the results of these studies don't tell us anything. On the contrary, they clearly tell us a lot about how chimps interpret the behavior of others. But what they don't seem to do is allow us to *precisely* characterize how chimps represent other minds. Of course, future experiments might succeed where past experiments have failed. They *might*. However, there is little evidence in the scholarly record over the past thirty years to suggest that our characterizations are trending towards precision.

As a second example, consider animals' widely studied capacity to track numerical information (Dehaene, 1997; Gallistel, 1990). Rats can be trained to press a lever eight times in return for food (Mechner, 1958). Pigeons can tell whether they've pecked a key 50 times or 40 times (Rilling & McDiarmid, 1965). Monkeys can spontaneously discriminate a sequence of six tones from a sequence of ten tones (Hauser, Tsao, Garcia, & Spelke, 2003). At some level, all of these creatures appear to exhibit the capacity to represent numerical information. But this capacity admits of considerable variability. A rat that is trained to press a lever six times in exchange for food will press exactly six times perhaps 25 percent of the time, five or seven times slightly less often, and so on. Moreover, the variability is scalar; it increases linearly with the mean. Thus, a rat trained to press a lever ten times might press it exactly ten times only twenty percent of the time, with the remaining responses spread even more widely in a bell-shaped curve around ten. As a result of this scalar variability, the ability of animals to discriminate two numerical values is ratio sensitive. As the ratio of two numbers approaches one, the ability to discriminate them deteriorates. Thus, it is harder to discriminate six from seven than four from five even though the difference in the two cases is just one.

Of all human concepts, our mathematical concepts admit of some of the clearest definitions. Thus, if there were any cognitive representations of animals that we should be able to precisely characterize, one would predict that these numerical representations would be among them. But matters are not so simple. A natural first suggestion is that animals represent the integers—whole numbers such as one, two, and seven—but as Carey (2009, pp. 294–295) argues, there are two reasons to be skeptical of this suggestion. First, it is hard to reconcile with the variability intrinsic in animals' number-oriented behavior. If animals are only able to respond to a given integer one quarter of the time (and less often as the integer increases), it seems awkward to say that they nevertheless represent exact integers. Second, the very notion of an integer is defined in terms of the successor relation. Six is the successor of five, which is the successor of four, and so on. But the ratio sensitivity in animals' numerical discriminations seems to indicate that they represent the difference between four and five as *greater* than the difference between five and six.

Instead of taking these numerical representations to stand for the integers, Gallistel and Gelman (2000) have argued that they represent the real numbers. In defense of this idea, they argue that the neural representations underlying animals' numerical capacities are likely to be dense and continuous, just like the real numbers themselves. But as Burge (2010, p. 481) observes, the integers form a proper subset of the reals. So if animals can't represent the integers, they can't represent the reals either.

Carey (2009, p. 127) instead proposes that animals represent *approximate* cardinal values—for example, that the rat believes that pressing the lever approximately seven times will yield food. But attributing representations such as APPROXIMATELY SEVEN to a rat fails to get the content of the rat's mental representations quite right, both because it fails to specify *how* approximate the representation is, and because it implies that the rat's representation is structured from the distinct representations SEVEN and APPROXIMATELY, which of course it can't be if it lacks the capacity to represent the integers. Perhaps there is a better way to characterize these numerical representations, but if so, it's not easy to envision what it would look like.[8] Of course, this doesn't show that some alternative characterization *won't* ultimately succeed, but nothing in the trajectory of current research on animal numerical cognition suggests that it is honing in on a precise characterization of animals' numerical representations.

As I read the empirical literature on animal cognition, the examples of theory of mind and numerical representation are typical. Even where researchers have devoted themselves to learning about how animals represent a particular aspect of the world, they do not seem to be converging on a precise linguistic characterization of animals' representations. To be clear, I do not claim to be able to prove that this trend will continue indefinitely. It remains possible that if animal researchers keep at it, fifty or a hundred years from now they will be able to precisely specify the contents of animals' cognitive representations. But I wouldn't bet on it. Certainly such an optimistic prediction can hardly be justified based on current trends. It is thus worth considering whether there might be a deeper explanation of Indeterminacy.

4. The Explanation from Muteness

Language is clearly a useful tool when it comes to ascertaining the mental states of our fellow humans. If I want to know whether my neighbor believes that it will rain tomorrow, all I have to do is ask her and she'll tell me. Granted, this process doesn't always work. One of my neighbors, an elderly woman from Portugal, returns all of my queries with an uncomprehending smile. However, there are fairly simple steps we can take to communicate about her mental states. I can learn Portuguese; she can learn English; or (as usually happens) we can employ the woman's bilingual daughter as our interpreter. But now suppose that I want to know whether my neighbor's *cat* believes that it will rain tomorrow. Try though I may, I cannot learn to speak Catese;

nor can the cat be induced to learn English; and there are no translators who can help us out.

These considerations lead to a natural suggestion. Perhaps it is the fact that animals cannot talk that prevents us from precisely characterizing their cognitive representations. According to this suggestion, and *pace* Wittgenstein, if God gave a lion the gift of the gab, we could know precisely what it was thinking. But without that gift, our evidence is simply too impoverished to allow us to pin down animals' cognitive contents.

There is no denying that the inability of animals to speak a language is inconvenient from the perspective of researchers who are trying to capture their cognitive contents. The contents would be so much easier to specify if we could only engage animals in dialogue. But to what extent is linguistic incompetence a more severe limitation? Even if it requires more ingenuity, can't we gain access to the cognitive contents of animals by studying their non-verbal behavior instead? After all, one way to find out whether you believe that it will rain today is to ask you. However, another is to present you with an umbrella as you leave the house and see whether you take it. Similarly, one way to find out whether you prefer apples to oranges is to ask you. Another is to present you with one of each and see which one you choose. Of course, all of these methods are fallible. You might refuse the umbrella not because you believe it won't rain, but because it clashes with your outfit, or because you're trying to trick me, or because you're not fussed about staying dry. But linguistic evidence is fallible too. You might say that you believe it won't rain not because that's what you really believe but because you think that's what I want to hear, or you're trying to trick me, or you're self-deceived. The fallibility of non-verbal evidence is thus not a decisive strike against it. So why can't we make use of non-verbal evidence to precisely characterize the cognitive contents of animals?

Davidson raises the following worry about such an approach:

> One can believe that Scott is not the author of Waverly while not doubting that Scott is Scott; one can want to be the discoverer of a creature with a heart without wanting to be the discoverer of a creature with a kidney. One can intend to bite into the apple in the hand without intending to bite into the only apple with a worm in it; and so forth. The intensionality we make so much of in the attribution of thoughts is very hard to make much of when speech is not present. The dog, we say, knows that its master is at home. But does it know that Mr. Smith (who is his master), or that the president of the bank (who is that same master) is home? We have no real idea how to settle or make sense of these questions. (Davidson, 1975/1984, p. 163)

The thrust of Davidson's contention is that without speech our evidence is too impoverished to decide between competing content attributions that differ only intensionally. While a language user might greet the man entering the door by saying, "welcome Mr. Smith," and thus indicate that he takes the man to be Mr. Smith, the dog will only wag its tail and bark, which provides no indication of the mode of presentation under which the dog conceives of the man entering the door. Davidson thus concludes that when it comes to expressing the cognitive contents of animals, we will have no grounds for choosing between any two attributions that differ only in co-extensional terms. Our attributions to animals will consequently be imprecise. We'll never have enough evidence to select between competing de dicto content attributions that are de re equivalent.

Davidson himself takes these considerations to support the view that animals don't have cognitive representations at all. But we can imagine a Realist who stops short of Davidson's eliminativist conclusion, yet agrees with Davidson that the muteness of animals prevents us from precisely characterizing their cognitive contents. Such a Realist might thus appeal to Davidson's argument to provide an epistemic explanation of Indeterminacy.

Is Davidson right that the muteness of animals prevents us from choosing among de dicto ascriptions that are de re equivalent? Is there really no evidence that could help one decide whether the dog believes that its master is home, that Mr. Smith is home, or that the President of the bank is home? Might not a suitably ingenious experimenter devise a way for us to select the appropriate content from our list of contenders?

Bermúdez (2003) has persuasively argued that animal researchers are far cleverer than Davidson imagines. They have developed experimental methods for deciding between various attributions that are de re equivalent. By using these methods, Bermúdez claims that when we are confronted with a set of competing sentences that are de re equivalent, we can have reasons to prefer some of those sentences to others. In fact, Bermúdez goes even further (too far, I will ultimately contend) by arguing that we can take one of those sentences "to express *the* determinate content" of the animal's belief (Bermúdez, 2003, p. 103). In support of these contentions, Bermúdez works through some actual examples and shows how the subtle methods of animal researchers are more powerful than they might at first appear. Here's one.

Consider a rat that has learned the location of food in a cross-shaped maze. The rat starts in one arm (the southern arm, say), and the food is placed in another arm (say, the western arm). The rat learns to scurry directly from its starting location to the location with the food. As Bermúdez observes, there are many different ways that the location of the food might be coded from the rat's perspective. He isolates four in particular:

(1a) Food is located at the end-point of movements $M_{1 \ldots n}$.
(1b) Food is located at coordinates (x, y) in egocentric space.
(1c) Food is located at coordinates (x', y') in maze space.
(1d) Food is located at coordinates (x'', y'') in environmental space. (Bermúdez, 2003, p. 100)

Bermúdez then argues that by altering the experimental situation, one can acquire evidence to choose among these sentences. For example, in one modification of the experiment, the rat is started from a different arm (say, the northern arm). If the rat coded the location of the food in terms of the movements needed to reach the food (1a), or in terms of egocentric coordinates (1b), it would turn left at the intersection of the cross and thus wind up in the eastern arm, where it would fail to find food.

But as it happens, rats turn right at the intersection and thus find the food immediately. This provides evidence against using sentences (1a) or (1b) to express the content of the rat's belief about the location of the food. Now suppose that we physically rotate the maze 180 degrees and start the rat from the same physical arm we have always started it from. So if we always started the rat from the southern arm, then when we rotate the maze 180 degrees, we will start the rat from the northern arm. Now, does the rat turn right or left at the intersection? If it turns right, we have reason to rule out (1c), since it travels to a new location in the maze. But if it turns left, we have reason to rule out (1d) since it is going to a new place in environmental space. As it happens, the rat turns left. We thus have reason to favor (1c) as best expressing the content of the rat's belief about the location of the food. *Pace* the worries of philosophers such as Davidson, careful experimentation and observation can provide evidence for choosing among sentences that are de re equivalent to express the content of an animal's belief.[9]

The considerations Bermúdez amasses here are sufficient to blunt Davidson's worry. We have reason to prefer (1c) despite its de re equivalence to (1a), (1b), and (1d). While greater experimental ingenuity may be required when dealing with nonlinguistic creatures, there is no obvious reason why a lack of language will, as a matter of principle, limit us from having sufficient evidence to precisely ascribe contents to their cognitive representations. Language is one (admittedly convenient) source of evidence about what a subject thinks, but there are others.

The assessment that Indeterminacy does not derive from animals' inability to speak is supported by two further considerations. First, we can be confident in the precision of certain ascriptions to humans even when linguistic communication is not possible. If I'm standing next to a monolingual Italian man, and we're both watching the snow come down in the light of day, then even though I don't speak Italian I can be reasonably certain that the English sentence 'It's snowing' would precisely characterize the content of one of his beliefs. Second, giving animals language does not obviate the difficulty of precisely characterizing their cognitive contents. For example, Matsuzawa (1985) taught his chimpanzee, Ai, to use Arabic numerals to label sets of up to nine objects. But given the considerable differences between Ai's numerical behavior and that of humans, we lack grounds for supposing that she means the same thing by those numerals that we do. In fact, because Ai's pattern of response latencies suggests that she simply mapped Arabic numerals onto the very same underlying numerical representations that rats, pigeons, and monkeys possess (Matsuzawa, 2009), we face all the same problems interpreting Ai's understanding of number that we face interpreting the numerical representations of these other animals. We know that Ai has mapped the Arabic numeral '7' to one of her numerical representations, but we don't know how to precisely characterize that representation. Thus, we don't know what Ai means by '7'.

The muteness of animals is thus unlikely to explain Indeterminacy. Before we consider other potential sources of our inability to precisely characterize the representations of animals, however, I want to stress one point. Although Bermúdez's examples tell against Davidson's suggestion that language is necessary for selecting among contents that are de re equivalent, they do *not* establish that we are capable of capturing the cognitive contents of animals with as much precision as we capture our own, or that for any given animal belief state we will be able to find a sentence that expresses its content "accurately and without remainder" (Bermúdez, 2003, p. 104). In fact, reflection on his examples suggests just the opposite. Bermúdez, recall, selects 'Food is located at coordinates (x′, y′) in maze space' as best expressing the content of the rat's belief about the location of food. Yet even this sentence, though better than the others Bermúdez canvasses, is far from ideal. For one thing, it seems inappropriate to attribute the representations MAZE SPACE and COORDINATES to the rat. Granted, the rat does manage to navigate mazes with impressive success, but it doesn't give any indication of knowing that it is a *maze* that it is navigating. Nor does it seem wholly accurate to describe the rat as representing *coordinates*. Humans plausibly acquired the capacity to represent coordinates only after formal treatments of geometry were developed. So what business do we have attributing such a capacity to the rat? Moreover, the representation COORDINATES (X, Y) can only pick out a particular location by virtue of assigning numerical values to X and Y. But what units are those numerical values going to take? Do rats represent coordinates in inches? Do French rats prefer the metric system? With no justification for choosing among various units of measurement, we are left with myriad competing ascriptions that could equally claim to capture "the" content of the rat's belief about where the food is located.

Of course, Bermúdez might counter that we could refine his sentence to precisely express the content of the rat's belief. Perhaps there are particular units of distance that rats employ, and perhaps the appropriate experiments would allow us to discover them. My point is not that we can rule this possibility out, but simply that Bermúdez does not establish it. Thus, while Bermúdez shows that muteness does not present the insuperable difficulties that Davidson supposes, he does not demonstrate that animal researchers have overcome Indeterminacy altogether.

5. The Explanation from Unfamiliar Contents

The presence of language does not always secure understanding. That much should be clear to anyone who has attempted to translate a text from ancient Greek, taken a course in physics, or tried to decipher the curious utterances of a three-year-old. The student of ancient philosophy who begins reading Aristotle and encounters words such as 'psuche' and 'eudaimonia' will find translations of these terms difficult to come by. The student of physics who hears her teacher use words such as 'mass' and 'inertial frame' struggles to grasp the thoughts her teacher is trying to express. And the mother who is informed by her daughter that tables are alive because you can eat on them realizes that it will take some work to figure out what, precisely, her daughter could be saying (Carey, 1985, p. 30).

In these cases, our failures of understanding plausibly stem from the difficulty of grasping representations with unfamiliar contents. This suggests a third explanation of our inability to precisely characterize animals' cognitive representations: perhaps the contents of their representations are fundamentally different from the contents of our own. Since it seems to be a general fact that it is difficult for us to grasp unfamiliar contents, no wonder we lack the linguistic resources to precisely specify animals' cognitive contents.

There are, however, two problems with this explanation of Indeterminacy. First, in ordinary cases when we are confronted with representations with unfamiliar contents, we are eventually able to grasp those contents and precisely express them in language. The scholar that has spent her life poring over the *Nicomachean ethics* is not in the same position as the freshman who first comes across the term 'eudaimonia'. The scholar is able to understand this term, and use it in her explanations of Aristotle's ethical theory. Similarly, the student of physics need not spend eternity wondering what the phrase 'inertial frame' means. Through careful study, she can learn its meaning and deploy it to make sensible claims about the physical world. Thus, if Indeterminacy were merely a function of the differences in content between human and animal cognitive representations, we would expect Indeterminacy to evaporate once humans closely studied the minds of animals. But on my reading of the empirical literature, while animal researchers are learning a lot about the cognitive representations of animals, they are not making serious progress in *precisely* characterizing the contents of those representations.

One might reply that the progress of animal researchers is hampered because animal cognitive representations are just *so* fundamentally different from our own that they are extremely difficult for us to grasp—much more difficult than the contents *eudaimonia* or *inertial frame*. Perhaps they are even impossible for us to grasp given our cognitive limitations. But in fact (and this brings me to my second point), the empirical evidence suggests that humans and animals actually *share* a great many cognitive representations. Evolution has given rise to a large number of domain-specific cognitive systems, present in humans and animals alike, that are specialized to represent such fundamental features of the world as objects, numbers, agents, durations, and locations. Carey (2009) calls these shared systems "core cognition," and contrasts them with linguistically encoded conceptual representations.[10]

As an example of core cognition, consider the numerical representations we discussed at the end of section 2. Not only are such representations found in rats, pigeons, and monkeys, they are also found in human beings. Indeed, they appear to be innate, since even six-month-old infants have the capacity to represent numerical information in this rough way (Xu & Spelke, 2000). Although children from most cultures eventually acquire more advanced mathematical concepts that appear to be tied to their use of language (Carey, 2009), these more primitive, core numerical representations are never jettisoned. For example, if you're presented with an array of dots too quickly to count, you get an approximate sense of how many there are. Moreover, if you're then presented with a second array of dots, you can reliably determine which array has more so long as the ratio is 7:8 or greater, even when non-numerical confounds (such as total dot area or dot density) are controlled for

(Barth, Kanwisher, & Spelke, 2003). Like animals, your ability to discriminate numerical information is thus ratio sensitive, suggesting that your numerical discriminations are made possible by a system of representations that is homologous to the system that animals deploy. Evolution, it would seem, is reluctant to throw away cognitive tools that have been proven to work.

The numerical representations that humans share with many animals are only one example of core cognition. Carey (2009) and Spelke (2000) detail many other systems of representation that are likewise shared. Notice that this is a highly substantive empirical discovery. It might have turned out that with the advent of language and a larger associative cortex, all human cognitive representations were fundamentally altered. But if Carey and Spelke are right, that is not what happened. Rather, humans grafted a host of new conceptual representations onto a large base of core representations. This suggests that our inability to precisely characterize the cognitive representations of animals does not primarily stem from the unfamiliarity of the contents of animals' cognitive representations. There is simply too much overlap with our own cognitive representations for such a proposal to be plausible.[11]

The thesis that humans and animals share core cognition has two significant further consequences for our understanding of Indeterminacy. First, it suggests that whether or not we are able to linguistically characterize animals' cognitive contents, there is a sense in which we are perfectly able to *understand* some of those contents. Since we share core cognition with animals, we can gain some insight into their cognition simply by thinking with the more primitive aspects of our minds. We can *use* our representations from core cognition whether or not we can precisely characterize them in language.

On the flip side, the fact that we share core cognition with animals means that Indeterminacy is not solely a problem we face characterizing the cognition of other species. When we try to characterize *each other's* core cognitive representations (as opposed to each other's linguistically encoded conceptual representations), our characterizations lose their precision. For example, the attempts of experimenters to specify the contents of humans' core number representations face the very same obstacles that confront attempts to specify animals' numerical representations. This shouldn't be surprising given the evidence that humans and animals have the very same system of core number representations. But it means that interpersonal ascriptions sometimes prove just as imprecise as interspecies ascriptions; and in fact, the same can be said of *intra*personal ascriptions. When each of us tries to linguistically characterize our *own* representations from core cognition, we are unable to find the right words. Indeterminacy is thus a problem not only of interpretation, but also of conceptualization. When we linguistically conceptualize our own core representations, we distort their contents. To borrow a Quinean maxim, Indeterminacy begins at home.

Does this picture of Indeterminacy conflict with my claim in section 2 that Indeterminacy is worse for animals than humans? No. On the picture I have been discussing, humans differ from animals in having a wide array of conceptual representations that complement their core representations, and it is compatible with everything I have said that humans' conceptual representations are perfectly determinate. In that case, there would be many best cases in which the attributions of contents to humans would be precise (since they would be aimed at humans' conceptual representations), but almost no cases in which the attributions of contents to animals would be precise (since they by and large lack conceptual representations). Indeterminacy thus would be worse for animals than humans. Still, any adequate explanation of Indeterminacy will need to accommodate the fact that Indeterminacy arises for core cognition in general, whether in animals or humans. One virtue of the explanation from nonlinguistic format, to which we now turn, is that it satisfies this desideratum.

6. The Explanation from Nonlinguistic Format

As a way of warming up to the point I want to make, consider the following question: how can the content of the Mona Lisa be precisely translated into language?

The answer, I take it, is that it can't be. That's not to say that some linguistic expressions won't be better than others. Clearly the sentence 'There is a woman with a subtle smile' is a better translation than 'There is a fireman riding a donkey'. But there doesn't seem to be any one sentence (or set of sentences) that translates the Mona Lisa precisely. The representational format of a picture is simply too different from the representational format of language. According to the explanation from nonlinguistic format, this is the same sort of barrier that gives rise to Indeterminacy. We lack the ability to precisely characterize animal representations—and the representations of core cognition more generally—because they have the wrong format to be precisely translated into natural language.

The idea that Indeterminacy is a result of translating between representations of different formats depends on two claims: the general claim that representations of different formats cannot be precisely translated into one another; and the empirical claim that the format of animal cognition differs from the format of natural language. I will discuss each claim in turn.

6.1. Translating across Formats

Even translations from one natural language to another can be difficult to render with precision (just what is *ennui*, anyway?), but the difficulties are clearly enhanced when the media of translation have different formats. Consider, again, whether the content of the Mona Lisa can be precisely characterized in language. We might first try to appeal to a ''macro'' description of the painting—i.e., a description at the level of ordinary objects and things along the lines of 'Lisa del Giocondo is looking straight ahead, with a sly smile, wearing black clothing . . . '. The main problem with this sort of approach, however, is that it seems hopeless to acquire the level of detail that the painting conveys. There are many ways to smile slyly, and the Mona Lisa has just one of them.

Perhaps we should instead attempt to render the painting with a "micro" description—for example, by digitizing the picture into thousands of pixels to which we assign location and color values. One virtue of this approach is that it would be fine grained enough to allow us to create a replica of the original, perhaps even an indiscriminable replica. But instructions for a replica are not the same thing as a translation. Suppose, for example, that I gave you instructions for reproducing the content of the front page of today's *New York Times* on your computer by writing down the sequence in which you should press the buttons on your keyboard. "First press the key marked 'T' while holding down the shift key. Now release the shift key and press the key marked 'H' . . . " Clearly, my instructions would not have the same content as the front page of the *New York Times*. My instructions are about what keys you should press, not about war, economic recession, or the minor tribulations of Manhattan's left-leaning upper-middle class. The same problem befalls a digitized micro description of the Mona Lisa, since the painting is not *of* or *about* colored pixels. It's a painting of a particular woman—by most counts, Lisa del Giocondo. Thus, although a micro description would allow us to replicate the painting, it would not amount to a translation of the painting's content.[12] Moreover, even if we were to waive this worry, we would still be left with the problem that there are many ways to digitize a picture—for example, many different sizes and shapes we might assign to each pixel. Which of these is supposed to correspond to "the" content of the original Mona Lisa? With no reason to prefer any one to the others, we are led to the conclusion that *none* of them precisely captures the painting's content.[13]

Representational formats differ in their *modes of composition*—that is, in their primitive parts and the rules that govern how those parts can be combined. For example, characterizations of the format of human language typically appeal to such discrete parts as singular terms, predicates, and logical connectives, along with familiar rules that dictate how such parts can be combined to form sentences. By contrast, pictures seem to have quite different modes of composition. For one thing, the parts of a picture are not well defined and discrete. Rather, pictures are compositionally homogenous (Fodor, 2008, pp. 172–177). Every part of a picture contributes to the content of the whole picture in the same way. Moreover, the rules that govern how the parts of a picture combine to determine the content of the whole are essentially spatial. The parts of a picture represent spatial parts of what the picture as a whole represents. By contrast, the 'not' in 'John is not tall' does not represent any spatial part of the sentence's content.

Although I have been contrasting the format of human languages with the format of pictures, I do not mean this dichotomy to be exhaustive, such that all representations must be either linguistic or pictorial. If a representation's format is individuated by its mode of composition, we should expect there to be many different representational formats, of which human language and pictures will be just two. Thus, the claim that animal cognition isn't formatted like human language does not entail that it is formatted like a picture. As we will see below, there are other options. For the moment, however, the crucial claim I want to make plausible is that translation between any two media of representation that have sufficiently different modes of composition cannot be rendered precisely. To lend credence to this general claim, it will be helpful to move from a discussion of specific examples to a more theoretical perspective that brings into focus the important role that compositionality plays in translation.

If we are presented with a novel language lacking in sub-sentential structure, it will be difficult, if not impossible, to discern differences in meaning between sentences that have the same truth conditions. For example, if R and S are both necessarily true sentences with different meanings from a language with which we are unacquainted, it is hard to see how we could figure out what they mean. Certainly appealing to the conditions under which they are affirmed or denied by speakers of the language (as a proponent of radical translation would recommend) would not provide us with the information we need. Notice how much easier our task becomes, however, if the sentences are structured. Suppose, for example, that the first sentence has the form '$1+1=2$' and the second has the form '$3^2=9$'. We can then see how the components of each sentence are used in other sentences, including sentences that are judged false by speakers of the language, and we can use that information to draw conclusions about how each of the constituents of the language map onto the constituents of our own language. Assuming that the language is compositional, we can thereby derive translations that distinguish sentences with identical truth conditions.

I do not mean to suggest that we should follow Quine in holding that the *only* evidence relevant to translation comes from the assent and dissent of speakers. As Chomsky (1992/2000) has persuasively argued, empirical theories of any stripe should not be guided by a priori limits on their domains of evidence. But I am suggesting that compositionality plays a critical role in our capacity to translate between languages. Part of the reason we can translate as well as we can from sentences of French to sentences of English is that we can translate the individual *words* of French into the individual *words* of English. Consequently, if two media of representation were to have different modes of composition, and thus involve different representational formats, we should expect precise translations between them to falter.

To further support the idea that similar modes of composition are crucial for precise translation, I want to recall Evans' response to Quine's famous gavagai argument. Quine (1960) imagines a group of speakers that assent to 'Gavagai!' when and only when a rabbit is present. Are we thereby justified in interpreting 'gavagai' to mean *rabbit*? No, argued Quine, since the speakers' dispositions to assent are equally compatible with several other translations, including *undetached rabbit parts*, *rabbit fusion*, *temporal stage in the life of a rabbit*, and *rabbithood*. Evans (1975) responds that Quine's argument fails to account for the full grammatical role that predicates play in sentences of a natural language, particularly in combination with other predicates and negation. For example, Evans argues that the word 'gavagai' couldn't mean *undetached rabbit parts* once we take into account how it interacts with other predicates. Suppose that the native speakers assent to 'F' when and only when something white is present, and that they assent to 'F gavagai' when and only when a white rabbit is present. In particular, they don't assent to 'F gavagai' when a brown rabbit with white feet is present, telling against the hypothesis that 'gavagai' means *undetached rabbit parts* (since the white feet of the rabbit are exemplary instances of white undetached rabbit parts). By considering how 'gavagai' composes with other terms, we can thus acquire evidence to constrain our interpretation.

Whether Evans' criticisms of Quine's argument are ultimately successful is a vexed question that I do not wish to consider here.[14] My aim in raising this debate is altogether more modest; I simply want to draw attention to the fact that similarities in the modes of composition of two media of representation help to reduce indeterminacy by constraining the range of plausible translations between them. Whether or not we take Evans' objections to ultimately succeed against Quine, it should be relatively uncontroversial that they illustrate the role that considerations of compositional similarity can play in making our translations more precise. Thus, where two media of representation have different modes of composition, we shouldn't be surprised if we cannot precisely translate between them.

6.2. The Format of Animal Cognition

Do the cognitive representations of animals really differ in format from human language? This is a difficult empirical question that I cannot hope to settle here. However, there are two reasons to think that the answer may well be affirmative.

First, there is evidence that specific systems of core cognition have a format that differs significantly from that of language. For example, the fact that the numerical representations discussed earlier are discriminated in accordance with their ratio suggests that they have an analog format. To see why, notice that it is no harder to visually discriminate the digit '7' from the digit '8' than it is to discriminate the digit '1' from the digit '2'. But as Carey (2009) observes, when we consider analog representations of these numerical values, the first pair (________ versus ________) is harder to discriminate than the second pair (_ versus __). Because analog representations of number are non-arbitrarily related to their referents, they become more difficult to discriminate as their ratio approaches one. As a result, most researchers take the existence of a ratio effect as evidence that core numerical representations lack the digital format of natural language and have an analog format instead. Of course, no one thinks that the brain uses line lengths to represent numerical values, but the existence of the ratio effect suggests that it likely uses some other quantity, such as neural firing rates, to serve as a direct analog of the numerical values it represents. If that is right, then it may well be that the reason that it has been so difficult to precisely characterize the numerical representations of animals is that they have a different format from natural language.[15]

There is also evidence that some of the representations that animals use for navigation are structured less like sentences, and more like maps. In a set of experiments performed by Cheng and Gallistel (Cheng, 1986; Cheng & Gallistel, 1984; Gallistel, 1990), rats were placed in a rectangular room and allowed to find food, which was hidden in (say) the northeast corner. The room had a variety of cues which could have been of use to the rats in remembering the location of the food. For example, the eastern wall was painted red, the western wall had a rough texture, and the northern wall smelled like lemon. (We know from other experiments that rats are able to represent color, texture and smell.) Having learned the location of the food, the rats were taken out of the room, disoriented, and then placed back in the center of the room. Being hungry, they sought out the corner with the food. But to everyone's surprise, they didn't only search in the northeast corner. Rather, they searched equally often in two corners: the northeast corner and the southwest corner. This is surprising because the rats could have used any of the cues to remember the location of the food. For example, they could have remembered that the food was in the corner to the left of the red wall, or in the corner to the right of the lemon-scented wall, but instead they relied purely on the geometry of the room. From a geometric perspective, the diagonally opposite corners of a rectangular room are metrically indiscernible. Whether you're facing the northeast corner or the southwest corner there will be a long wall to your left and a short wall to your right. The rats, it would seem, only code the location of the food relative to the geometric properties of the room. On the assumption that the rat is deploying a cognitive map—a representation with geometric structure—this result makes sense. Moreover, if this assumption is right, it can help to explain why even the best of the sentences Bermúdez proposed to describe the content of the rat's beliefs about where food is located in its environment seem imperfect. If the rat's representation is map-like, no wonder we can't precisely express it in a sentence.

In a number of recent papers, Rescorla (2009a, 2009b, 2009c) has explored the notion of cognitive maps. According to Rescorla, the distinguishing feature of cognitive maps is that they are governed by the axioms of a recognized geometry, such as that of a metric space. A map is then accurate when the geometric relations among its parts are isomorphic to the spatial relations among the objects it represents.[16] While maps are compositional (and Rescorla even provides a compositional semantics for them), the parts from which they are composed and the rules that govern how they compose are quite different from the parts and rules possessed by sentences of natural language. The components of a map do not include predicates, quantifiers, or logical connectives. Roughly put, the mode of composition of a map is geometric rather than logical. Rescorla's discussion thus illustrates how a compositional system of representation can support sophisticated forms of cognition without being structured like human language.

There have been other suggestions in the philosophical literature about nonlinguistic forms of animal cognition. For example, Camp (2009) argues that representations of social dominance relations by baboons have a diagrammatic tree-like structure that contrasts with the logical format of natural language. Similarly, Bermúdez (2003) contends that animal inferences involve a "proto-logic" that lacks the full-blown logical apparatus that is presumably present in natural language. The possibility thus arises that animal cognition (as well as human core cognition) involves myriad forms of nonlinguistically structured representations.

A second reason to think that animal cognition has a nonlinguistic format derives from considerations of what animals *cannot* do. While humans have the capacity to execute logical inferences, perhaps even innately (Crain & Khlentzos, 2010), the evidence that animals can do the same is at best equivocal (Bermúdez, 2003; Penn & Povinelli, 2007; Rescorla, 2009a). Philosophers have occasionally suggested otherwise. For example, there is the famous old rumor, deriving from the stoics, that a dog that is chasing its prey and arrives at a three-way fork will sniff the first two paths and, not finding a scent, run immediately down the third path, thus providing evidence that it has executed a disjunctive syllogism. But there are two problems with taking such rumors at face value. First, it is far from obvious that the behavior they describe can only be explained through an appeal to logical inferences. For instance, Rescorla (2009a) argues that Bayesian inferences over cognitive maps can explain the behavior the stoics report. Second, such rumors do not always withstand the scrutiny of controlled experiments, where the standards of evidence require more than anecdotes. For example, when one group of researchers tested the stoics' hypothesis by hiding a desirable object behind one of three screens, they found that while human children from four to six years of age quickly searched behind the third screen after failing to find the object behind either of the first two, dogs displayed the opposite pattern. They actually took *longer* to search behind the third screen after failing to find the object behind either of the first two (Watson et al., 2001). On the assumption that animal cognitive representations are language-like, and thus have something like predicative and logical structure, the inability to perform logical inferences is rather hard to understand. After all, sentence-like representations are the ideal vehicles for logical inference. If animals are really incapable of performing logical inferences, their cognitive representations are thus unlikely to have the sort of logical form that is familiar from natural language.

I want to emphasize that the extent to which animals are capable of logical inference is still hotly contested among animal researchers. It would thus be premature to claim that animals are definitely incapable of all forms of logical inference. But even on the most sanguine interpretation of what animals can do, it seems doubtful that they are capable of anything like the full array of logical inferences that humans routinely carry out, suggesting that the format of their representations is almost certainly very different from the format of human languages. I have argued that if that is right, it can help to explain our inability to precisely translate animal cognition into natural language, thus reconciling Realism and Indeterminacy.

The explanation from nonlinguistic format makes a prediction. As we get closer to discovering the exact formats of animal cognition, we should get closer to precisely characterizing the contents of animal cognition. Our characterizations just need to be couched in the nonlinguistic formats we discover. Some support for this prediction can be gleaned from particular instances in which nonlinguistic representations are used to characterize animal cognition, such as Carey's use of the analog medium of line lengths to characterize core numerical representations, and Rescorla's use of maps to characterize animals' representations of the spatial layout of their environments. These characterizations really do seem to get closer to capturing the contents of animals' cognitive representations. But that is not to say that such characterizations are perfect. The formats of animal cognition are surely highly complex, and the line lengths and maps to which Carey and Rescorla appeal are only first approximations in the direction of those formats. As we gain greater insight into the compositional mechanisms that typify animal cognition, our characterizations of animal cognition should improve.

Before leaving this section, I want to consider an objection to the explanation from nonlinguistic format. One might worry that if animal representations cannot be characterized in language, they must lack determinate contents, leading to a conflict with Realism, which requires animal representations to have determinate contents. It would be a mistake, however, to suppose that representations that cannot be characterized in language must lack determinate contents. Such a supposition would only hold if all contents had to be linguistically expressible. There is, however, no obvious reason to accept this latter claim. In fact, many philosophers have rejected the hegemony of linguistic content by arguing that there exists an important class of contents—*nonconceptual* contents—that by their very nature are not linguistically expressible. The possibility thus remains open that animals' representations have perfectly determinate nonconceptual contents—a possibility that some philosophers have defended on independent grounds (Beck, forthcoming; Bermúdez, 1998). If that is right, then Indeterminacy is inevitable. No amount of studying animals will allow us to express the contents of their representations in language. At the same time, Indeterminacy is no knock on animals. Their representations have contents that are every bit as real and determinate as our own. Yes, it is indeterminate how to translate animal cognition into human language. But so too is it indeterminate how to translate human language into animal cognition.

7. Conclusion

We started with a puzzle: if animals really have cognitive representations with determinate contents, why can't anyone say precisely what contents their representations have? This puzzle is especially pressing for Realists because so many philosophers have brandished it as a weapon against Realism. It is thus incumbent on Realists to show how Indeterminacy can be explained away.

As we have seen, Realists are not without options. First, they can maintain that we simply haven't spent enough time studying the representations of animals. As animal researchers carry out additional observations and experiments, perhaps they will zero in on precise characterizations of animal representations. While I granted that this possibility couldn't be ruled out, I also argued that reflection on animal research seems to tell against it. Researchers are learning a lot about animals' representations of, say, other minds and numbers. But there is little evidence that they are converging on precise linguistic characterizations of these representations. I thus suggested that there might be a more principled explanation of Indeterminacy.

One type of principled explanation is epistemic. Perhaps there are linguistic expressions that precisely characterize animals' cognitive representations, but we humans are blocked from discovering those expressions because of our epistemic limitations. For example, perhaps the fact that animals don't talk prevents us from knowing precisely what they are thinking. Or perhaps the contents of animals' representations, though expressible in language, are too unfamiliar for us to grasp. Although such epistemic explanations would allow the Realist to accommodate Indeterminacy, I argued that they are implausible. While language is a convenient source of information about a thinker's cognitive contents, *pace* Davidson there is little reason to think that it is indispensable to interpretation. Moreover, the fact that humans and animals share core cognition suggests that Indeterminacy is not simply a product of unfamiliar contents. Many of the representations that we struggle to characterize in animals turn out to be present in us too.

The explanation from nonlinguistic format succeeds where the explanation from unfamiliar contents does not: it explains why intimately familiar representations sometimes elude precise expression. For example, if we find it difficult to characterize animals' numerical representations because they have a nonlinguistic format, and our core numerical representations have the same format as those of animals, then of course we will have difficulties characterizing our core numerical representations too. Moreover, the explanation from nonlinguistic format naturally explains why we seem to be making so little progress in overcoming Indeterminacy since it alone maintains that Indeterminacy is inevitable. It is just what you get when you try to translate across representations with disparate formats. Thus, if the explanation from nonlinguistic format is correct, we are not somehow *blocked* from discovering the linguistic expressions that characterize animal and core cognition. Rather, there are no such expressions to discover.

Acknowledgments

For comments on earlier drafts, I am grateful to John Doris, David M. Kaplan, Peter Langland-Hassan, Alison Simmons, and two anonymous referees for this journal. Thanks also to Gillian Russell for the title, to Edouard Machery for a helpful discussion of the Mona Lisa, and to the audience at Texas Tech University where I presented an early version of this material. This paper was completed with the assistance of a postdoctoral fellowship from the James S. McDonnell Foundation and Washington University in Saint Louis.

Notes

[1] In fact, Dennett and Jamieson embrace interpretivism for most of *human* cognition as well, though Dennett (1991) carves out an exception for what he calls "opinions," which are roughly dispositions to assent to sentences. Because humans have opinions in addition to their other cognitive states, there is a sense for Dennett in which interpretivism is less fully true of humans than of animals.

[2] A very similar example is discussed by Malcolm (1972–3) and Davidson (1982/2001).

[3] Similar worries can also be found in Davidson (1975, 1982/2001), Dennett (1987, pp. 103–116, 1991), Jamieson (2009, pp. 17–18), and Tetzlaff and Rey (2009, pp. 74–75).

[4] Another reason to be unhappy with holism as an explanation for Indeterminacy is that it only pushes the explanation back one step. It tells us that Fido lacks one belief because he lacks another belief, but it doesn't tell us why his beliefs *in general* seem to be subject to Indeterminacy.

[5] Notice that this justification for accepting Indeterminacy (that animals' behavior evinces little evidence that they possess the concepts expressed by our words) is compatible with different perspectives on the reality of animal cognitive representations, and with different explanations of the source of Indeterminacy (including the explanations from insufficient time, muteness, unfamiliar contents, and nonlinguistic format). Thus, we can agree with Stich, Davidson, Dennett, and other early participants in this debate that the behavior of animals is evidence of Indeterminacy without necessarily endorsing either their anti-realism about animal cognitive representations or their preferred theoretical explanations of Indeterminacy.

[6] This stance obviously commits me to rejecting the soundness of various arguments that purport to show that human cognition is marked by considerable indeterminacy (e.g., Quine 1960). Such arguments deserve attention, but cannot be considered here.

[7] Compare Povinelli and Vonk (2003, 2004) to Tomasello et al. (2003, 2004). For a nice summary of the dispute, see Rosati, Hare, and Santos (2010).

[8] Burge (2010, p. 482) suggests that an appeal to Eudoxes' theory of pure magnitudes might help, but he doesn't fill in the details, and I'm not aware of anyone else who has tried.

[9] Bermúdez (2003) embeds his approach to animal cognition in a controversial success semantics. But his argument that animal researchers can distinguish between contents that are de re equivalent is independent of this feature of his view.

[10] Spelke (2000) uses the term 'core knowledge' to isolate the same systems of representation, but Carey's 'core cognition' avoids the misleading implication that the representations are always justified and accurate.

[11] This is not to deny that the explanation from unfamiliar contents will apply in some cases. For example, five-year-old children, adults with Alzheimer's disease, and people from exotic cultures plausibly have language-like conceptual beliefs that are difficult to characterize because their contents are so unfamiliar to us. But given the considerable overlap of core cognition across humans and animals, another explanation is needed to accommodate our inability to precisely characterize many of the contents of animals' cognitive representations.

[12] Haugeland (1998, chapter 8) makes a similar point.

[13] Heck (2007) argues that similar problems befall any attempt to translate a map into a sentence.

[14] For different takes on the success of Evans' objections, see Wright (1997) and Burge (2010, pp. 216–223).

[15] This point is developed in Beck (forthcoming).

[16] The maps we use on an everyday basis are literally spatial. You can post them on a wall, or fold them up and put them in your pocket. It is thus tempting to suppose that maps need to be spatial in character, and this can make the idea of a cognitive map seem psychologically implausible. Are we to suppose that cognitive maps are literally spread across the hippocampus? Who would read them? But as Rescorla (2009b) observes, this worry reflects an insufficient grasp of the abstractness with which mathematicians treat geometric notions such as that of a metric space. Any set of objects can constitute a metric space so long as there is some distance relation with respect to which they satisfy the axioms for a metric. The distance relation does not have to be one of spatial distance. For example, even the set of ordered pairs of real numbers satisfies the axioms for a metric and can thus constitute a map.

References

Allen, C. (1992). Mental content. *The British Journal for the Philosophy of Science*, *43*, 537–553.

Barth, H., Kanwisher, N., & Spelke, E. (2003). The construction of large number representations in adults. *Cognition*, *86*, 201–221.

Beck, J. (forthcoming). The generality constraint and the structure of thought. *Mind*.

Bermúdez, J. (1998). *The paradox of self-consciousness*. Cambridge, MA: MIT Press.

Bermúdez, J. (2003). *Thinking without words*. New York: Oxford University Press.

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

Call, J., & Tomasello, M. (1999). A nonverbal false belief task: The performance of children and great apes. *Child Development*, *70*, 381–395.

Camp, E. (2009). A language of baboon thought? In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 108–127). New York: Cambridge University Press.

Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press.

Carey, S. (2009). *The origin of concepts*. New York: Oxford University Press.

Cheng, K. (1986). A purely geometric module in the rat's spatial representation. *Cognition*, *23*, 149–178.

Cheng, K., & Gallistel, C.R. (1984). Testing the geometric power of an animal's spatial representation. In H. Roitblat, T. Bever, & H. Terrace (Eds.), *Animal cognition* (pp. 409–423). Hillsdale, NJ: Erlbaum.

Chomsky, N. (2000). Language and interpretation: Philosophical reflections and empirical inquiry. In his *New horizons in the study of language and mind* (pp. 46–74). Cambridge: Cambridge University Press. (Original work published 1992.)

Crain, S., & Khlentzos, D. (2010). The logic instinct. *Mind & Language*, *25*, 30–65.

Davidson, D. (1984). Thought and talk. In his *Inquiries into truth & interpretation* (pp. 155–170). New York: Oxford University Press. (Original work published 1975.)

Davidson, D. (2001). Rational animals. In his *Subjective, intersubjective, objective* (pp. 95–106). New York: Oxford University Press. (Original work published 1982.)

Dehaene, S. (1997). *The number sense: How the mind creates mathematics*. New York: Oxford University Press.

Dennett, D. (1969). *Content and consciousness*. New York: Humanities Press.

Dennett, D. (1987). *The intentional stance*. Cambridge, MA: MIT Press.

Dennett, D. (1991). Two contrasts: Folk craft versus folk science, and belief versus opinion. In J. Greenwood (Ed.), *The future of folk psychology: Intentionality and cognitive science*. New York: Cambridge University Press.

Evans, G. (1975). Identity and predication. *Journal of Philosophy*, *72*, 343–363.

Fodor, J. (1975). *The language of thought*. Cambridge, MA: Harvard University Press.

Fodor, J. (2008). *LOT 2: The language of thought revisited*. New York: Oxford University Press.

Fodor, J., & Lepore, E. (1992). *Holism: A shopper's guide*. Oxford: Blackwell.

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

Gallistel, C.R., & Gelman, R. (2000). Non-Verbal numerical cognition: From reals to integers. *Trends in Cognitive Sciences*, *4*, 59–65.

Haugeland, J. (1998). *Having thought: Essays in the metaphysics of mind*. Cambridge, MA: Harvard University Press.

Hauser, M., Tsao, F., Garcia, P., & Spelke, E. (2003). Evolutionary foundations of number: Spontaneous representation of numerical magnitudes by cotton-top tamarins. *Proceedings of the Royal Society, B*, *270*, 1441–1446.

Heck, R. (2007). Are there different kinds of content? In J. Cohen & B. McLaughlin (Eds.), *Contemporary debates in philosophy of mind* (pp. 117–138). Oxford: Blackwell.

Jamieson, D. (2009). What do animals think? In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 15–34). New York: Cambridge University Press.

Malcolm, N. (1972–3). Thoughtless brutes. *Proceedings and Addresses of the American Philosophical Association*, *46*, 5–20.

Matsuzawa, T. (1985). Use of numbers by a chimpanzee. *Nature*, *315*, 57–59.

Matsuzawa, T. (2009). Symbolic representation of number in chimpanzees. *Current Opinions in Neurobiology*, *19*, 92–98.

Mechner, F. (1958). Probability relations within response sequences under ratio reinforcement. *Journal of the Experimental Analysis of Behavior*, *1*, 109–121.

Nagel, T. (1974). What is it like to be a bat? *The Philosophical Review*, *83*, 435–50.

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

Povinelli, D., & Eddy, T. (1996). What young chimpanzees know about seeing. *Monographs of the Society for Research in Child Development*, *61*(3), 1–152.

Povinelli, D.J., & Vonk, J. (2003). Chimpanzee minds: Suspiciously human? *Trends in Cognitive Sciences*, *7*, 157–160.

Povinelli, D.J., & Vonk, J. (2004). We don't need a microscope to explore the chimpanzee's mind. *Mind & Language*, *19*(1), 1–28.

Premack, D.G., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, *1*, 515–526.

Putnam, H. (1962). It ain't necessarily so. *Journal of Philosophy*, *59*, 658–671.

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

Quine, W.V.O. (1960). *Word and object*. Cambridge, MA: MIT Press.

Rescorla, M. (2009a). Chrysippus' dog as a case study in non-linguistic cognition. In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 52–71). New York: Cambridge University Press.

Rescorla, M. (2009b). Cognitive maps and the language of thought. *British Journal for the Philosophy of Science*, *60*, 377–407.

Rescorla, M. (2009c). Predication and cartographic representation. *Synthese*, *169*, 175–200.

Rilling, M., & McDiarmid, C. (1965). Signal detection in fixed ratio schedules. *Science*, *148*, 526–527.

Rosati, A., Hare, B.A., & Santos, L.R. (2010). Primate social cognition: Thirty years after Premack and Woodruff. In M. Platt & A.A. Ghazanfar (Eds.), *Primate neuroethology* (pp. 117–143). New York: Oxford University Press.

Spelke, E.S. (2000). Core knowledge. *American Psychologist*, *55*, 1233–1243.

Stich, S. (1979). Do animals have beliefs? *Australasian Journal of Philosophy*, *57*, 15–28.

Stich, S. (1983). *From folk psychology to cognitive science*. Cambridge, MA: MIT Press.

Tetzlaff, M., & Rey, G. (2009). Systematicity and intentional realism in honeybee navigation. In R. Lurz (Ed.), *The philosophy of animal minds* (pp. 72–88). New York: Cambridge University Press.

Tomasello, M., Call, J., & Hare, B. (2003). Chimpanzees understand psychological states—The question is which ones and to what extent. *Trends in Cognitive Sciences*, *7*, 153–156.

Tomasello, M., Call, J., & Hare, B. (2004). Chimpanzees versus humans: It's not that simple. *Trends in Cognitive Sciences*, *7*, 239–240.

Watson, J.S., Gergely, G., Csanyi, V., Topal, J., Gacsi, M., & Sarkozi, Z. (2001). Distinguishing logic from association in the solution of an invisible displacement task by children (*Homo sapiens*) and dogs (*Canis familiaris*): Using negation of disjunction. *Journal of Comparative Psychology*, *115*, 219–226.

Wittgenstein, L. (1973). *The philosophical investigations*. Oxford: Blackwell.

Wright, C. (1997). The indeterminacy of translation. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language* (pp. 397–426). Oxford: Blackwell.

Xu, F., & Spelke, E. (2000). Large number discrimination in 6-month-old infants. *Cognition*, *74*, B1–B11.