Philosophy and
Phenomenological Research

Philosophy and Phenomenological Research
doi: 10.1111/phpr.12455

Rational Inference: The Lowest Bounds

CAMERON BUCKNER
University of Houston

A surge of empirical research demonstrating flexible cognition in animals and young infants has raised interest in the possibility of rational decision-making in the absence of language. A venerable position, which I here call "Classical Inferentialism", holds that nonlinguistic agents are incapable of rational inferences. Against this position, I defend a model of nonlinguistic inferences that shows how they could be practically rational. This model vindicates the Lockean idea that we can intuitively grasp rational connections between thoughts by developing the Davidsonian idea that practical inferences are at bottom categorization judgments. From this perspective, we can see how similarity-based categorization processes widely studied in human and animal psychology might count as practically rational. The solution involves a novel hybrid of internalism and externalism: intuitive inferences are psychologically rational (in the explanatory sense) given the intensional sensitivity of the similarity assessment to the internal structure of the agent's reasons for acting, but epistemically rational (in the justificatory sense) given an ecological fit between the features matched by that assessment and the structure of the agent's environment. The essay concludes by exploring empirical results that show how nonlinguistic agents can be sensitive to these similarity assessments in a way that grants them control over their opaque judgments.

A rising tide of empirical and philosophical work has revealed remarkable intelligence in the judgments of nonlinguistic animals, young children, and adult intuition (for recent reviews see Andrews, 2014; Banaji & Gelman, 2013; Brownstein & Saul, 2016a, 2016b; Gigerenzer, Hertwig, & Pachur, 2011; Gilovich, Griffin, & Kahneman, 2002). The judgments surveyed in this literature fall into what Nagel (2012) has recently called the "murky zone" between the rational heights of explicit deduction and the a-rational lows of mindless reflex. The latest cognitive science shows that these judgments can be highly flexible and reliably successful in even novel and dynamic circumstances, but their structure appears to be non-classical in form and is often opaque to introspection. We thus need to decide when such judgments should be pushed down to the non-inferential level by pointing out their automaticity and lack of integration with other epistemic states (e.g. Gendler, 2008), upgraded to inferential status by emphasizing their sensitivity to evidence and inferential promiscuity (e.g. Mandelbaum, 2013), or declared a sui generis middle ground that fits into neither of these better-understood categories (e.g. Levy, 2015).

In this paper, I aim to establish the lowest bounds of rational inference, arguing that many (though not all) of these opaque judgments in nonlinguistic animals should be counted as inferential. To clarify, there is a trivial sense of 'inference', common in cognitive psychology, on which any information processing counts as inferential; this is not the sense of inference explored here. This more philosophical notion begins instead with the idea that, at minimum, rational inference (hereafter, just 'inference') is the mental process of arriving at a conclusion on the basis of reasons which support it. Theoretical inference involves the process of arriving at one belief on the basis of others; for example: Fluffy is a cat, all cats love liver, therefore Fluffy loves liver. Practical inference, on which I shall focus here, involves instead the selection of an action to perform on the basis of beliefs and desires; for example: Fluffy desires to jump on the counter, Fluffy believes—after sufficient crouching and wiggling—that leaping from the top of the chair will land her on the countertop. . .so straightaway Fluffy leaps.

But wait—many readers may balk at the idea that the judgment of an animal like Fluffy is even eligible to count as an inference. It may appear that Fluffy is weighing her evidence, trying to determine whether she can make the leap or not; but any theory here must articulate criteria to distinguish inferences—drawn by agents on the basis of reasons—from processes in which psychological states non-mentally cause a behavioral outcome. For example, students may answer a test question incorrectly due to a pounding headache, or say that Tide is their favorite brand of laundry detergent because a sneaky professor primed them earlier in a lecture with the words 'moon', 'shifting', 'waves', and 'sea'; but since in neither of these cases are the processes determining the behavioral outcomes sensitive to any support that the headache or priming may or may not provide for their choice, these processes are not inferential in nature. These causal connections are a-rational, given that the headaches and lexical primes do not bring about these behavioral effects because they possess some particular representational contents. A theory of inference must minimally solve this "demarcation problem" by identifying a shared character that inferences possess and non-inferential judgments lack.

In this paper, I evaluate two popular ideas which have been deployed to police this boundary between inferential and non-inferential psychological causation: 1) the "Taking condition", which holds that inferring necessarily involves thinkers somehow *taking* their premises to support their conclusions (Boghossian, 2014; Valaris, 2016), and 2) a view about the form of practical inference that I will hereafter call "Classical Inferentialism" (CI), which holds that one's premises (inferentially) support their conclusion when reached by syntactically fitting those premises and conclusion together in the schema of a formal rule, such as those provided by deductive logic or decision theory. Each idea is independently popular, and they are commonly paired—despite frequently acknowledged tensions. Taken together, these ideas preclude the possibility that nonlinguistic agents could draw genuine inferences, because using formal inference rules probably requires language.[1]

I begin in Section 1 by justifying the need for a new theory of nonlinguistic inference and offering some caveats about its scope. Section 2 offers a series of putative examples of nonlinguistic inference to open the discussion. Section 3 provides a deeper exposition of Taking and CI to clarify the bar such judgments need to clear, and Section 4 reframes the role that rationality plays in this debate by understanding practical inference more broadly as a species of intuitive categorization judgment. I then conclude in Section 5 by considering a series of objections to the effect that the proposed model should still not count as genuinely inferential, the rebuttals highlighting the forms of inference-like control that nonlinguistic agents can exert over their judgments.

[1] The requirement could be satisfied by either a natural language or a language of thought; but nonhuman animals likely possess neither (see below).

1. Motivations and Caveats

By denying inferential status to nonlinguistic judgments, the combination of Taking and CI stands in tension with trends in cognitive science, where researchers increasingly describe the judgments of nonlinguistic animals and infants in inferential terms. Tapping into a methodology in comparative psychology that can be traced back to Darwin and Morgan (Clatterbuck, 2016), these researchers treat these judgments as inferential because they enable subjects to reliably succeed in perceptually novel circumstances where other forms of behavioral explanation—especially in terms of reflex, instinct, and basic forms of associative learning—are inadequate (Buckner, 2011; Heyes & Dickinson, 1990; Ristau, 1991). Instead, they explain these judgments by appeal to the representational contents of those agents' cognitive and conative states, especially their tracking of abstract, higher-order, spatial, causal, or psychological relations in their environment (Buckner, 2015). They thus seem to satisfy at least one core idea associated with Taking and CI: that some nonlinguistic judgments can be sensitive to the representational contents of the cognitive and conative states governing the agent's actions.

Recent philosophical accounts of opaque judgments in adult humans, however, do not address the needs of the scientists who are here attempting to mark out the lower bounds of inferential judgment in nonlinguistic agents. An instructive comparison here is Mandelbaum (2013), because he also thinks that many opaque judgments in humans should be considered inferential. His view begins as a criticism of a view attributed to Gendler (2008) that implicit attitudes are associatively-structured representational states dubbed "aliefs". Specifically, Mandelbaum argues against the need for a novel category of alief with a two-pronged attack: by 1) pointing out that old-fashioned propositionally-structured beliefs can be associated with valences and motor responses, accounting for some of the key phenomena that aliefs were introduced to explain, and 2) arguing that the representational contents of implicit attitudes cannot be purely associative—that is, at least some must have internal logical structure—because they can feature in deductively valid inferences and are responsive to evidence in sophisticated ways. In more recent work, Mandelbaum further argues that if the internal structure of implicit attitudes were purely associative then they should be nothing more than "concatenations of mental states such as the coactivation of two concepts…or of a concept and a valence" (Mandelbaum, 2015, 633). If so, then they should be eradicated by behaviorist training techniques like counterconditioning and extinction, where the contingencies between a cue and a reward are reversed or eliminated; but recent empirical work has shown that many implicit attitudes are not so easily eliminated, so Mandelbaum concludes that they cannot have a purely associative structure.

Mandelbaum's strategy offers a powerful argument for the conclusion that many opaque judgments in adult humans involve logical form, but it is not the right instrument for the demarcation cut currently needed in nonlinguistic agents. The problem is that the argument relies on a dichotomy between the simplest forms of association on the one hand (which should i.e. be eradicated by counterconditioning and extinction) and logical structure on the other; but we already know that many of the candidates for inference in animals and infants do not fit neatly into either of these two categories. These judgments cannot be classified as merely associative on Mandelbaum's dichotomy, because nearly every experiment in the field of comparative cognition has been designed to rule out the possibility that the behaviors observed can be explained by the simplest forms of classical and operant conditioning (Buckner 2015).[2] Neither do the animal and infant judgments fit on the propositional side of this dichotomy, because animals and nonlinguistic infants probably cannot represent either logically-structured contents (Bermúdez, 2003; Rescorla, 2009a) or the syncategorematic logical relations required to deploy those structures in classically-valid judgments (Watanabe & Huber, 2006).

Much of the difficulty here is that associative learning theory has progressed in leaps and bounds since the days of Watson and Skinner. It had already gone beyond counter-conditioning and extinction with the introduction of the Rescorla-Wagner model of classical conditioning in the 1970s, which supposes that associative learning is not automatic but rather involves an active role for attention and prediction (Rescorla, 1972). On this model, conditioning only occurs if animals actively attend to the relevant stimuli and have not already learned that certain cues predict the observed outcomes.[3] Furthermore, occasion-setting and context-sensitivity—wherein some associations "gate" the activation of others, allowing different responses or valences to be associated with the same stimulus in different contexts—have been stock components of the associative toolkit since the 1980s (Bouton & Swartzentruber, 1986). Occasion-setting could notably explain failures of extinction or counterconditioning without appealing to logical structure, for it could allow subjects to treat the original training situation and counterconditioning/extinction situation as different contexts.

Most relevantly, configural learning models—which take the internal structure of categorical representations to be context-sensitive configurations of cues, wherein the effect of any particular cue depends upon the values of many other cues in context—have taken associative learning theory beyond mere concatenations of mental states without clearly implementing any classical logical structure (Gluck & Myers, 1993; Pearce, 1987). In short, if inferential "upgrading" strategies are to be applicable to animal judgments, they must move beyond old dichotomies and more clearly articulate when associative inferential connections become complex enough to merit inferential explanation, but without requiring that in doing so they must connect fully propositional contents.

Other upgrading accounts also fail to mark out the needed boundary; I briefly review three other popular alternatives. First, Boghossian (2014) simply stipulates that genuine inference begins at "System 1.5 and above", referencing the dual system/process literature in human psychology (Evans, 2008); but what exactly counts as System 1.5, and why the boundary should be located precisely there, is left unexplained. Second, ethicists and action theorists trying to mark out the lower bounds of inference tend to be more concerned with questions of responsibility and ownership than psychological explanation, and these issues may not be directly related. For example, in a recent cluster of philosophically astute and empirically informed papers, Levy argues that implicit attitudes in

---

2 Even staunch skeptics in comparative psychology concede that experiments on dozens of species have succeeded in this regard (Penn, Holyoak, & Povinelli, 2008).

3 One associative effect that recommended this revision to associative theory is known as 'blocking'. A standard blocking design has three phases: 1) the animal is exposed to a single cue, A, plus a reward (A+), 2) the animal is exposed to a compound cue plus reward (AB+), and finally tested on the second cue in isolation (B?). When blocking occurs, diminished association takes place between the second cue (B) and the reward. The effect is typically explained in terms of attention and prediction; since the animal already predicts the reward as a result of A alone, it does not attend to B as a possible target for a new association. Blocking has been observed in dozens of different species, but could not be accounted for by the behaviorist theories of the day.

adult humans are structured neither purely associatively nor purely propositionally, but rather fall into a sui generis middle ground of "patchy endorsements" that defy easy categorization on questions of ownership or responsibility (Levy, 2014, 2015, 2017). Yet even if Levy is right for the purposes of ethical evaluation, this may not tell us much about animals and infants, for they are not morally responsible for anything they do. Finally, Bermúdez' approach—which is closest in methodological aims to the present inquiry—marks the boundary at the presence of "protological" structures involving nonlinguistic analogues of logical operators like the negation and conditional. Yet, I argue below that a great many flexible, representation-driven judgments that lack even protological structure should count as inferences, so this still places the lower boundary too high.

In this paper, I explore an alternative, broadly Lockean strategy for articulating the lower bounds of inference. This strategy takes the relevant rationalizing connection to be intensional, but intuitive rather than formal in nature. To provide a historical contrast with present debates, classical empiricists like Locke and Hume saw no tension between association and inference; indeed, they argued that demonstrative reasoning was just a rarified form of association. Locke in particular heaped derision on the Classical Inferentialists of his day, the Scholastics who held that inference involved thoughts being fit to the forms of the Aristotelian syllogisms. The problem is that flexible thought precedes learning these forms, which are mastered only with effort. As Locke put it in the *Essay*:

> "God has not been so sparing to men to make them barely two-legged creatures, and left it to Aristotle to make them rational . . .Rather, he has given them. . .a native faculty to perceive the coherence or incoherence of its ideas, and can range them right, without any such perplexing repetitions." (Locke, 1975, IV.17.4)

Many later thinkers found this strategy unintelligible; when evaluating Lowe's (1993) Lockean suggestions, Bermúdez worries that "it is hard for us to make much sense of what Locke and his contemporaries found so easy to take for granted, namely, that we can directly perceive connections between thoughts" (2003, 110). I here answer this charge by offering a psychologically-plausible model of the relevant faculty of nonlinguistic judgment. On this model, nonlinguistic inferences can be causally sensitive to particular modes of presentation of an agent's mental states, though the details of the rationalizing connection may be opaque to introspection. In short, this model rejects CI and finesses Taking—the idea being that nonlinguistic practical inference can be understood as a form of learned, similarity-based categorization judgment, operating over complex, configural associative structures, in which one behavioral option is selected, on intensionally-sensitive grounds, as matching the type of action which is desired.

A difficulty that such "intuitive" models of inference must overcome is that the way in which these decisions are intensionally sensitive is opaque to the agent's introspection: the particular features assessed may not be directly available for reflection and scrutiny. Below, I argue that this concern can be overcome by noting that these processes may be accompanied by epistemic feelings (such as hunches and feelings of knowing) that are available for introspection, enabling a greater degree of executive control than is typically acknowledged for such judgments in nonlinguistic agents. Since agents can choose to follow hunches or not, and compare different hunches according to their analogue strengths, these feelings provide the agent with the minimally required level of epistemic control.

Three more caveats before moving on: first, readers may question the exclusive focus on practical inference here; if animals and infants are capable of rational practical inferences, shouldn't they be capable of rational theoretical inferences as well? The short answer is that yes, they likely are; but since the outcomes of nonlinguistic theoretical inferences are neither directly observable nor reportable, empirical evidence for them will be even more tenuous and indirect. My strategy is thus to first mark out the base case of nonlinguistic practical inference as it manifests itself in these empirical literatures, and approach the question of theoretical inference later. Second, I will here focus on very simple practical inferences made in contexts where the environment actively presents behavioral options for consideration, rather than more complex instances of "planning" or "insight" that require agents to actively work out what they might do. Again, my goal here is only to mark the lowest bounds of practical inference, with the hope that more sophisticated forms of inference can be bootstrapped up from this base case by adding additional components. The bootstrapping project is itself a substantial task, and so tackled in another essay (Buckner, 2017). Third, I will not speculate here as to whether this model is ultimately the correct one for the debate over adult human implicit attitudes or intuitive judgment. Though some lessons will no doubt be applicable, it is possible that language fundamentally rewires even the tacit/implicit part of cognition in adult humans (Anderson, 2010; Bermúdez, 2005; Clark, 1998), so that opaque judgments in humans implicate logical structure just as Mandelbaum suggests. Early exposure to language in infants may also complicate the case, as rewiring may begin earlier than explicit language use. Therefore, I will focus exclusively on nonhuman animals in what follows.

2. Putative Examples on Nonlinguistic Inference

Many authors take for granted a tight relationship between inference and logical entailment, so it may help to open with some candidate non-logical inferences in order to loosen those intuitions. Though my eventual target is even more basic kinds of judgment, map-based reasoning may serve as an instructive gateway example (Camp, 2007; Rescorla, 2009b). Maps are complex, structured representations that can ground judgments about unobserved spatial relationships and guide novel navigation behaviors without relying on logical structure.[4] Let us consider an example from Camp (2007, 161-162), who has argued persuasively that maps offer a kind of genuine inference that does not depend on natural language. Suppose I have seen Alice at the café, Bob at the grocery store, and Ted at the park, and need to know their locations relative to one another. In a clear sense, if I can place their locations with respect to one another on a two-dimensional map, then I can reason about their relative locations in a way which goes beyond anything I observed directly in my personal experience. This works in large part because much of the inferential work has been done in the construction of the map itself, the structure of which—if the agent has built it correctly—will be isomorphic to these individuals and locations' relative positions in the external world. Merely by surveying the map (Figure 1), for example, I can reason that Alice and Ted are north of

[4] Camp (2007) argues that a map-based representational system could instantiate a "weak" language of thought by satisfying some associated properties, such as Evans' Generality Constraint, and there is live debate as to how much logical structure could potentially be captured by maps (Rescorla, 2009c). Either way, the point stands that maps can ground substantial inferences even when they do not replicate logical structure.

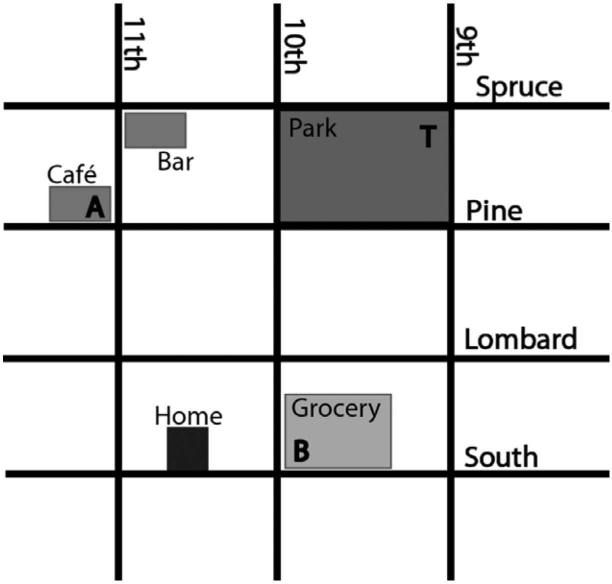

Figure 1. Map reproduced from Camp (2007, 162).

Bob, or that Bob and Ted are east of Alice; and if I were at home and needed to pick up all three in the car, I could chart a novel, efficient course to do so, using the map to guide me along my way. In fact, such novel, map-based shortcutting has long been considered one of the best places to look for sophisticated cognition in nonlinguistic agents, and research has demonstrated spatial integration and novel route-finding in animals as diverse as chimpanzees, rats, corvids, and even honeybees (Gibson & Kamil, 2001; Menzel et al., 2005; O'Keefe & Nadel, 1978; Tolman, 1948).[5]

It is hard to deny that something relevant to Taking has occurred here, especially when such maps are represented entirely "in the head". An agent using the map can obtain a variety of new conclusions about previously unknown spatial relationships by taking the structure of the map to correspond to the structure of the world. If the map were taken in a different way, for example—if the map were reconsidered with north facing down, or if icons on the map were moved to different locations—then the agent would take the map to support different conclusions. Just as one takes the logical structure of propositions to support logical inference in classical deduction, the spatial structure of the map supports spatial inference in map-based reasoning. The latter kind of reasoning, however, is much less plausibly rule-like; while one could formulate rules for map-based inference, typically agents just "see" the right answer by considering the map in the relevant ways.

If we are willing to consider map-based judgments to be inferential, then the two relevant features about how the map is taken—1) the matching between relations in the map and relations in the world that 2) allows the agent to reliably arrive at the correct solutions to novel problems—generalizes to even more basic forms of representation. I

[5] Note that there is an ongoing debate as to whether it is correct to say that animals have maps in their heads or can understand external maps (Bennett, 1996; Gould, 2002); but since I intend to focus on even simpler judgments, we need not enter that debate here.

will focus particularly here on configural learning, an advanced form of associative learning that occurs paradigmatically in cases like facial recognition. Configural learning is associative in the sense that the acquisition of the patterns depends upon statistical tracking of spatiotemporal contiguities between cues and that cue configurations lack explicit logical form, but is more advanced than basic associative learning in that multiple cue associations can occur in parallel and can be relational and context-sensitive in nature.[6] Specifically, the associative influence of any particular feature can depend upon the value of many other features in spatial, temporal, or conceptual context. For example, if I show you a picture of a friend's face and then another picture with the same features scrambled in random locations, then you may recognize your friend in the first presentation but not in the second. Many other mundane types of tasks can be understood as configural learning problems; for example, it is the obvious way that we can easily distinguish a tree from a pile of kindling, which possesses all the same features arranged in a different configuration. This kind of configuration cannot be reduced to a mere concatenation of stimulus generalizations along individually conditioned stimulus-response links.

Many animals are capable of this kind of configural judgment specifically with faces, including chimpanzees, pigeons, horses, and sheep, who can distinguish hundreds of distinct individuals in their flocks by not only subtle facial cues, but even multi-modally by integrating visual cues with olfactory information and auditory calls (Huber et al., 2013; Kendrick et al., 2001). Configural learning extends to a variety of other kinds of task, however; for example, pigeons can be trained with multiple exemplars to discriminate abstract geometric categories (Figure 2), which requires not just recognizing three lines and three angles, but also that those features are arranged in the right way (Watanabe, 2006). Pigeons have even been trained to discriminate Impressionist artwork from Japanese paintings in a way that generalizes to novel paintings, and to judge "good" and "bad" paintings composed by children (Watanabe, 2010). Even bees have demonstrated configural learning, correctly discriminating different flower and facial patterns on

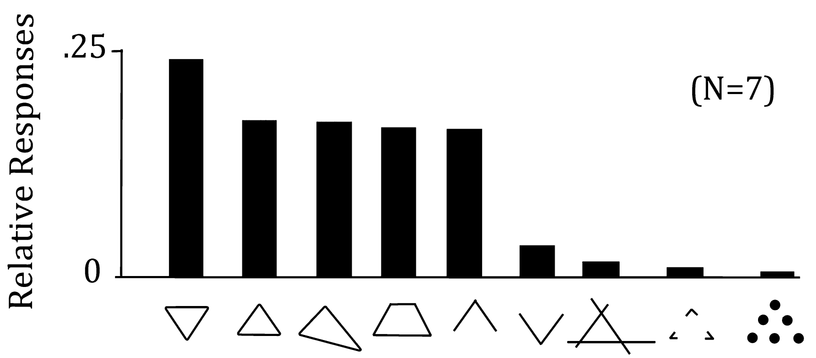

Figure 2. Pattern of responding to triangle-like stimuli after learning from multiple exemplars, reproduced from Watanabe (2005, 620).

[6] There are several different grades of configurality—merely the ability to form compound stimuli that can serve as the subject of association, detecting first-order relations amongst multiple stimuli (i.e. to eyes above a nose and a mouth), holistic processing which glues features into a gestalt, or second-order relations amongst stimuli (i.e. spacing amongst features)—see Maurer et al. (2002).

repeated discrimination tasks (though they can only remember a few configurations at a time—Avarguès-Weber et al., 2010). In fact, cognitive mapping may be a special case of configural learning, one that involves coordinating two different kinds of configurations —egocentric perspectives of nearby cues and spatial orientation information from distant landmarks—in the hippocampus (O'Keefe & Nadel 1978).

To evaluate the relevance of configural learning to inference, I suggest that instead of starting with classical logic or decision theory and revising downward—by trying to extract protological structure from configurations or idealize these judgements as somehow ranking disparate options in an abstract, fungible currency like utility—we should try a more biologically-inspired, bottom-up approach. Specifically, we should focus on the kinds of practical decisions animals make on a day-to-day basis. These decisions have a different structure than a typical deductive or decision-theoretic task: the environment presents multiple options of the same target category for evaluation—such as prey, predators, or mates—and the task of the agent is to quickly categorize these individuals and/or select a good exemplar.

While this kind of task may sound simple, in nature there is nothing easy about these judgments. The stakes here are often high; there can be many possible combinations of cues across different sensory modalities that need to be configured and considered in context; and yet the decision may need to be made quickly in a dynamic situation. Competition, moreover, will induce a pressure for individuals to learn and explore diverse strategies, so the individual's particular learning history and distinctive take on these judgments may be crucial. Where decision-making is collaborative—as is common in animals that live in social groups—there may even be an epistemic division of labor involving deference to the most skilled categorizers. This kind of specialized configural selection is the kind of judgment where we should expect to see a Taking-like inferential sensitivity take shape, where the individual's own epistemic character begins to guide decision-making.

To take one illustrative example, it used to be common knowledge in Africa that lions did not consider giraffes to be viable prey. Giraffes have strong legs to support their tall and unwieldly bodies, and have an enviable height from which to deliver skull-crushing kicks to any lioness which gets too close. However, a population of lions in South Africa's Selous Game Reserve have learned that if giraffes are encountered in a sandy river bed—where their hooves can get stuck, causing them to trip—then the lioness can outmaneuver a giraffe's attempts to kick and go in for the kill.[7] In short, these lions have learned from their individual experience to flexibly interpret an ambiguous cue—that a particular configuration of giraffe appearances is deadly in one context, but can be safely pursued as prey in another. Other lions without this experience would not pursue these giraffes as viable prey, even if they were encountered on the sandy river bed. The strategy, moreover, appears to be picked up by the young when they begin to engage in hunts with the pride. This community of lions has come to see giraffes in a way that others have not, and as a result act more successfully.

To consider another example, elephants form matriarchal groups, where herd-mates follow the lead of one of the eldest females, especially when evaluating threats (McComb

---

[7] Unpublished observations from Dr. Dennis Ikanda, Tanzania Wildlife Research Institute (though for videos of these encounters, see *Nature Shock*, Episode 2, "Lions vs. Giraffes" http://www.smithsonianchannel.com/shows/nature-shock/lions-vs-giraffes/1003142/3383147).

et al 2001). When approached by a possible predator or an out-herd conspecific, the elephants need to quickly decide whether to go about their business or interrupt grazing to gather up the young in a "bunching" reaction, a defensive circle where the older adults protect the young and prepare to charge onlookers. Deference to the eldest females appears to be highly adaptive in guiding these defensive reactions, for they have decades of experience identifying different classes of threats. Not only can these matriarchs distinguish hundreds of different individual elephants and categorize their threat level by distant calls, olfactory cues, and physical features, but matriarch age is highly correlated with the likelihood that their troops would selectively bunch in response to unfamiliar intruders. The eldest matriarchs are also more likely to seek out additional cues—for example, by engaging in exploratory smelling—when calls were recognized as unfamiliar, perhaps demonstrating some degree of active epistemic control. Playback experiments—where researchers cleverly position speakers in Kenya's Amboseli National Park to test reactions to a variety of recorded human voices—have shown that experienced matriarchs can even determine the threat level of human intruders by categorizing them according to ethnicity, gender, and age by using acoustic cues, olfactory cues, and dress—for adult Maasai tribesmen sometimes spear elephants in competition for grazing or in retaliation for attacks against humans, whereas Kamba tribesmen, or women/children of either tribe, usually pose no threat (Bates et al., 2007; Shannon et al., 2013). These matriarchs thus play a critical epistemic role in their herds; and the fitness of the herd is greatly diminished when matriarchs are killed by poachers, who target them for their large tusks.

The picture of nonlinguistic practical judgment that I am offering here is that most animals have a relatively small set of motivational modes divided by basic biological needs and drives that serve as primitive reasons for acting (Hurley, 2003), reasons that can motivate a range of behavioral responses that can perhaps be flexibly chained. The animals we have just been discussing also possess an advanced associative learning system that allows them to change between modes, select from behaviors appropriate to that mode, and guide the performance of that behavior towards a particular option by matching configural patterns individually acquired from their environment. This integration of biological and psychological perspectives on animal behavior is aptly-captured in Timberlake's behavioral systems theory, which systematizes animal action in just these terms (Timberlake, 1994, and see Figure 3). Advanced forms of associative learning enhance the practical agency of animals by modulating and guiding natural appetitive responses towards functional ends—rather than appearing to de-rationalize behavior through automatic linking of arbitrary stimuli to functionless responses, an attitude encouraged by an exclusive focus on inflexible elemental associations and artificial laboratory environments. Configural patterns serve as individually-acquired gates and guides, constructed and monitored by active agents over time in response to environmental feedback to better direct natural motivated responses, which achieve their biological purposes when learning guides them well, and fail when learning guides them poorly (Allen & Bekoff, 1999).

This package forms an attractive set of views that can explain flexible and sophisticated animal behaviors in psychological and biological context. However, the classical inferentialist may still be shrugging her shoulders and sighing, "Yes, yes, all very interesting—and perhaps I will even concede that these decision processes are too flexible to be explained in terms of reflex, instinct, or elemental association. But I still don't see

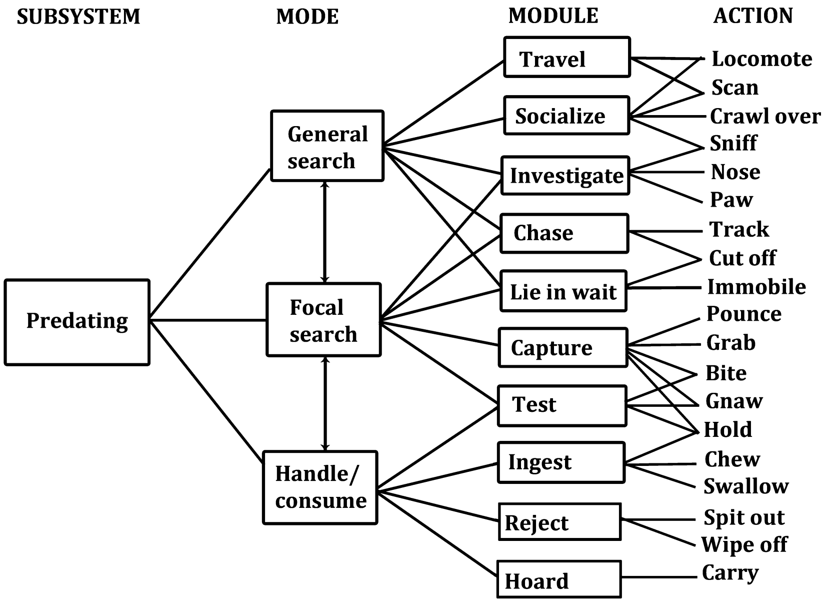

Fig 3. A decomposition of the rat's predation subsystem reproduced from Timberlake (1994, 408).

anything like rules or logical form; so why should we think of them as *inferences*?" To determine how much of this lingering skepticism can be addressed, we shall have to return to the considerations that brought classical inferentialists to rules and logical form in the first place.

3. Taking and Classical Inferentialism

Let us thus reflect on the long tradition of classical inferentialism. Skeptics of nonlinguistic rationality, from Aristotle to Davidson, have thought reason-based explanation to depend upon a formal inferential connection linking the performed action to an agent's epistemic and motivational states. On this picture, the rationalizing connection takes the form of a syllogism, deductive rule, or decision-theoretic calculation. CI is aptly summarized in an oft-quoted passage from Davidson, who was influential in laying the view's contemporary foundations:

> If someone acts with an intention then he must have attitudes and beliefs from which, had he been aware of them and had he the time, he could have reasoned that his act was desirable…If we can characterize the reasoning that would serve we will, in effect, have described the logical relations between descriptions of beliefs and desires and the description of the action, when the former gives the reasons with which the latter was performed. We are to imagine, then, that the agent's beliefs and desires provide him with the premises of an argument. (Davidson, 1978, 85-86)

In short, CI asks us to imagine the agent as considering a course of action and inferring, using the rules of logic, decision theory, or other formal method, that performing the action will satisfy a desire.

While for Davidson the attributions of these sentences and deductions remained an interpretive, virtual, and possibly anomalous matter—what Glock (2000) calls a "holophrastic" approach—the picture acquired new relevance to cognitive science with the shift towards realism recommended by the computational theory of mind. This theory held that the brain is a mechanism that can physically implement these rules in a causally realist fashion—it really possesses symbols with causal structures corresponding to the syntactic form of these sentences, and really possesses inferential mechanisms that could reliably, physically perform the relevant operations in virtue of their form. Fodor (1975) in particular linked this view with the Language of Thought thesis, which holds that all thought occurs in a compositional, systematic, language-like medium of mental representations.[8] The attraction of this computationalist package should not be underestimated, for it offered an interlocking set of solutions to the problems of representation, mental causation, and rationality.

Though both the interpretive and realist forms of lingualism about inference may still have much to offer when it comes to explicit adult human cognition, especially forms of thought that involve inner speech—(Carruthers, 2011), though see also (Langland-Hassan, 2014)—they unfortunately can no longer be regarded as empirically plausible when it comes to nonlinguistic animals. While animals may often judge in ways consistent with classical logic and decision theory in limited domains—a mode of thinking that Brunswik (1955) dubbed "ratiomorphic"—they lack the mastery of syncategorematic and mathematical relations required to apply such rules in a domain-general way (Hurley, 2003; Watanabe & Huber, 2006). Moreover, it is unlikely that any animals are capable of explicit metarepresentational awareness of the contents of their mental states or the inferential connections between them. In particular, they cannot conceptualize these feelings in terms of concepts like *truth-preservation* or *likelihood,* which would be required to understand why these rules recommend a particular course of action—leaving it unclear how they could satisfy Taking on the classical paradigm (Boghossian, 2014). Thus, the application of either a literal or virtual linguistic medium for nonlinguistic judgment appears inappropriate.

An important insight about inference emerges from the role played by Taking in these views, however: that intensionality is central to the distinction between rational inferences and non-rational forms of behavioral causation. Davidson stresses that the rationalizing connection does not hold between an agent's reasons directly, but rather between the agent's own descriptions of those beliefs, desires, and actions. For a rationalizing explanation of behavior to identify real causes of behavior, it cannot be phrased in terms of what the agent could or should have recognized. The agent must have actually performed the action because of those particular contents, as those contents are presented in its own psychology. This is in part what sent so many theorists searching for language-like

[8] For example, consider Fodor's (1975, pp28-29, paraphrased by Bermúdez, 2003, 134) model of practical inference:

i. A given creature finds itself in a certain situation S.

ii. It believes that a certain set of behavioral options, $B_1 \ldots B_n$, is available in S.

iii. The creature predicts the probable consequences of performing each of those behavioral options by computing a series of conditionals of the form: If B_i is performed in S then consequences C_i will occur with a certain probability.

iv. A preference ordering is assigned to the consequences.

v. The creature's choice of behavior is determined as a function of the assigned preferences and probabilities.

representations, because the obvious candidate for such individualized, causally potent modes of presentation are literal, sentential descriptions.[9]

It is worth remembering at this point that there are venerable reasons—related to Lewis Carroll's story of Achilles and the tortoise (Carroll, 1895)—why we should hesitate to require that the agent's understanding of the fit between reasons and action take the form of an explicit metarepresentational belief or further rule, at least in the base case of inference. The worry is that if the agent's awareness of the inference's correctness took the form of a second-order belief or rule, then the agent would simply have another inference to draw that requires additional justification. Specifically, we should need some third-order belief or rule connecting the second-order one to the original inference, and so on to vicious regress. The moral that many draw from this worry is that an inferential agent's understanding of the fit between its reasons and action must ultimately bottom out in "a certain complicated disposition or competence or practical capacity" (Stroud, 1979), and a capacity for assessing configural fit might be just the right kind of disposition for this job. There are equally famous rebuttals to this worry, so I will not press the advantage here; though I will return to this point in the final section, where I review empirical work showing that many animals seem to possess awareness of the fit between their reasons and actions.

While I agree that some form of Taking is a non-negotiable component of rationalizing explanation, I argue below that procedures operating on other non-language-like representations like configural patterns are sufficiently intensional to satisfy this demand.[10] Before saying precisely how, though, we must first open up some conceptual space for a different type of practical inference that could operate on such representations in a way that respects their distinctive modes of presentation.

4. Practical Inference as Categorization

The possibility of alternative models of practical inference can be located in yet earlier foundations of Davidson's own approach. Davidson begins his landmark article "Actions, Reasons, and Causes" by noting that

> . . .a reason rationalizes an action only if it leads us to see something the agent saw, or thought he saw, in his action—some feature, consequence, or aspect of the action the agent wanted, desired, prized, held dear, thought dutiful, beneficial, obligatory, or agreeable. (1963, 685)

Davidson obtains his famous thesis that reasons are causes by following Aristotle in positing these pro-attitudes as psychologically causal forces. The key transition in practical inference is thus a directing of this causally-efficacious pro-attitude towards a particular behavioral option.

> Davidson unpacks this transition by holding that. . .

---

[9] For more on this assumption and alternatives, see Glock (2000).

[10] There has been intense debate as to whether similarity-based concepts like prototypes can be considered concepts in the classical sense—mostly focused on whether they are compositional or not (Fodor & Lepore, 1996; Prinz, 2002)—but since two prototypes could pick out the same extension via different features, even if non-compositional they would still be intensional in the relevant sense.

> ...whenever someone does something for a reason, therefore, he can be characterized as (a) having some sort of *pro attitude toward actions of a certain kind* and (b) believing...that his action *is of that kind*. (1963, 685, my emphasis)[11]

Thus the task in question can be understood as identifying the action to-be-performed as an instance of a type of action towards which the agent has a pro-attitude. Practical inference can then be understood as a categorization judgment; and once the judgment is so broadly construed, models of inference and rationality other than those provided by CI become available for discussion.

A distinction is commonly drawn in psychology between rule-based and similarity-based categorization (Juslin & Olsson, 2004). The operations of classical logic and decision theory are fine examples of rule-based categorization procedures, assuming we have formal descriptions of the conditions of membership in the target category, such as necessary and sufficient conditions for classical logic and/or probability estimates or units of utility currency for decision theory. But similarity-based categorization heuristics are also eligible. A difficulty with this approach, as Goodman reminded us, is that if we are too liberal in tallying similarities, then everything is similar to everything else in an infinite number of ways. Our ability to apply reason-based explanation to such judgments will thus depend upon our ability to discern, for the task, which similarities count as relevant; but this is precisely where the agent's individualized modes of presentation become paramount. Animals with different configural patterns for the same goal-categories (i.e. predator or prey) may categorize the same exemplar differently, because they will focus on different similarities.[12]

Because such similarity judgments can be opaque to introspection and structurally more complex than logical deductions—in the sense that facial recognition is a harder problem in artificial intelligence than playing the challenging but neatly-defined game of chess—identifying the relevant similarities will require us to do more empirical homework. First, we should bring to bear our models of the creature's attentional demands, perceptual abilities, working memory constraints, and representational capabilities to

---

11 A similar, optional transition from categorization to decision theory can also be found in Fodor, who explains the explanatory payoff of his decision-theoretic model of practical inference in the following way:
"The point to notice is that it is built into this pattern of explanation that agents sometimes take their behavior to be behavior of a certain kind; in the present case, it is part of the explanation of *a*'s behavior that he believed it to be of the...kind for which highly valued consequences are predicted" (1975).

12 Compare Bermúdez (2003, 95): "The essence of perception under a particular mode of presentation comes because different nonlinguistic creatures will perceive different types of similarity between [the same objects]." Importantly, this kind of intensional sensitivity can hold even if the configural assessment lacks what Bermúdez calls "protological structure", such as the forms of protodisjunction, protonegation, and protoconditional that he thinks minimally necessary for reasoning (2003, 140-147). For example, Bermúdez suggests that "protoconditional" reasoning based on a causal belief could support a nonlinguistic form of reasoning akin to *modus ponens*, and it might be tempting to analyze the lioness/giraffe example discussed above along these lines. However, it is not clear that the lioness need possess anything like a causal belief in this case, that approaching a giraffe in the grassland would cause it to be kicked, but approaching in the riverbed would cause it to obtain a meal. Minimally required for such a causal belief is some sense in which the animal distinguishes causation from mere correlation, and it is not clear that the lioness needs to draw any such distinction. Some animals appear sensitive to this difference (Blaisdell, Sawa, Leising, & Waldmann, 2006; Taylor, Miller, & Gray, 2012), and the addition of protological structure to configural judgments may bootstrap them to more advanced forms of reasoning; but these capacities go beyond the basic forms of configural pattern-matching I have discussed above.

develop a model of the sorts of regularities the creature is capable of perceiving, learning, and applying to decision problems. Second, we need to carefully scrutinize our normative analysis of the cognitive task the creature is performing; while we can describe my attempt to choose a cereal in the grocery store in decision-theoretic terms as an attempt to select the optimal cereal in terms of global utility, it is more accurately described as an attempt to find a merely satisfactory instance of *healthy cereal* in a crowded and misleading aisle under realistic time constraints. Thirdly, we need to understand the informational structure of the environment in which the animal makes the decision to know what sort of similarities it will have had the opportunity to learn and will have reliably enough led to success. With these pieces in place, we can determine the features that the agent could and should evaluate to reliably succeed on a particular categorization task.

While similarity-based models of inference might have been considered dead in the water several decades ago as a result of Goodman's critiques (Goodman, 1972), there is no longer good reason for this default skepticism. Since Goodman's day, psychology and philosophy have developed a variety of empirically-supported, formal models of similarity, including Tversky's influential feature-matching model, Gärdenfors new geometric model, and popular mathematical models of prototype and exemplar-based categorization (Ashby & Maddox, 1993; Decock & Douven, 2011; Gärdenfors, 2004; Tversky, 1977). However, a lingering worry is that similarity judgments are not assessable by norms, and so could not support rational inferences. I conclude this section by considering and rebutting some specific forms of this worry.

One of Goodman's concerns here is that similarity assessment varies too much by context to be subject to normative assessment. Granted, similarity varies by context in a way that deductive validity judgments do not. However, many other straightforward normative evaluations vary by context—what counts as good basketball player for a neighborhood pickup game might not cut it in the Olympics—so it cannot be context-sensitivity alone that renders similarity assessment ineligible for normative evaluation.

A more targeted worry about context-sensitivity might be that similarity judgments could not be stably sensitive to the intensions of representations on which they act; but this is no problem if the representations themselves have context-sensitive intensions. First, consider that there is a straightforward intensional/extensional distinction to be made for human similarity-based categories like prototypes—different lists of prototypical features can extensionally pick out the same list of exemplars. Moreover, the intensions or modes of presentation for configurations or prototypes often vary by context—just as the lion's classification of a giraffe as prey might vary depending on whether the ground is rocky or sandy, our idea of the prototypical police officer for rural Texas is quite different from our idea of the prototypical police officer in Trafalgar Square. This observation importantly implies a concession to particularism about practical judgment—what counts as relevantly similar will depend upon the category's intension in that particular context—but it is not clear why this is inappropriate, especially for the kinds of nonlinguistic judgments we are considering.[13] The Selous lions get a meal while avoiding concussions precisely because their judgments adapt to context.

The final objection we might consider in this line is that such particularist judgments are not assessable by norms: if the criteria for having made a good judgment are different

[13] In fact, these categorization procedures often outperform classical alternatives precisely because they more flexibly adapt to subtle changes of context (Gigerenzer & Brighton, 2009).

for every situation, how can they be assessed for rationality? Here, we need only note that there are ready and appropriate evaluative criteria, but they are ecological rather than rule-based in nature. We should here simply accept the standards of ecological rationality recommended by empirical researchers who study these intuitive judgments. Specifically, these judgments will be ecologically rational if the features used to assess target category membership are highly valid in that context, where 'validity' measures the conditional probability that an exemplar falls under the target category given the cues assessed (Gigerenzer, Todd, & ABC Research Group, 1999). For any given categorization problem, the goal needs to remain stable, but the "right" features to assess can and should adapt to context.

From this perspective, the evaluative component of practical rationality concerns the "bet" that agents make about the informational structure of their environment, and the bet is justified if it achieves a good fit with environmental circumstance. These bets will typically be informed by a structured processing of statistical information over time; however, we must give up on the idea that one can be guaranteed to act rationally just by "following the rules". In fact, heuristics researchers have demonstrated that classical approaches prizing coherence norms like truth-preservation or utility optimization will be routinely outcompeted by frugal, ecological strategies that are more economical in their collection of evidence and avoid the dangers of optimization, especially in common environments that are "non-compensatory" (in which a more valid cue cannot be outweighed by the less valid cues even if all the latter disagree with the former) or exhibit high noise that optimizing strategies overfit (Gigerenzer & Brighton, 2009; Gigerenzer et al., 1999). As these researchers have argued, these are just the kinds of environments in which animals often live—in large part because animals broadcast highly-valid cues (like colorations, antlers, or behaviors such as roars or stotting) that save each other the sometimes expensive exploration of other cues through chases or fights.

To take stock, we are thus left with the following new package: intuitive practical inferences are a) categorization procedures that are b) causally sensitive in well-understood ways to the context-sensitive intensions of the representational states that drive them and c) epistemically rational according to their ecological fit to the informational structure of the agent's environment. We thus here reject CI entirely, but finesse Taking by noting the intensional sensitivity of these judgments for the purposes of psychological explanation, allowing us to draw a firm border between such similarity judgments on the one hand and the intensionally-insensitive causes of behavior like headaches, lexical priming, and "mere concatenation" associations on the other. In slogan form, inference begins at configural conception.

It is worth noting that this view comprises a novel hybrid of traditional internalism and externalism: judgments are internally rational given an intensional match between a target category and an exemplar, but externally rational according to an ecological match to informational context. The former secures the explanatory rationality of these judgments for the purposes of psychological explanation and Taking—they tell us what the agent saw in its action—while the latter determines whether the judgment made an epistemically rational use of available information. Too often these two questions have been conflated; but dealing with each separately, I suggest, is the key to understanding the role of psychological explanation in cognitive ethology and comparative psychology.

Before moving on, it is worth demonstrating that this account of categorization-based practical inference is not so liberal that it cannot draw principled exclusions.

Interestingly, not all categorization heuristics studied by psychologists of intuitive judgment will satisfy the intensional interpretation of Taking offered here. For example, a favorite of the Simple Heuristics group is the recognition heuristic, according to which agents simply choose the candidates they recognize (Gigerenzer, Todd, and the ABC Research Group, 1999). Though Gigerenzer and colleagues have stridently defended the ecological rationality of this heuristic—given that recognition is a highly valid predictor of category membership in many environments and can be computed automatically and efficiently—it is not rational in the sense relevant to Taking. The reason is that the attribute assessed (simple recognition) is causally insensitive and semantically unrelated to the mode of presentation of the agent's goal category; it simply does not matter what the subjects think of the target category, they will just choose the option they recognize. Thus there is good reason to think that the current perspective offers principled and powerful means of inclusion and exclusion, providing a novel way to taxonomize inferential heuristics and understand their relevance to philosophy of mind and action.

5. Objections

I conclude by considering three tempting objections. I will first review all three, and then argue that they invite a shared rebuttal—that recent work on nonlinguistic inference demonstrates that configural judgment in many animals is under a form of executive control that is sensitive to an animal's takes on its reasons for action.

Objection 1: Configural categorization is perceptual, not inferential

A common skeptical reaction to what has come thus far might be that these associative categorization judgments are perceptual rather than inferential in nature; that if subjects somehow just "see" the right answer, then these judgments cannot be inferences. This option would only be attractive to those who suppose perception and inference to be mutually exclusive, and thus is not for more expansive neo-empiricists who take both inference and reason to be perceptual in nature (e.g. Prinz, 2002). However, a number of theorists have recommended just this kind of perceptual strategy to tackle this class of opaque judgments (see Bermúdez 2003, Ch3 for review and criticism), so we should discuss it here.

Objection 2: Configural categorization does not allow comparisons between options

Even if all the above has been granted, it may yet be objected that a key motivation for Taking has been not yet been satisfied: that rational decision procedures must allow agents to compare different options in an intensional way. Bermúdez, for example, suggests that it is a "minimal requirement that the selection of a particular course of action. . .be made on consequence-sensitive grounds" which includes a "comparison of that outcome with the likely outcomes of other possible courses of action" (2003, 124). To illustrate such comparison, suppose that during my visit to the cereal aisle I narrow my options to a few final candidates, and choose amongst the remaining options by comparing their cost and selecting the least expensive option. While some sophisticated organisms like us may be able to explicitly rank options in this fungible, consequence-sensitive manner, non-linguistic agents using the intuitive judgments I describe above admittedly cannot. We should thus explore whether animals in fact have another way of ranking multiple options in decision-making.

Objection 3: Animals lack inferential control over their categorizations

Some writers will still be dissatisfied with the permissive, intensional interpretation of Taking offered above, because they believe a more significant form of awareness is required if we are to be held epistemically responsible for our practical inferences. For example, upon considering actions performed on the basis of an intuitive appraisal of another's emotional state computed from subtle facial and behavioral cues—a process much like the configural judgments highlighted here—Valaris (2016) writes that:

> Reasoning is an expression of the sort of control that we have over our cognitive lives; it is our way of making up our own minds. One way to bring this fact out is by noting that it has distinctive normative import: if you make a bad inference, we can legitimately criticize you as having been hasty, irresponsible, biased, and so on. By contrast, there is only a very thin sense in which the subject herself is responsible for her immediate judgments about another's emotional state. If you misread another's facial expressions, your mistake is more akin to a perceptual illusion than a case of bad reasoning. A natural way to explain this difference is to say that inferring reflects the subject's take on what her evidence requires, whereas our system for judging other people's emotional states is a lower-level system whose workings are opaque to us. Marking this difference is the point of the Taking Condition. (Valaris 2016, 4).

For these reasons, Valaris interprets Taking to require that subjects "be aware of the grounds on which [they] have made [judgments]". This restrictive interpretation of Taking would rule out the deflationary intensional construal I have recommended above, so we should confront it below.

Rebuttal: Animals can exert control over their inferences that is guided by their takes

First, realistic expectations: the present view cannot convince all skeptics of nonlinguistic inference, especially those who consider it akin to an analytic truth that genuine inference requires logical structure or an explicit, metarepresentational awareness of one's reasons for acting. The best that could probably be said in response to such definitional skepticism is that I have here articulated an important natural kind of judgment—probably with the kind of homeostatic property-cluster structure that characterizes nearly any interesting psychological posit, whether modules, concepts, emotions, or cognition itself (Buckner, 2015; Fodor, 1983; Griffiths, 2004; Machery, 2005)—that merely shares many properties in common with "genuine" inferences.

The three objections just reviewed, however, do not resort to this definitional nuclear option, but rather appeal to important aspects of inferential control that animals could in principle possess. I thus grant that the properties highlighted by these objections—the abilities to inhibit, guide, and compare initial takes on a decision problem, components of judgment often studied in psychology under the heading of "executive control"—should be considered for inclusion in this inferential cluster. However, the empirical evidence strongly suggests that many nonlinguistic animal judgments do in fact possess these forms of executive control. If skeptics avoid definitional table-pounding, then, these three objections to nonlinguistic inference can be empirically overcome.

While I argued above for a separation of internalist explanatory rationality and externalist justificatory rationality, these objections may highlight a lingering dissatisfaction with this division that is hard to put into words. Perhaps the ideal achievement of a theory of rational inference would be to show how inference bridges the internal and the external—how through the process of rational inference, internal takes guide external success. We could thus understand dissatisfaction with pure internalism and pure externalism as an appreciation of their failures to complete this bridge. Purely internalist theories insist on following the rules, even when doing so systematically leads to ruin; and purely externalist theories heap praise on merely lucky agents, even if their internal judgments are disordered and mutually incoherent (Taylor, 2003). Probably no theory can complete this bridge in a way that satisfies the aims of both pure internalists and pure externalists. However, reflecting on the three objections above in the context of empirical work on nonlinguistic executive control shows that the present view supports a sturdy bridge indeed.

Beginning with Objection 1, let us start with the idea that configural categorization is a kind of automatic perceptual judgment, and see what must be added to accommodate the empirical results. First, we should note that these configural judgments could only be subsumed under perception if perception were enriched with the categorical, intensional structure we have been discussing above; to do otherwise is to simply mystify their workings without explanatory payoff. Whether we call these judgments perceptual or not, there would remain a significant sense, relevant to Taking, in which they would simply count as inferential, because if a different configuration of cues were associated with the same goal-category, it would lead the animal to different categorizations. What must be at issue is thus not the mere sensitivity of the judgment to intensions, but whether there is an additional agent-level awareness of this sensitivity that can be deployed to guide judgment.

Indeed, many animals do seem to possess an awareness that would be difficult to subsume under "automatic" perception. In adult humans, related kinds of similarity judgment—such as the "prototyping heuristics" long studied in the psychology of intuitive judgment, like representativeness (Kahneman & Frederick, 2002)—are typically accompanied by and acquire their motivational potency through "epistemic feelings," such as hunches and feelings of knowing, confidence, and uncertainty (Arango-Muñoz, 2014). When asked why we behaved a certain way when we act on intuition, we are often aware of being motivated by these feelings, even if we are dumbfounded when asked to explain why we feel as we do. Crucially, these feelings are not experienced in the format of any sensory modality, and so it would be difficult to classify them as just a further aspect of sensory perceptions. Notably, many animals perform much like humans do in experiments that elicit epistemic feelings, such as in "uncertainty monitoring" paradigms. When asked to make difficult discriminations on dot-pattern paradigms, for example, a variety of animals—dolphins, chimpanzees, monkeys, and rats, to name a few (Beran, 2012; Foote & Crystal, 2007)—make adaptive use of a "bail-out" option (which quickly provides them a new trial without the time penalty due an incorrect answer) on borderline trials where their discrimination reliability drops to chance. These response curves fit those of humans when they report deliberating over feelings of uncertainty when completing these tasks (Shea et al., 2014; Smith, Couchman, & Beran, 2014).

To be clear, by appealing to the animal uncertainty monitoring literature here, I am not disputing well-known arguments by Carruthers or Bermúdez that language is required for truly metarepresentational metacognition; I instead aim to establish only that animals have internally-generated signals that are sensitive to their takes. Even these skeptics concede that this research now implicates a form of executive control over decision-making that cannot be wholly subsumed under basic perception or associative learning (Carruthers, 2014). Here as elsewhere, simple associative explanations need to be ruled out, and some of the experimental designs that do so reveal much about the rich internal control animals possess over these forms of judgment.

In particular, for a time the most popular skeptical response to these experiments was to hold that the bail-out option serves as a third response that can acquire a mediated reinforcement value throughout the course of the experiment. Specifically, perhaps the animals could learn that they received better rewards by creating a third "middle" category of stimulus presentations between "sparse" and "dense", and treating what researchers supposed was the uncertainty response as the key for that middle category. To overcome this interpretation, Beran and colleagues re-did the original experiments on macaques by only delivering rewards in a lump sum at the end of four-trial blocks (Smith et al., 2014). This should have prevented the monkeys from associating that uncertainty response with its own reward value, for there was no trial-by-trial reward feedback received for each bailout choice. One of the macaques in this experiment—Murph—used the bailout option in an adaptive manner trial-by-trial, despite the fact that there was no reward signal present on each trial to tell him whether his answer was "correct". The authors argue that Murph's use of the bailout option on each trial must have thus been guided by some other, internally-generated signal.

Even more persuasively, Murph's response curves on the blocked trial tasks were not centered on objective reward value for each trial, which would be predicted by low-level associative models of choice. As they put it,

> "Murph clearly did not track the associative structure. . .his [uncertainty responses] uncoupled from *objective* difficulty over a broad range of performance accuracies. But his [uncertainty responses] were elegantly coupled to the task's *subjective* difficulty. . .across the whole range of the task. Subjective difficulty is nonassociative because it is independent from reward, punishment, or reinforcement history. The organism can respond uncertain notwithstanding all of those factors. Murph did, suggesting that he was assessing difficulty subjectively" (Smith, Couchman, & Beran, 2014, 120).

These data suggest that not only do some animals have a subjective take on the suitability of the option they are evaluating for their goal, they possess a subjective, internal signal regarding their confidence in this take that can be deployed to select amongst different options. There is no evidence that this signal carries information about the particular features assessed in each categorization judgment; but it does systematically vary in an adaptive fashion with the subjective strength of each take, which by hypothesis is generated by the strength of the similarity assessment to the agent's intension for that category.

While evidence for the existence of this subjective signal is strongest in this kind of blocked-reward uncertainty-monitoring task, once its existence is established, it increases the likelihood that other experimental results are also best explained as forms of take-guided executive control, including experiments on inhibition, goal maintenance, and cue-seeking behavior. For example, another popular executive control paradigm in humans is the Stroop task, where words for certain colors are written in a different color (i.e. the word "blue" written in green) and subjects are asked to indicate the color in which the word is written. There is a strong tendency for subjects to respond with the color word that has been written, and subjects must become aware of this tendency and inhibit it to succeed. Such inhibition has been shown to correlate with other forms of executive control, and subjects do poorly on this task when distracted, tired, or under cognitive load. Several versions of Stroop-like tasks have been conducted on chimpanzees and monkeys, and they show many of the same patterns of responding that in humans are taken as proxies for effortful inhibition (Beran, Washburn, & Rumbaugh, 2007). In light of the background findings, it is reasonable to conclude that these animals have detected a clash between two different takes on the same stimulus, and with cognitive effort can inhibit their initial take in favor of a more considered one.

Supposing that animals do have these forms of implicit, take-guided executive control, let us now move to evaluate Objection 2. Recall that the question here is whether agents can compare multiple options along goal-relevant and intensionally-sensitive criteria. It is doubtful that we should require an intensionally-sensitive comparison of each option with one another—that is, say, directly comparing a box of Cheerios to a box of Wheaties qua mutual instances of cereal, as opposed to independently evaluating each serially—but even if we do, such a requirement has likely been overintellectualized. In each case, the rational criteria that the comparison must satisfy should be constrained by biological and psychological context, which often—even in sophisticated language-users like us—precludes an exhaustive or fungible comparison of options. As we noted above, utility maximization and other forms of optimization are not only not required for adaptive action, they are typically a hindrance.

Moreover, the deeper question should be whether animals can do much of the work that this intensional comparison was supposed to do without explicit awareness of the features being assessed in the categorization judgment. The hunches and feelings that accompany intuitive judgments provide some semantically-sensitive information about relevant options—the strength of the internal signal is sensitive to the strength of the similarity match to individual's pattern for that goal-category. Even if it is only the sort that allows comparison up to the point of satisficing rather than optimality, these hunches can be sorted according to analogue strength which, insofar as the features that drive them are valid indicators of category membership, can allow agents to pass by or inhibit responses to dissatisfactory options in favor of options supporting a stronger take. While this form of comparison amongst multiple options will never rise to the level of the "optimal and exhaustive", it provides all that is needed for flexible, adaptive, intensionally-sensitive satisficing.

Finally, let us approach the most difficult Objection 3, which raises the question of epistemic responsibility; can animal intuitive categorizations be legitimately criticized as good or bad reasoning, as Valaris requires? We have already specified the epistemic criteria that should be used to evaluate these judgments, which appeal to principles of ecological validity rather than internal coherence. This kind of reasoning has the epistemic goods when it is efficient, reliable, and operates on ecologically-valid cues, and lacks them when it does not. Note that as interpreted here, this is not a raw, case-by-case consequentialism; a thinker could get lucky every time with a highly-invalid procedure without their inferences counting as ecologically rational.

What remains perhaps to be satisfied here is an "ought implies can" type principle; is a tacit sensitivity to the strength of a take sufficient to support the idea that the animals could have chosen otherwise if they had reasoned better? Of course, if what is sought here is libertarian epistemic freedom, none of the empirical evidence adduced above could suffice. We should have to remain agnostic on questions of epistemic responsibility, albeit perhaps with a strong supposition of analogy to human freedom grounded by similar behavior on similar tasks. However, I doubt Valaris should be interpreted as demanding anything so strong.

Instead, I suggest that at issue is the assumption that if the specific features assessed in a configural take are opaque to us, then the take cannot grant us control over any judgments caused by it. Against this assumption, a significant sense of control can be satisfied here by noting the range of counterfactual situations in which the agent could have done better. I suggest that the range of counterfactual situations relevant to this question are those where the agent applies a greater amount of executive resources to the decision —which could be operationalized in terms of empirically-assessable variables pertaining to cognitive load, attention, and mental effort. The empirical results suggest that when agents do apply more of these resources, they tend to perform better on these tasks by inhibiting riskier judgments and seeking out more information when uncertain (Beran, 2015; Davidson, Amso, Anderson, & Diamond, 2006; Smith, Coutinho, Church, & Beran, 2013). Recall the Stroop-like inhibition tasks, for example; the animals perform better on these tasks when they are able to marshal more executive resources to guide their judgments, and poorly when these resources are not brought to bear. In short, what might be successfully argued in the passage from Valaris above is that one cannot be held epistemically responsible for taking a rapid configuration of facial cues to suggest a particular emotional state; but contra Valaris' apparent conclusion, when executive resources are available, one does have control over whether one acts on that take.

The role of executive control in these rebuttals raises a final tricky question: if the kind of intuitive inference I have been articulating here is a natural kind, but executive control is not present in all animals capable of configural takes, then should the lower boundary of rational inference be placed at the advent of mere configural categorization, or rather at advent of executive control, which allows these takes to flexibly guide judgment? At issue may be the inclusion or exclusion of large swathes of the Animal kingdom, should configural takes be common but executive control rare. We cannot definitively settle this question now, because it is ultimately an open empirical question about the clustering of these capacities in nature. My take on the current state of the empirical literature is that we should lean towards the more inclusive demarcation, because we should expect executive control to come in degrees and be modulated by a great many general cognitive resources that we already know vary widely across the Animal kingdom, like working memory capacity, intelligence, and attention.

There have already been several times in this young literature that consensus has supposed executive control to be notably absent from some species, only to later find that we were simply not presenting animals with the right task to elicit it. For example, until recently it seemed likely that while chimpanzees and Old World monkeys like macaques excelled at executive control tasks, New World monkeys like capuchins routinely failed. It was thus suggested that executive control might have uniquely evolved in a common ancestor with the Old World monkey lineage, and from there passed down to us. Against this evolutionary hypothesis, however, researchers have since found evidence for the convergent evolution of executive control in distant species with a reputation for intelligence like corvids, as well as suggestive findings in distant cousins like rats. Moreover, the same lab that repeatedly delivered capuchin failures has recently produced some successes, and now hypothesizes that capuchins only eschewed the bailout options in earlier experiments because capuchins are natural daredevils with a high tolerance for risk. When the amount of risk was increased in the bailout trials by offering six categorization options instead of just two (to lower the likelihood of getting the right answer by guessing), capuchins suddenly began using the bailout options just like macaques, chimps, and humans (Beran et al., 2016). I thus suggest that while we should remain agnostic between these two more precise locations for the lower boundary, we should conclude that it dips well into nonlinguistic territory, perhaps much deeper than even optimistic researchers recently supposed.

6. Conclusion

This essay offered a novel model of nonlinguistic practical inference, based in the relevant empirical psychology and philosophical tradition. This model defends something that much philosophical orthodoxy has regarded as incoherent: that opaque inferences can be internally rational in the sense relevant to Taking. The key to this solution was to treat the questions of intensional fit and epistemic evaluation separately, offering an internalist solution to the former and an externalist solution to the latter. Rational inference, it was suggested, bridges these two forms of rationality by allowing intensional takes to guide adaptive decision-making—in short, it is the process whereby agents exert some control in placing their own ecological bets.

This model is offered in the hope of reinvigorating philosophical interest in the psychology of nonlinguistic and intuitive inference. For one, focusing on the Taking condition allows us to move beyond older, hopefully settled debates about the epistemic evaluation of intuitive judgment (e.g. the "Rationality Wars" of Samuels, Stich, & Bishop, 2002) and focus instead on how an agent's implicit and intuitive processing flexibly govern its decisions. For another, it suggests that empirical work on intuitive judgment may have overlooked important philosophical distinctions amongst different kinds of categorization procedures—namely between those that are inflexible and insensitive to an agent's learned intensions, and those which are flexible and can only be explained by appealing to the agent's internal modes of presentation and executive control. If we are careful to avoid falling back into old dichotomies and tendencies to overintellectualize, we may close a rift between mutually relevant areas of philosophical and empirical research that have walked separate paths for too long.[14]

References

Allen, C. and Bekoff, M. (1999). *Species of Mind: The Philosophy and Biology of Cognitive Ethology*. MIT Press.

Anderson, M. L. (2010). Neural reuse: a fundamental organizational principle of the brain. *The Behavioral and Brain Sciences*, 33(4), 245–66; discussion 266–313. https://doi.org/10.1017/s0140525x10000853

[14] This work has undergone many revisions and has received comments in several forms; I would like to thank in particular Colin Allen, Michael Beran, Susan Blake, Maria Botero, Michael Brownstein, Hayley Clatterbuck, Bryce Huebner, Mark Kaplan, Elanor Taylor, Jonathan Weinberg, two anonymous reviewers, and audiences at the Southern Society for Philosophy and Psychology, the Eastern American Philosophical Association, Rice University, and at the University of Houston. This work was also supported in part by a fellowship from the Alexander von Humboldt Foundation.

Andrews, K. (2014). *The animal mind: an introduction to the philosophy of Animal Cognition*. Routledge.

Arango-Muñoz, S. (2014). The nature of epistemic feelings. *Philosophical Psychology*, 27(2), 193–211.

Ashby, F. G. and Maddox, W. T. (1993). Relations between prototype, exemplar, and decision bound models of categorization. *Journal of Mathematical Psychology*, 37(3), 372–400.

Avarguès-Weber, A., Portelli, G., Benard, J., Dyer, A. and Giurfa, M. (2010). Configural processing enables discrimination and categorization of face-like stimuli in honeybees. *Journal of Experimental Biology*, 213(4), 593–601.

Banaji, M. R. and Gelman, S. A. (2013). *Navigating the social world: What infants, children, and other species can teach us*. Oxford University Press.

Bates, L. A., Sayialel, K. N., Njiraini, N. W., Moss, C. J., Poole, J. H. and Byrne, R. W. (2007). Elephants classify human ethnic groups by odor and garment color. *Current Biology*, 17(22), 1938–1942.

Bennett, A. T. (1996). Do animals have cognitive maps? *Journal of Experimental Biology*, 199(1), 219–224.

Beran, M. J. (2012). *Foundations of metacognition*. Oxford University Press.

—— (2015). Chimpanzee cognitive control. *Current Directions in Psychological Science*, 24(5), 352–357.

——, Perdue, B. M., Church, B. A. and Smith, J. D. (2016). Capuchin monkeys (Cebus apella) modulate their use of an uncertainty response depending on risk. *Journal of Experimental Psychology: Animal Learning and Cognition*, 42(1), 32.

——, Washburn, D. A. and Rumbaugh, D. M. (2007). A stroop-like effect in color-naming of color-word lexigrams by a chimpanzee (Pan troglodyte). *The Journal of General Psychology*, 134(2), 217–228.

Bermúdez, J. L. (2003). *Thinking without words*. Oxford University Press.

—— (2005). *Philosophy of Psychology: A Contemporary Introduction*. Psychology Press.

Blaisdell, A. P., Sawa, K., Leising, K. J. and Waldmann, M. R. (2006). Causal reasoning in rats. *Science*, 311(5763), 1020–1022.

Boghossian, P. (2014). What is inference? *Philosophical Studies*, 169(1), 1–18. https://doi.org/10.1007/s11098-012-9903-x

Bouton, M. E. and Swartzentruber, D. (1986). Analysis of the associative and occasion-setting properties of contexts participating in a Pavlovian discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*, 12(4), 333.

Brownstein, M. and Saul, J. (2016a). *Implicit Bias and Philosophy, Volume 1: Metaphysics and Epistemology*. Oxford University Press.

—— and —— (2016b). *Implicit Bias and Philosophy, Volume 2: Moral Responsibility, Structural Injustice, and Ethics*. Oxford University Press.

Brunswik, E. (1955). "Ratiomorphic" models of perception and thinking. *Acta Psychologica*, 11, 108–109. https://doi.org/10.1016/s0001-6918(55)80069-8

Buckner, C. (2011). Two approaches to the distinction between cognition and "mere association." *International Journal of Comparative Psychology*, 24(4).

—— (2015). A property cluster theory of cognition. *Philosophical Psychology*, 28(3), 307–336.

—— (2017). Scaffolding intuitive rationality. In Newen, A., de Bruin, L. and Gallagher, S., eds. *The Oxford Handbook of 4E Cognition*. Oxford: OUP.

Camp, E. (2007). Thinking with maps. *Philosophical Perspectives*, 21(1), 145–182.
Carroll, L. (1895). What the tortoise said to Achilles. *Mind*, 4(14), 278–280.
Carruthers, P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. OUP Oxford.
—— (2014). Two concepts of metacognition. *Journal of Comparative Psychology*, 128 (2), 138–139.
Clark, A. (1998). Magic words: How language augments human computation. *Language and Thought: Interdisciplinary Themes*, 162–183.
Clatterbuck, H. (2016). Darwin, Hume, Morgan, and the verae causae of psychology. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 60, 1–14. https://doi.org/10.1016/j.shpsc.2016.09.002
Davidson, D. (1963). Actions, Reasons, and Causes. *Journal of Philosophy*, 60(23), 685–700.
—— (1978). Intending. In Y. Yovel (Ed.), *Philosophy of History and Action* (pp. 41–60). Springer Netherlands. https://doi.org/10.1007/978-94-009-9365-5_5
Davidson, M. C., Amso, D., Anderson, L. C. and Diamond, A. (2006). Development of cognitive control and executive functions from 4 to 13 years: Evidence from manipulations of memory, inhibition, and task switching. *Neuropsychologia*, 44(11), 2037–2078.
Decock, L. and Douven, I. (2011). Similarity after Goodman. *Review of Philosophy and Psychology*, 2(1), 61–75.
Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annu. Rev. Psychol*., 59, 255–278.
Fodor, J. A. (1975). *The Language of Thought*. Harvard University Press.
—— (1983). *The modularity of mind: An essay on faculty psychology*. MIT press.
Fodor, J. and Lepore, E. (1996). The red herring and the pet fish: Why concepts still can't be prototypes. *Cognition*, 58(2), 253–270.
Foote, A. L. and Crystal, J. D. (2007). Metacognition in the rat. *Current Biology*, 17(6), 551–555.
Gärdenfors, P. (2004). *Conceptual spaces: The Geometry of Thought*. MIT Press.
Gendler, T. S. (2008). Alief and Belief. *The Journal of Philosophy*, 105(10), 634–663.
Gibson, B. M. and Kamil, A. C. (2001). Tests for Cognitive Mapping in Clark's Nutcrackers (nucifraga columbiana). *Journal of Comparative Psychology*, 115(4), 403–417. https://doi.org/10.1037/0735-7036.115.4.403
Gigerenzer, G. and Brighton, H. (2009). Homo Heuristicus : Why Biased Minds Make Better Inferences. *Cognitive Science*, 1, 107–143. https://doi.org/10.1111/j.1756-8765.2008.01006.x
——, Hertwig, R. and Pachur, T. (2011). *Heuristics: The Foundations of Adaptive Behavior* (1st ed.). New York, NY, USA: Oxford University Press, Inc.
—— and Todd, P. M., & ABC Research Group. (1999). *Simple heuristics that make us smart*. Oxford University Press.
Gilovich, T., Griffin, D. and Kahneman, D. (2002). *Heuristics and biases: The psychology of intuitive judgement*. Cambridge: Cambridge University Press.
Glock, H.-J. (2000). Animals, Thoughts And Concepts. *Synthese*, 123(1), 35–64. https://doi.org/10.1023/a:1005295521736
Gluck, M. A. and Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, 3(4), 491–516.

Goodman, N. (1972). Seven strictures on similarity. In *Problems and Projects* (pp. 437–447). Indianapolis: Bobbs-Merrill.

Gould, J. L. (2002). Can honey bees create cognitive maps. *The Cognitive Animal: Empirical and Theoretical Perspectives on Animal Cognition*, 41–45.

Griffiths, P. E. (2004). Emotions as Natural and Normative Kinds. *Philosophy of Science*, 71(5), 901–911. https://doi.org/10.1086/425944

Heyes, C. and Dickinson, A. (1990). The intentionality of animal action. *Mind & Language*, 5(1), 87–104.

Huber, L., Racca, A., Scaf, B., Virányi, Z. and Range, F. (2013). Discrimination of familiar human faces in dogs (Canis familiaris). *Learning and Motivation*, 44(4), 258–269.

Hurley, S. (2003). Animal Action in the Space of Reasons. *Mind & Language*, 18(3), 231–257. https://doi.org/10.1111/1468-0017.00223

Juslin, P. and Olsson, H. (2004). Note on the rationality of rule-based versus exemplar-based processing in human judgment. *Scandinavian Journal of Psychology*, 45(1), 37–47. https://doi.org/10.1111/j.1467-9450.2004.00376.x

Kahneman, D. and Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin and D. Kahneman (Eds.), *Heuristics and biases: The psychology of intuitive judgment* (pp. 49–81). New York, NY, US: Cambridge University Press. https://doi.org/10.1017/cbo9780511808098.004

Kendrick, K. M., da Costa, A. P., Leigh, A. E., Hinton, M. R. and Peirce, J. W. (2001). Sheep don't forget a face. *Nature*, 414(6860), 165–166.

Langland-Hassan, P. (2014). Inner Speech and Metacognition: In Search of a Connection. *Mind & Language*, 29(5), 511–533. https://doi.org/10.1111/mila.12064

Levy, N. (2014). Consciousness, implicit attitudes and moral responsibility. *Noûs*, 48(1), 21–40.

—— (2015). Neither Fish nor Fowl: Implicit Attitudes as Patchy Endorsements. *Noûs*, 49 (4), 800–823. https://doi.org/10.1111/nous.12074

—— (2017). Implicit Bias and Moral Responsibility: Probing the Data. *Philosophy and Phenomenological Research*, 94(1), 3–26.

Locke, J. (1975). *An essay concerning human understanding*. (1690), P.H. Nidditch (ed.), Oxford: Clarendon Press.

Lowe, E. J. (1993). Rationality, deduction and mental models. In K. I. Manktelow and D. E. Over (Eds.), *Rationality: Psychological and philosophical perspectives* (pp. 211–230). Florence, KY, US: Taylor & Frances/Routledge.

Machery, E. (2005). Concepts Are Not a Natural Kind*. *Philosophy of Science*, 72(3), 444–467. https://doi.org/10.1086/498473

Mandelbaum, E. (2013). Against alief. *Philosophical Studies*, 165(1), 197–211. https://doi.org/10.1007/s11098-012-9930-7

—— (2016). Attitude, inference, association: On the propositional structure of implicit bias. *Noûs*, 50(3), 629–658.

Maurer, D., Le Grand, R. and Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6(6), 255–260.

Menzel, R., Greggers, U., Smith, A., Berger, S., Brandt, R., Brunke, S., . . . Watzl, S. (2005). Honey bees navigate according to a map-like spatial memory. *Proceedings of* *the National Academy of Sciences of the United States of America*, 102(8), 3040–5. https://doi.org/10.1073/pnas.0408550102

Nagel, J. (2012). Gendler on alief. *Analysis*, 72(4), 774–788.

O'Keefe, J. and Nadel, L. (1978). *The hippocampus as a cognitive map*. Clarendon Press.

Pearce, J. M. (1987). A model for stimulus generalization in Pavlovian conditioning. *Psychological Review*, 94(1), 61.

Penn, D. C., Holyoak, K. J. and Povinelli, D. J. (2008). Darwin's mistake: explaining the discontinuity between human and nonhuman minds. *Behavioral and Brain Sciences*, 31(2), 109–130; discussion 130–178.

Prinz, J. J. (2002). *Furnishing the mind: concepts and their perceptual basis*. MIT Press.

Rescorla, M. (2009a). Chrysippus' dog as a case study in non-linguistic cognition. In Lurz, R. (Ed). *The Philosophy of Animal Minds* (pp 52–71). Cambridge University Press.

—— (2009b). Cognitive maps and the language of thought. *The British Journal for the Philosophy of Science*, 60(2), 377–407.

—— (2009c). Predication and cartographic representation. *Synthese*, 169(1), 175–200. https://doi.org/10.1007/s11229-008-9343-5

Rescorla, R. A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black and W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.

Ristau, C. A. (1991). Aspects of the cognitive ethology of an injury-feigning bird, the piping plover. In *Cognitive ethology the minds of other animals* (pp. 91–126). Lawrence Erlbaum.

Samuels, R., Stich, S. and Bishop, M. (2002). Ending the Rationality Wars: How to Make Disputes About Human Rationality Disappear. In R. Elio (Ed.), *Common Sense, Reasoning and Rationality* (pp. 236–268). Oxford University Press.

Shannon, G., Slotow, R., Durant, S. M., Sayialel, K. N., Poole, J., Moss, C. and McComb, K. (2013). Effects of social disruption in elephants persist decades after culling. *Frontiers in Zoology*, 10(1), 62.

Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C. and Frith, C. D. (2014). Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences*, 18(4), 186–193.

Smith, J. D., Couchman, J. J. and Beran, M. J. (2014). Animal metacognition: A tale of two comparative psychologies. *Journal of Comparative Psychology*, 128(2), 115.

——, Coutinho, M. V., Church, B. A. and Beran, M. J. (2013). Executive-attentional uncertainty responses by rhesus macaques (Macaca mulatta). *Journal of Experimental Psychology: General*, 142(2), 458.

Stroud, B. (1979). Inference, belief, and understanding. *Mind*, 88(1), 179–196.

Taylor, A. H., Miller, R. and Gray, R. D. (2012). New Caledonian crows reason about hidden causal agents. *Proceedings of the National Academy of Sciences*, 109(40), 16389–16391.

Taylor K. A. (2003). *Reference and the rational mind*. CSLI Publications.

Timberlake, W. (1994). Behavior systems, associationism, and Pavlovian conditioning. *Psychonomic Bulletin & Review*, 1(4), 405–420.

Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*, 55(4), 189–208.

Tversky, A. (1977). Features of similarity. *Psychological Review*, 84(4), 327.

Valaris, M. (2016). What reasoning might be. *Synthese*, 1–18. https://doi.org/10.1007/s11229-016-1034-z

Watanabe, S. (2006). The neural basis of cognitive flexibility in birds. In E. A. Wasserman and T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 619–639). Oxford: Oxford University Press.

—— (2010). Pigeons can discriminate "good" and "bad" paintings by children. *Animal Cognition*, 13(1), 75–85.

—— and Huber, L. (2006). Animal logics: Decisions in the absence of human language. *Animal Cognition*, 9(4), 235–245. https://doi.org/10.1007/s10071-006-0043-6