

Are epistemic emotions metacognitive?

Peter Carruthers

To cite this article: Peter Carruthers (2017) Are epistemic emotions metacognitive?, *Philosophical Psychology*, 30:1-2, 58-78, DOI: [10.1080/09515089.2016.1262536](https://doi.org/10.1080/09515089.2016.1262536)

To link to this article: <http://dx.doi.org/10.1080/09515089.2016.1262536>



Published online: 02 Dec 2016.



Submit your article to this journal [↗](#)



Article views: 173



View related articles [↗](#)



View Crossmark data [↗](#)

Are epistemic emotions metacognitive?

Peter Carruthers

Department of Philosophy, University of Maryland, College Park, MD, USA

ABSTRACT

This article addresses the question whether epistemic emotions (such as surprise, curiosity, uncertainty, and feelings of knowing) are in any sense inherently metacognitive. The paper begins with some critical discussion of a recent suggestion made by Joelle Proust, that these emotions might be implicitly or procedurally metacognitive. It then explores the theoretical resources that are needed to explain how such emotions arise and do their work. While there is a perennial temptation to think that epistemic emotions are somehow *about* the cognitive states of the person undergoing the emotion, we will see that such views can and should be resisted.

ARTICLE HISTORY

Received 10 January 2016
Accepted 9 November 2016

KEYWORDS

Implicit; memory;
metacognition; procedural;
surprise; uncertainty

1. Introduction

Psychologists who study metacognition define it as involving metarepresentations of one's own states and processes. Metacognition is the kind of cognition that monitors, models, and controls regular first-order cognition – where the “monitoring and modeling” involves representations of first-order cognitive states and processes involved in learning, remembering, emoting, and knowing (Dunlosky & Metcalfe, 2009; Flavell, 1979; Nelson & Narens, 1990). Indeed, metacognition is thought to involve *judgments* about one's own mental states and processes; it involves *thinking about thinking*. These judgments aren't necessarily conscious, of course; and it has certainly never been required that they should be expressed in speech (whether overt or covert). Nor is it thought that these metacognitive judgments need reflect direct awareness of the states and processes they are about. On the contrary, the consensus in the field is that metacognitive judgments are cue-based, relying on factors such as processing fluency or disfluency, behavioral hesitation, and a variety of epistemic emotions (feelings of knowing, feelings of uncertainty, feelings of surprise, tip-of-the-tongue feelings, and so forth; Dunlosky & Metcalfe, 2009).

Likewise, comparative psychologists who study metacognition in nonhuman animals have thought that they were finding evidence of simple forms of self-directed metarepresentation in animals. The main paradigm in the field employs tests of so-called “uncertainty monitoring” (Smith, Beran, Couchman, & Coutinho, 2008; Smith, Shields, & Washburn, 2003), and the findings have been heralded as demonstrating the phylogenetic origins of full-blown human-like self-awareness (Couchman, Beran, Coutinho, Boomer, & Smith, 2012; Couchman, Coutinho, Beran, & Smith, 2009). Some have gone further, and have claimed support for introspection-based accounts of the origins of third-person social cognition – or “mindreading” – since the species in question appear to be incapable of the latter (Couchman et al., 2009). It is thus claimed that metarepresentational capacities initially emerged in evolution to subserve first-person metacognitive functions. Only later were they exapted for use in understanding the mental states and behaviors of others (perhaps when combined with imaginative or so-called “simulation” abilities; see Goldman, 2006). Since on this view metacognition is held to be prior to and independent of third-person mindreading, I shall refer to it as the “two-systems account.”

These comparative claims have been critiqued by others, who have claimed that the kinds of uncertainty-monitoring and memory-monitoring tests that have been employed with animals fail to demonstrate the presence of metarepresentational capacities (Carruthers, 2008; Carruthers, Fletcher, & Ritchie, 2012; Carruthers & Ritchie, 2012; Kornell, 2014; Perner, 2012). Although the tasks in question may require executively controlled decision-making of various sorts, these can be regular forms of *first-order* (non-metarepresentational) decision-making. Success in the tasks need not require self-awareness or modeling of one’s own mental processes. Moreover (and in line with the views of most developmental psychologists), those advancing these critiques have claimed that in humans the same mental faculty, involving the same basic conceptual resources, is involved both in metacognition of one’s own mental processes and in reading the minds of other people. Accordingly, I shall refer to this as the “one-system account.”

In reply to these criticisms, some two-system researchers have allowed that animals might only possess a sort of *implicit* sensitivity to their own mental states and processes, without explicitly representing them (Couchman et al., 2012). And Proust (2012, 2014), in particular, has developed a view according to which the metacognitive processes of animals, young human children, and even human adults (for much of the time) are *procedurally* rather than explicitly metarepresentational. It is claimed that epistemic feelings, especially, have the monitoring-and-control functions that are thought to be distinctive of metacognition, carrying information about mental states and processes, and guiding targeted interventions, while not explicitly representing those processes. These ideas will provide an initial critical focus for our discussion (in Sections 2 through 4). Thereafter (in Sections 5 through 7) we will examine the nature of epistemic emotions more closely, showing that they can be accounted for in first-order

(non-metarepresentational) terms. Overall, our question is whether there is any good sense in which epistemic emotions are genuinely *metacognitive* in nature. It will be argued that they are neither implicitly or explicitly metarepresentational.

Before we get to that, though, a few words of clarification are in order. First, the goal is not to provide an account of what epistemic emotions, as such, are. And I likewise propose to remain neutral on the question of the essential properties of emotions in general. All I take for granted is that a range of properties regularly co-occur with emotional states. Emotions generally involve some sort of appraisal of a thing or event as relevant to one's goals or values; they generally activate related motor plans (facial expressions, bodily postures, approach or retreat behaviors); they often result in some degree of bodily arousal; they generally co-occur with conscious feelings of some sort; and they normally involve positive or negative valence that gets directed at the object of the emotion. The same clustering of properties seems also to hold true of epistemic emotions.

In fact, for present purposes, epistemic emotions can be characterized extensionally, as the set that includes surprise, interest, curiosity, uncertainty, and feelings of knowing. Our question is whether these states (however they should be individuated) in any way represent (that is, metacognize) the cognitive states of the agent (whether past, present, or future). Does surprise represent falsified belief, for example? And does curiosity represent one's current ignorance of something, causing one to want to know it? Or is there a simpler, first-order, way of characterizing such states?

We shouldn't assume at the outset, however, that there are no differences among epistemic emotions that are relevant to our inquiry. On the contrary, surprise, for instance, seems to constitutively involve expectation or belief (which is contradicted by the surprising object or event), whereas nothing similar seems to be true of uncertainty. If surprise involves awareness of cognitive conflict, then that might provide an argument for its metacognitive status which won't work for other epistemic emotions. For this reason, Sections 5–7 will discuss each kind of epistemic emotion separately.

Moreover, nothing that I say should be taken as denying that we often have metacognitive awareness *of* our epistemic emotions. Nor do I deny that such awareness can play important roles in our mental lives. Often when one is surprised one is aware *that* one is surprised, for example, and one may respond differently as a result (suppressing one's initial facial expression, perhaps). But the fact that epistemic emotions are often *accompanied* by metacognitive awareness doesn't make them metacognitive. Our question is about the nature of the underlying epistemic emotion itself that enables it to fulfill its core functions (and whether that nature is first-order or metacognitive), not the processes that often surround such states in human beings.

In addition, it should be noted that the distinction between first-order and metacognitive processes cross-cuts the distinction between conscious and unconscious ones. First-order states (e.g., perception of a cube) can be either conscious

or unconscious. And likewise, metacognitive states (e.g., knowledge that one is surprised) can be either conscious or unconscious. Nor does the distinction between first-order and metacognitive processes bear any direct relationship to the empirical distinction sometimes drawn between “System 1” (intuitive) and “System 2” (reflective) reasoning processes (Kahneman, 2011). For the latter, distinction is best understood as a contrast between processes that don’t constitutively depend on working memory and those that do (Carruthers, 2015; Evans & Stanovich, 2013). And both first-order and metacognitive processes can be of either sort.

Although System 2 reflection is always conscious (because implicating working memory) and is often accompanied by metacognitive awareness of the contents of reflection, it is another matter to claim that the states and processes that constitute reflection always *depend on* metacognitive awareness. It is plain that they don’t. As we will note in passing later, for example, prospective decision-making (which is reflective in nature and thus “System 2”) depends just upon valenced responses to represented future options. We can thus set aside the System 1/System 2 distinction as not directly relevant to our concerns.

2. Notions of implicitness

In what sense might metacognitive processes in animals be implicit, then? Sometimes “implicit” is used to mean “unconscious.” But that cannot be what is at stake here, if only because it is quite unclear how often *human* metacognitive processes are conscious ones. Moreover, people who have objected that uncertainty-monitoring experiments conducted with animals fail to demonstrate the presence of metacognitive processes have never relied on claims about consciousness. In some other contexts “implicit” is used to mean “not reportable in speech.” But again this cannot be what is at issue. For, of course, it is trivially true that *all* cognitive processes in animals are implicit in this sense, metacognitive or not. Nor have skeptics about the metacognitive status of uncertainty-monitoring ever insisted on knowing what the animals in question *would* report about their mental processes if only they could speak, or anything of the sort.

Yet another possibility is that something is implicitly represented if it is tracked by a processing-principle or inference-rule, but without there being any symbol-like mental structure encoding it. For example, consider the mindreading principle, *seeing leads to knowing*. This is explicitly represented if the inference that leads from a representation like, *John sees what is in the box* to the conclusion, *John knows what is in the box* passes through a major premise of the form, *seeing leads to knowing*. It is implicitly represented, in contrast, if there is a special-purpose mindreading inference-mechanism which, when provided with *John sees what is in the box* as input, spits out the representation, *John knows what is in the box* directly, without consulting any such major premise. Implicit metacognition, then, would comprise cognitive processes that are sensitive to and somehow track

those occurring elsewhere in the mind without representing them, while also functioning to modulate and correct the latter.

One initial problem with this idea when transposed to the animal metacognition literature, however, is that it presupposes explicit metarepresentations at the level of input and output. While the processing principle that mediates between the former and the latter might be left implicit in the processing-rule itself, it nevertheless seems that there must be metarepresentational input and output in order for it to qualify as an implicit *metarepresentation*. For considering what constitutes the processing-rule of a mindreading system as implicitly metarepresenting that seeing leads to knowing. This is that it plays the same role as an explicit premise with that content might do, mediating between premises and conclusions that are themselves metarepresentational. By parity of reasoning, then, an implicitly metacognitive self-directed processing-rule would also need to play a role like that of an explicit metarepresentational generic belief; and that would mean mediating among premises and conclusions that are themselves explicitly metarepresentational.

In an attempt to respond to this difficulty, it might be said that the input for implicit metarepresentational processing can consist of *nonconceptual* representations of the mental processes that get modulated as a result; and the output can be nonconceptual actions that do the modulating. Sometimes Proust (2012) seems tempted by this move, since she stresses that epistemic feelings are nonconceptual evaluative representations of first-order processes. Such feelings provide information about first-order processing while initiating procedures for modulating or responding adaptively to the results of that processing. It seems clear, however, that this combination of views fails to avoid a commitment to explicit metarepresentation. For the nonconceptual representations of color, shape, and motion that constitute the output of the early visual system are fully, and explicitly, representational, of course (albeit without yet employing *concepts* of color, shape, or motion). By analogy, then, one would think that the nonconceptual metacognitive content of one's epistemic feelings should also be explicitly representational: in this case, explicitly (albeit nonconceptually) *metarepresentational*.

Proust (2014) has other things to say in characterizing what she calls "procedural metacognition," as we will see shortly. But unfortunately she also saddles the competing one-system view (which she calls the "attributivist" account) with commitments that it need not (and does not) have. She claims that one-system theorists hold that any cues and principles relied on in first-person metacognition must likewise be available for third-person mindreading, and are of a purely doxastic (belief-like) sort. Partly for these reasons, she also claims that the one-system account is committed to thinking that first-person metacognition isn't constitutively embedded in, and dependent on, the ongoing mental processes of the agent. But all the one-system account is really committed to is that it is the same core conceptual capacities that are employed in the first-person as in the third.

Consistent with this claim, a one-system theorist can allow that first-person applications of mental-state concepts make use of internal cues that are unavailable for attributing mental states to others, and that these cues can involve both affective as well as doxastic elements. Moreover, a one-system theorist can allow that many of these first-person cues are produced by the very mental process being evaluated, in such a way that metacognition is deeply embedded in, and dependent on, one's ongoing mental life. Indeed, it can also be allowed that many of the cues that underlie metacognitive judgments (especially the epistemic emotions that will form our main focus here) do much of their work in cognition *independently* of the metacognitive judgments that they ground.¹ In fact, this is the core of the position to be defended in this paper: while Proust is right that epistemic emotions don't presuppose metarepresentational judgments of any kind, she is wrong to think that there is any interesting sense in which epistemic emotions are themselves metacognitive in nature.

Unfortunately, too, Proust relies on definitions of metacognition in general (whether implicit or explicit) that are too broad to be helpful. In Proust (2012), she adopts a definition from Hampton (2009) which is entirely negative, and is designed to exclude purely associative and stimulus-driven explanations of animal behavior. Since this definition is consistent with the sorts of executively controlled but nevertheless first-order (non-metacognitive) processes that have been proposed to explain the animal data (Carruthers, 2008; Carruthers & Ritchie, 2012), it cannot be considered adequate. Proust (2014), in contrast, offers the following definition:

Metacognition is the set of capacities through which an operating cognitive subsystem is evaluated or represented by another subsystem in a context-sensitive way. (p. 4)

The phrase "or represented" is included here to allow for the kinds of explicit metarepresentational metacognition that are claimed to be distinctive of all forms of metacognition by the one-system view. In contrast, Proust's form of *procedural* metacognition is held to involve one subsystem evaluating the performance of another *without* explicitly metarepresenting it. But as we will see, this is much too broad. It would encompass many kinds of process that no one would want to claim to be metacognitive in nature. In particular, there are a range of predictive models, used at many different levels in cognition, that fit Proust's definition of procedural metacognition. This is where we go next.

3. Predictive models

Predictive uses of implicit models are ubiquitous in human and animal cognition. Sensorimotor coordination, in particular, is thought to make use of three types of model: physical models to predict changing properties of the world, inverse models of the dynamical properties of one's own body to generate the motor commands needed to achieve desired outcome states, and forward models to predict the sensory consequences of one's own movements, used for swift online

correction and control (Jeannerod, 2006). There is evidence that even dragonflies use internal models of these types when pursuing and capturing prey (Mischiati et al., 2015).

Is there any good sense in which internal models are implicitly or procedurally metacognitive in nature? One might be tempted to think that forward models, in particular, could be so.² For the forward model uses an efferent copy of the motor instructions for movement to create a representation of what one should experience (both proprioceptively and visually) if the movement in question should unfold as planned, and it does so using processing-principles that contain an implicit model of the kinematics of the body. This representation is then received by a comparator mechanism where it can be aligned with the afferent sensory feedback from the movement itself. If the two match, then no adjustments are required. But if they do not (for example, because the object one is lifting is heavier than expected), then changes can be made in the execution of the motor sequence to correct for the discrepancy (e.g., by tightening one's grip).

Notice that these forward models seem to fit Proust's (2014) definition of metacognition. For they involve one operating cognitive subsystem (the motor-production system) being evaluated by another (the comparator/controller) in a context-sensitive way. And one might be tempted to think, too, that forward models are metacognitive in nature because they contain representations of one's own future experiences. They need not explicitly represent them *as* one's visual or proprioceptive experiences, of course. But in representing the expected contents of those experiences it might be said that they are *implicitly* metacognitive.

A similar case can be made with respect to visual processing, on the widely held assumption that the latter makes use of predictive coding of various sorts (Clark, 2013; Kok, Brouwer, van Gerven, & de Lange, 2013). On the account outlined by Bar and colleagues (2006), for example, fast magnocellular pathways are used to project a "gist" representation of the stimulus to orbitofrontal cortex where it activates a set of candidate stored concepts. These are prioritized for fit and relevance and projected back to higher areas of visual cortex, getting there some 50 ms before the arrival of information contained in the slower parvocellular processing stream. There they are matched against the contents of that stream, helping to resolve ambiguity and compensate for degraded or incomplete images. Since the process seeks a best match between initial expectations and further incoming information it, too, seems to fit Proust's definition of implicit metacognition. For perceptual contents are both evaluated and modulated by top-down expectations (again, in a context-sensitive way).

Notice, however, that if predictive models are implicitly metacognitive, then implicit metacognition is ubiquitous throughout the animal kingdom and at all levels of cognition. For, as noted above, even invertebrates employ them, as does processing within the visual system. Hence, there would be no reason whatever to think that implicit forms of metacognition constitute the first steps on the road to introspective self-awareness. Or if they are such a step, that step was taken quite

early in phylogeny, with the evolution of invertebrate life. And in truth there seems little plausibility in the suggestion that the use of predictive models to modulate the contents of either motor intentions or perception is any kind of step toward possession of a concept of intention, or toward possession of explicit metarepresentations of one's perceptual experiences as such. For the predictive models in question concern the world-directed or body-directed contents of perceptual experience. It is predictions of what the world should be like, or of what one's own body should be doing, that drive the modulating process, not representations of what one's *experience* of the world or body should be like. Although these predictions are matched against the content of incoming experience, those contents are likewise world-representing or body-representing. The processes involved are entirely first-order in nature.³

4. Central signaling

It might be replied that there is at least one significant difference between the epistemic emotions that are claimed to be procedurally metacognitive in nature (surprise, uncertainty, feelings of knowing, and the rest) and the predictive models discussed above. This is that predictive models operate *within* mental faculties (vision, action, and elsewhere), whereas epistemic emotions serve to guide the personal-level decisions that we take concerning what to do, or what to choose. For example, someone who is uncertain about the category of a stimulus (whether an image is densely or sparsely pixelated, say) might opt out of taking the test, choosing to avoid the mild penalty of a time-out for a mistaken answer. Here, the epistemic feeling of uncertainty seems to provide the agent with a *signal* concerning the likelihood of an error, leading to strategic (and adaptive) opting out.

We will return to the alleged signaling function of epistemic emotions shortly. For the present, let us focus on the point that these emotions serve to influence personal-level cognition. It is plain that this, by itself, cannot be enough to make any process that fits Proust's (2014) overly broad definition of metacognition genuinely metacognitive. For consider *attention*. This comprises two different (interacting) networks. One is the top-down attentional network linking dorsolateral prefrontal cortex, frontal eye-fields, and intraparietal sulcus, which selects some representations over others for "global broadcasting" and entry into conscious experience. The other is the bottom-up saliency system identified by Corbetta, Patel, and Shulman (2008), which links right-hemisphere ventral parietal cortex with the basal ganglia and right-hemisphere ventral prefrontal cortex. The saliency system continually monitors perceptual and mnemonic contents that are not currently in the focus of attention and appraises their relevance to both current goals and standing values, competing for control of the top-down attentional system via anterior cingulate cortex. When items of sufficient interest are identified, a decision is taken to switch the direction of top-down attention, resulting in global broadcasting and conscious availability of the contents in question.

Note that the operations of this system are quite flexible. Many different variables can influence whether one becomes conscious of one's name spoken in the background at a party, for example (including not just the clarity and volume of the speech, but the identity of the speaker, the degree of interest of the conversation that forms the current focus of attention, and the hearer's trait-like ability to control attention in general). And the resulting shift of attention, it can be argued, is genuinely action-like, resulting, like other actions, from a regular (albeit unconscious) decision (Carruthers, 2015).

Since the attentional networks monitor and evaluate unconscious perceptual and mnemonic contents, while controlling the contents of conscious experience, and since they perform their control-functions through flexible forms of decision-making, they appear to satisfy the criteria for so-called implicit metacognition according to Proust's (2014) account.⁴ And hence again it will follow that implicit metacognition is rife throughout the animal kingdom. For we know that attentional networks are highly conserved in mammals and birds (and perhaps in all vertebrates), and that they operate according to similar principles across species (Allen & Fortin, 2013; Knudsen, 2007; Mysore & Knudsen, 2013). So once again there will be no interesting respect in which implicit forms of metacognition (so understood) are the first steps toward the kinds of explicit self-awareness enjoyed by humans.

It seems, then, that if epistemic emotions are to be shown to be implicitly metacognitive in nature (to any interesting degree and in any interesting sense), then all the weight must fall on their alleged signaling functions. This is where we go next, examining the emotions of surprise, interest, and curiosity (Section 5), uncertainty (Section 6), and feelings of knowing (Section 7). Although it is perennially tempting to construe such emotions as being *about* one's own cognitive states, in fact these emotions can be characterized and explained in first-order terms.

5. Surprise, interest, and curiosity

Consider surprise. Some philosophers have claimed that it is an *explicitly* metacognitive emotion. They have said that to be surprised is to realize that one's experiences are in conflict with one's prior expectations or beliefs, thus giving one reason to take actions (such as looking closer) that might resolve the conflict (Davidson, 1982). But most would now agree that this over-intellectualizes surprise. Surprise is a state that is *caused* by conflicts between experience and expectation without representing those conflicts explicitly. (Of course, humans can become aware *that* they are surprised, and can thus not only *be* surprised, but be *aware* of themselves *as* surprised. This is an explicitly metarepresentational state. But it might well turn out to be uniquely human, depending on our distinctive mindreading capacities. And in any case awareness of surprise *as surprise* is distinct from the surprise itself.) This leaves open, however, that conflicts of belief are represented *implicitly* by feelings of surprise.

It is surely not sufficient for feelings of surprise to be implicitly metacognitive that they should *carry information about* conflicts of belief, however. For almost all mental states carry information about the occurrence of others. Any conscious experience, for example, carries the information that attention has been directed toward its contents; and any decision carries the information that belief-states and desire-states have just previously interacted to lead to the decision. What surely matters, if epistemic emotions are to be implicitly metacognitive, is that they should be used by downstream processes to modify other mental states (perhaps including those emotions themselves) in a way that depends on the information carried. And such uses had better happen in a controlled, flexible manner, too. For almost all mental states cause modifications of others. Rather, the modification either needs to result from a decision to bring about such a change, or it needs to belong among the *functions* of the metacognitive state that it should normally bring about such changes.

Surprise (sometimes more broadly, *interest*) is often listed among the basic emotions, common to many other creatures besides ourselves (Ekman, 1992; Izard, 1977). Although classified as an emotion, surprise seems not to have any intrinsic valence, however. There are, as one says, both good surprises and bad surprises. The valence of a surprise depends on its content, and not on the mere fact of the surprise itself. But surprise is always arousing, to some degree. And like many other basic emotions, it directly motivates not only a characteristic facial expression (especially widening of the eyes), but also a range of behaviors, including sustained attention to the surprising object or event, together with a set of information-gathering movements (approaching the object of surprise, looking more closely, sniffing it or tasting it, and so forth). Indeed, one might plausibly claim that the function of surprise is to facilitate the acquisition of new beliefs. Is this sufficient to show that surprise, and one's reactions to surprise, is implicitly metacognitive in nature? Only if one is prepared to say the same about vision, and about perception generally. For of course one might say that one of the functions of vision, too, is to facilitate the acquisition of new beliefs. It seems better to say that surprise is a *cognitive* emotion (having epistemic functions) without being (even implicitly) metacognitive.

Consider the flexibility of the actions that can be motivated by epistemic emotions like surprise and curiosity, however. If a novel object disappears out of sight, for example, one might, if one is curious, move in such a way that one can see it again. Is one's goal in such a case to *see* the object again, or to *find out* (i.e. learn) what it is? These are metarepresentational goals. If such goals form parts of epistemic emotions, or are at least caused by such emotions on a regular basis, then that might warrant describing epistemic emotions as metacognitive.

It should be granted that the only way we have of explaining such actions from the perspective of common sense is in terms of metarepresentational goals. But this may be because common sense knows little of the sensorimotor route to action (as opposed to goal-directed intentional action; see Milner & Goodale, 1995) and

the ways in which actions can be directly motivated by emotions independently of decision-making processes (Panksepp, 1998). An emotion of curiosity can motivate looking closer, or looking behind the barrier, in the same sort of direct manner that an emotion of fear can motivate running away, or an emotion of anger can motivate attacking. That is to say, actions of the appropriate sort are *primed* or *initiated* by the emotion combined with features of the circumstances, and need to be *inhibited* if they are not to be performed. No decision to act need be made, and the action need not have any explicitly represented goal (assuming that goal-states are the products of decision-making, rather than internal to a directly primed motor-plan).

Moreover, many animals need to explore their environments, of course. For example, honey bees spend their first few days as foragers making exploratory flights around the vicinity of the hive, constructing a mental map of their surroundings (Cheeseman et al., 2014; Menzel et al., 2005). And many foraging animals will check potential sources of food *to find out* (as we would say from the perspective of common sense) whether the fruit has ripened or whether the flowers have opened. These actions, too, might appear to have metarepresentational goals (to find something out, or to learn something). But no one really thinks that they do. Rather, the animals are, as it were, addressing a first-order question to the world: What is there around here? Are the fruits on that tree ripe yet? Such questions can motivate search behavior, guided in detail by the affordances of the environment (moving around obstacles and so forth), and drawing on background knowledge (such as the location of the tree in question on one's mental map).

The same sort of thing is true, I suggest, when one is curious about the identity of a now-hidden object. One's behavior can be motivated by a first-order question-like state: What is *that*? This question can motivate search behavior that can be unlimitedly flexible, including looking around corners, lifting up occluders, searching in likely hiding places, and so on. Note that common-sense psychology seems to require one to cite a metarepresentational goal when explaining such behavior. ("I *wanted to know* what it was, so I walked over to look.") But in reality, curiosity is a first-order affective state, which includes a first-order motivating question-like state, together with a suite of primed information-gathering behaviors.

Two things need to be done if this account is to be defended, however. First, can we cash out the metaphor of a "question-like state" in an acceptable way? Can we give a merely first-order characterization of the contents and functional role of the state that motivates search behavior? And second, what argument can be given for believing in the reality of such a state in humans and other animals? Let us take these issues in turn.

Questions are prompted by ignorance, and so is search. In general, when one asks a question, that is because one is currently ignorant of the answer. Likewise, when one searches for something or looks closer at something, that is because one is ignorant of its location or identity. Does this require that one be *aware that*

one is ignorant (which would be a metacognitive state)? Surely not. One can easily imagine a mechanism (in this case a sort of seeking-motivation) that is activated by failures of recognition, say, without needing to *represent that* recognition has failed. The motivation would be triggered whenever the processes that normally issue in recognition fail to generate an output.

Similarly, questions specify the kinds of things that could count as an answer. If one asks, “Where is home from here?” then a satisfying answer could take the form, “Home is half-a-mile in *that* direction.” And likewise we can suppose that when a bee gets lost (transported in a black box to a novel location by an experimenter, for example), it enters a motivational state that issues in a looping flight pattern, prompted by its failure to recognize its immediate surroundings. This continues in gradually widening circles until the bee locates a familiar landmark, from which it can compute the distance and direction to the hive (Menzel et al., 2005). At that point the state that motivated search behavior is extinguished/satisfied, just as a question can be answered, satisfying curiosity.

Granted that search behavior *could* be motivated by a first-order question-like state, what reason is there to think that it *is*? We have noted that flexible forms of search behavior are well nigh ubiquitous in the animal kingdom. Bees do it. Birds do it. All other mammals do it. And humans do it. Our choices are then: (1) to accept that the same sort of first-order motivation is present throughout; (2) to accept that the same metacognitive abilities are present in all these creatures; or (3) to maintain that humans and perhaps some other primates are motivated metacognitively, whereas invertebrates, birds, and all other mammals are not, despite the close similarities in their behavior. I submit that (1) is the most likely option. The second alternative is intrinsically implausible, and option (3) requires us to believe in an unmotivated evolutionary discontinuity.

It is worth stressing once again, however, that humans can of course also be aware *that* they are surprised, and they can experience themselves *as* curious about something. The emotions in question thereby acquire a metarepresentational aspect. But it can still be the first-order question-like state that leads to many of the curiosity-manifesting actions humans perform. This is not to say that the metarepresentational component is causally inert, however. On the contrary, it can lead one to do things to modify the emotion itself, for example. Realizing that one is curious about someone’s strange behavior, one might tell oneself, “It is none of my business; look way.” Or one might simply decide to direct one’s attention elsewhere as a result. Such behavior genuinely deserves to be described as metacognitive. For one is aware of one’s emotion as such and intervenes to alter it.

It is worth noting that the common-sense assumptions critiqued here are widely taken for granted in the comparative and developmental literatures. For example, it is assumed that primates who move and orient themselves appropriately to find out which of three differently oriented opaque tubes have been baited with food thereby display metacognitive awareness (Krachun & Call, 2009). It may be true that these animals have an understanding of vision and the conditions for visual

access. But we don't need to suppose that this is so to explain the behavior. Rather, the animals are motivated by a first-order question-like state directed at the world ("Which one has the food?"), and activate the appropriate procedures to find out. One can know *how* to achieve visual access without understanding (cognitively) anything about vision.

Similarly, it is widely assumed in the developmental literature that when children ask questions of their parents and other adults this, too, displays metacognitive awareness (Goupil, Romand-Monnier, & Kouider, 2016; Mills, Legare, Bills, & Mejias, 2010). The child asks the question because she realizes that she is *ignorant* of the answer, and wants *to know* it. Now in this case, I have no doubt that children have the conceptual resources to form such metacognitive beliefs and goals (Carruthers, 2013). But we don't need to appeal to those resources in order to explain the behavior. In reality, question-asking may manifest a more primitive epistemic emotion of curiosity, which involves a world-directed (non-metarepresentational) question-like state. Although use of verbal behavior (asking questions) in pursuit of such a goal is uniquely human, the attitude that it manifests may not be.

6. Uncertainty

Consider another epistemic emotion, this time one that is at the heart of the metacognitive research program in comparative psychology: uncertainty. There are seemingly two different forms of it: one can be uncertain of some fact, and one can be uncertain of one's ability to complete an action. There are good reasons to think that the latter notion is the basic one. This is because not any kind of affective state caused by failure to access knowledge constitutes that state as a form of uncertainty. In many cases the affective state is, rather, one of surprise or curiosity. It only counts as uncertainty when there is some task whose successful execution depends on the fact in question. If one is simply unaware of the category of a stimulus, then this might give rise to curiosity. But if one is required to take an action of some sort (like answering a question) whose success depends on correct categorization, then this is uncertainty.

I suggest that uncertainty is the affective state that arises when one appraises an action as unlikely to succeed. (Or at any rate, as not likely *enough* to succeed, given what is at stake and one's own tolerance for risk.) The result is some degree of arousal, and some degree of negative valence targeted at the thought of the action in question. This might make one hesitate, or it might lead one to select an alternative action instead. Or in cases where the appraisal depends on ignorance of some fact that is necessary for success, uncertainty may issue in information-seeking behavior (looking closer, for example, or in the case of humans, asking a question).⁵

Consider a cat pacing back and forth on the edge of a roof while contemplating a leap to a nearby tree. The cat is uncertain of its ability to do the jump. This

motivates it to examine the distance from different angles. And the cat might even assume the crouch position that would initiate a leap, seemingly rehearsing the action itself. We should take seriously, I think, that this is what the animal is actually doing: engaging in a form of mental rehearsal of action. For a cat facing a leap from a roof, or a squirrel contemplating a jump from one branch to another in the forest canopy, may be confronted with a combination of heights, distances, and surfaces never before encountered. The animal thus needs to make a *prediction* of likely success in these particular circumstances. The action is not a routine one, nor can it draw on any specific memory of previous success. Mentally rehearsing the action can enable the normal predictive processes that operate in perception (drawing on and extrapolating from previous experience) to generate an answer (Seligman, Railton, Baumeister, & Sripada, 2013). If the animal appraises the action as unlikely to succeed, then this will result in negative valence directed at the thought of it, making the action in question seem like a bad option. The resulting state of uncertainty will motivate the animal to attempt something else instead (finding another way off the roof).

Now consider a monkey engaged in one of the tasks that have been said to require uncertainty monitoring (Smith et al., 2003, 2008). On a given trial, the monkey is presented with a stimulus and a pair of primary response options. It knows that it should make the D response to get a reward if the stimulus is densely pixelated, and that it should likewise make the S response if the pixelation is sparse. It also knows that an error will result in an unwelcome time-out. But the stimulus in question may be one that the experimenters have placed close to the dense/sparse boundary. So the discrimination is a difficult one. The monkey is thus uncertain about doing D and equally uncertain about doing S. Both are appraised as unlikely to succeed, making them seem like bad options to take. But the monkey also knows that there is an opt-out response, which carries no direct reward, but which takes the animal immediately to the next trial (and a further chance to receive a reward) without a time-out. This alternative option may then seem best to the monkey in the circumstances, leading it to opt-out.

Nothing here needs to be monitored except the likelihood of success, together with the resulting seeming-goodness of the various alternatives. And although the monkey is *in* a state of uncertainty, it doesn't need to be aware that it is. It just needs to be aware that the two primary response options seem like bad ones in the circumstances, whereas the opt-out response does not. This is just regular first-order valence-based decision-making of the sort that humans engage in all the time (Carruthers, 2015; Damasio, 1994; Seligman et al., 2013).

It should be noted, however, that different authors make differing assumptions about the nature of valence, and some of these bear on our topic. Damasio (1994) famously claims that humans rely on “somatic markers” when deciding among options for action. Although he is not fully explicit about the nature of these markers, they seem to be representations of one's own actual or predicted *bodily* states. They would thus qualify as first-order in character. Others regard

valence as a common currency of *value*, involving nonconceptual representations of subjective value, which are likewise first-order in nature (Levy & Glimcher, 2012). In contrast, Gilbert and Wilson (2005, 2007) think of valence in terms of hedonic properties of experience. It then follows that prospective decision-making would be metacognitive, requiring one to represent one's own future experiences as such. Carruthers (forthcoming) reviews these differences, arguing in support of the nonconceptual-value account – pointing out, among other things, that the hedonic approach faces many of the notorious problems that confront motivational hedonism (Sober & Wilson, 1999).

Defenders of animal metacognition have often argued that the differential performance of old world monkeys (macaques) and new world monkeys (capuchins) in tests like these is a problem for first-order explanations, including ones of the sort just offered (Smith, 2005; Couchman et al., 2012). For if the animals that are successful in these tasks are so merely by appraising the chances of success, then why should we see species differences? Why is it that macaques make adaptive use of the opt-out response, whereas capuchins in the same experiments do not? But these authors fail to note that bees, too, pass the same tests of “uncertainty monitoring” as do macaques (Perry & Barron, 2013). So the species differences in question are more likely to have to do with differing species approaches to risk, rather than overall cognitive sophistication.

The same conclusion is supported by the finding that there are large individual differences in “uncertainty monitoring” *within* species. Among humans, macaques, and also bees there are some individuals who never use the opt-out response whereas others make adaptive use of it. Rather than postulating large intra-species differences in cognitive sophistication, it seems better to explain the finding in terms of something already known to be the case, namely that there are large individual differences in risk-tolerance.

What, then, might explain the species-difference between macaques and capuchins in such tests? Notice that macaques are generalist foragers that have managed to find a niche for themselves in many diverse ecologies (some of them quite challenging), whereas capuchins are leaf and fruit-eating forest-dwellers. We can therefore suppose that macaques have lower tolerance for risk, and are perhaps more finely calibrated in their estimates of risk. Indeed, just as Carruthers and Ritchie (2012) predicted on such grounds, it has now been found that capuchins *do* make adaptive use of the opt-out response when they are required to make riskier six-way discriminations as opposed to the usual bimodal ones (Beran, Perdue, Church, & Smith, 2016).

7. Metamemory and the feeling of knowing

Finally, consider the sorts of tests that are alleged to show that monkeys have metacognitive access to their own memories. In a match-to-sample task where the difficulty can be varied by altering the wait-time, monkeys will choose to opt

out of taking the memory test more often in cases where they have to wait longer, where remembering is more difficult (Hampton, 2001, 2005, 2009). It is said that they are *aware* that they no longer remember the original stimulus, which is why they opt out. Of course, it is true that the monkeys, just like humans, *might* have such metacognitive awareness. Like us, they might be aware of themselves *as* remembering or failing to remember something. But the data fail to support this interpretation.

Many creatures *have* memories, of course, both long-term and short-term. We know that many species of food-caching birds have remarkable recall for the locations of their caches (Bednekoff & Balda, 1997). And we know that jays and other corvids can recall the *what*, *where*, and *when* components distinctive of human-like episodic memory (Clayton, Dally, Gilbert, & Dickinson, 2005; Clayton, Yu, & Dickinson, 2001, 2003). Moreover, much of what we know about the neuroscience of working memory derives from work done with monkeys, using match-to-sample tasks among others (Goldman-Rakic, 1995; Goldman-Rakic, Funahashi, & Bruce, 1990).

Now consider what a food-caching bird will do when hungry. It will consult its memory, of course. This does not mean that it *tries to remember*, where remembering forms the metacognitive goal of its mental activity – although note, once again, that common-sense psychology pretty much forces such an explanation on us. Rather, a first-order signal is sent to the long-term memory system, asking that system a first-order question (as it were): Is there a food cache near here? Perhaps a representation of food, or of a specific type of food-item, is projected to the long-term memory system where it activates related representations. These are evaluated for relevance (especially in light of the existing hunger-motivation), with the most relevant becoming a focus of attention and entering the bird's working memory.

We can suppose that if a representation of a location (perhaps one that passes some threshold for determinacy and vividness) enters the bird's working memory as a result of the query, then this is used to guide an appropriate direction of flight. And note that the representation of a location doesn't have to be categorized *as* a memory in order to acquire the causal role of a memory. Indeed, it would be quite odd if the first-order causal role distinctive of a given type of mental state (in this case, memory) required a higher-order representation to be present specifying that a token of that type of state is currently active.

But what does the bird do if it *fails* to retrieve a cache-location from memory? Presumably (since it is hungry) it sets out to do some exploratory foraging instead. But there are likely to be intermediate cases. There will be cases where some fragment of a memory of a location comes to the bird's mind. Perhaps, it recalls a knot-hole in a tree, but without the location of the tree being specified. In these circumstances the bird might persist in querying its memory. What the bird may experience, in fact, is what humans would call a "feeling of knowing." For such feelings tend to be grounded in retrieval of related memories or components of the target memory (Dunlosky and Metcalfe, 2009). When the animal needs to

make a decision about whether to continue querying memory or to do something else instead, it will engage in something resembling a cost–benefit calculation, not unlike the calculations we know many species of foraging animal engage in when they confront the problem of stay-and-exploit vs. shift-and-explore.⁶ There will be cues (such as fragments of memory) that indicate to the bird that it should stick with its present (memory-querying) strategy, which will need to be weighed against the likelihood of success if it engages in exploratory foraging instead.

Now return to the monkeys involved in a so-called “memory monitoring” task. Presumably the monkeys, like humans, solve the task by retaining in working memory a representation of the target stimulus. But for them (just as for humans) such a task can be hard, since distraction will result in that representation being lost. So the more time elapses between presentation and test, the more chance there is that the information in working memory has become fragmentary and incomplete, or lost altogether. The monkey knows that its task is to touch the item on the test screen that matches the one presented earlier. If it retains in working memory a representation of that item, then it should appraise the chances of success in the test to be good, and therefore opt to take it. In contrast, if the monkey no longer has a working-memory representation at all, or only has one that is highly incomplete, then it should appraise the chances of success to be low, and opt to move immediately to the next trial instead. Success in this task requires the monkey to *have* a memory, not to *monitor* its memory. And opting out adaptively requires the monkey to respond appropriately if it *lacks* any specific memory of the stimulus. This doesn’t require it to be aware of itself *as* lacking a memory.

Might these ways of using memory amount to a form of *implicit* meta-memory? By acting in one way in the presence of a memory, and in another in its absence, does the animal display some sort of procedural meta-memory? It is hard to discern any grounds for saying so. For that one is disposed to act in one way if one has a memory, and in another if one does not is just what it *is* to have or lack a memory. This just describes the normal first-order causal role of memory in the cognitive and decision-making processes of creatures that possess memory-states. There is nothing *meta* about it.

8. Conclusion

There is a perennial temptation to regard epistemic emotions as metacognitive in nature, involving some form of representation of one’s own cognitive states. Indeed, such an interpretation is seemingly forced on us by the limitations of common-sense psychology. But in fact, epistemic emotions may be neither implicitly nor explicitly metarepresentational. Moreover, one can only continue to regard them as implicitly metacognitive by defining “metacognition” in such weak terms that we would have to allow that human and animal cognition is almost ubiquitously metacognitive. This drains such a claim of any of the interest that might otherwise attach to it.

Notes

1. This is even true, I think, of what psychologists call processing *fluency*. While many of the effects of fluency are known to be mediated by naïve metacognitive theories (Alter & Oppenheimer, 2009), others are more direct, caused by the positive valence that is produced by fluent processing (Topolinski & Strack, 2009). Although processing fluency is treated at some length in Proust (2014), it will not be discussed here. This is because it is not normally regarded as an epistemic emotion. Nor is the notion familiar to common sense in the way that epistemic emotions like surprise and uncertainty obviously are.
2. Note that the “double accumulator” account that Proust (2012) uses as her central example of a procedurally metacognitive system is a predictive model not unlike those described here. I think she can only be tempted to believe that the system is a metacognitive one because it is used to explain uncertainty-monitoring behavior, which is itself intuitively metacognitive in nature, as we will see. Once control systems of the “double accumulator” sort are seen to be well nigh ubiquitous in cognition, all temptation to see them as metacognitive should fall away.
3. Nagel (2014) uses a very different example to make essentially the same point as the one made in this section. She points out that cross-modal sensory integration appears to qualify as metacognitive by Proust’s (2014) criteria. But no one seriously thinks that it is.
4. Attention is not *explicitly* metacognitive, of course. For it doesn’t represent experiences and memories *as such* when deciding how to direct attentional signals, nor is there an explicit intention to render the resulting states conscious. Rather, the intentions take the form, *switch [attention] to that*, where the indexical picks out the first-order content represented by the percept or memory in question.
5. Recall, however, that information-seeking is not really metarepresentational in nature; it is motivated, rather, by a first-order question-like state.
6. Indeed, a number of researchers have noted the commonalities – and likely partially shared mechanisms – involved in both spatial search and memory search; see Hills (2006); Hills and Dukas (2012).

Acknowledgments

I am grateful to Joelle Proust and two anonymous referees for their comments on a previous draft of this paper.

Disclosure statement

No potential conflict of interest was reported by the author.

References

- Allen, T., & Fortin, N. (2013). The evolution of episodic memory. *Proceedings of the National Academy of Sciences*, *110*, 10379–10386.
- Alter, A., & Oppenheimer, D. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, *13*, 219–235.
- Bar, M., Kassam, K., Ghuman, A., Boshyan, J., Schmid, A., Dale, A., ... Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences*, *103*, 449–454.

- Bednekoff, P., & Balda, R. (1997). Clark's nutcracker spatial memory: Many errors might not be due to forgetting. *Animal Behaviour*, 54, 691–698.
- Beran, M., Perdue, B., Church, B., & 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, 32–43.
- Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind & Language*, 23, 58–89.
- Carruthers, P. (2013). Mindreading in infancy. *Mind & Language*, 28, 141–172.
- Carruthers, P. (2015). *The centered mind: What the science of working memory shows us about the nature of human thought*. New York, NY: Oxford University Press.
- Carruthers, P. (forthcoming). Valence and value.
- Carruthers, P., Fletcher, L., & Ritchie, J. B. (2012). The evolution of self-knowledge. *Philosophical Topics*, 40, 13–37.
- Carruthers, P., & Ritchie, J. B. (2012). The emergence of metacognition: Affect and uncertainty in animals. In M. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of metacognition* (pp. 76–93). New York, NY: Oxford University Press.
- Cheeseman, J., Millar, C., Greggers, U., Lehmann, K., Pawley, M., Gallistel, C., ... Menzel, R. (2014). Way-finding in displaced clock-shifted bees proves bees use a cognitive map. *Proceedings of the National Academy of Sciences*, 111, 8949–8954.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36, 181–204.
- Clayton, N., Dally, J., Gilbert, J., & Dickinson, A. (2005). Food caching by western scrub-jays (*Aphelocoma californica*) is sensitive to the conditions at recovery. *Journal of Experimental Psychology: Animal Behavior Processes*, 31, 115–124.
- Clayton, N., Yu, K., & Dickinson, A. (2001). Scrub jays (*Aphelocoma coerulescens*) form integrated memories of the multiple features of caching episodes. *Journal of Experimental Psychology: Animal Behavior Processes*, 27, 17–29.
- Clayton, N., Yu, K., & Dickinson, A. (2003). Interacting cache memories: Evidence for flexible memory use by western scrub-jays (*Aphelocoma californica*). *Journal of Experimental Psychology: Animal Behavior Processes*, 29, 14–22.
- Corbetta, M., Patel, G., & Shulman, G. (2008). The reorienting system of the human brain: From environment to theory of mind. *Neuron*, 58, 306–324.
- Couchman, J., Beran, M., Coutinho, M., Boomer, J., & Smith, J. D. (2012). Evidence for animal metaminds. In M. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of metacognition* (pp. 21–35). New York, NY: Oxford University Press.
- Couchman, J., Coutinho, M., Beran, M., & Smith, J. D. (2009). Metacognition is prior. *Behavioral and Brain Sciences*, 32, 142.
- Damasio, A. (1994). *Descartes' error*. New York, NY: Putnam.
- Davidson, D. (1982). Rational animals. *Dialectica*, 36, 317–327.
- Dunlosky, J., & Metcalfe, J. (2009). *Metacognition*. Thousand Oaks, CA: Sage.
- Ekman, P. (1992). Are there basic emotions? *Psychological Review*, 99, 550–553.
- Evans, J., & Stanovich, K. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8, 223–241.
- Flavell, J. (1979). Metacognition and cognitive monitoring: A new area of cognitive-developmental inquiry. *American Psychologist*, 34, 906–911.
- Gilbert, D., & Wilson, T. (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science*, 14, 131–134.
- Gilbert, D., & Wilson, T. (2007). Prospection: Experiencing the future. *Science*, 317, 1351–1354.
- Goldman, A. (2006). *Simulating minds*. New York, NY: Oxford University Press.
- Goldman-Rakic, P. (1995). Cellular basis of working memory. *Neuron*, 14, 477–485.

- Goldman-Rakic, P., Funahashi, S., & Bruce, C. (1990). Neocortical memory circuits. *Cold Spring Harbor Symposia on Quantitative Biology*, 55, 1025–1038.
- Goupil, L., Romand-Monnier, M., & Kouider, S. (2016). Infants ask for help when they know they don't know. *Proceedings of the National Academy of Sciences*, 113, 3492–3496.
- Hampton, R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, 98, 5359–5362.
- Hampton, R. (2005). Can Rhesus monkeys discriminate between remembering and forgetting? In H. Terrace & J. Metcalfe (Eds.), *The missing link in cognition* (pp. 272–295). New York, NY: Oxford University Press.
- Hampton, R. (2009). Multiple demonstrations of metacognition in nonhumans: Converging evidence or multiple mechanisms? *Comparative Cognition and Behavior Reviews*, 4, 17–28.
- Hills, T. (2006). Animal foraging and the evolution of goal-directed cognition. *Cognitive Science*, 30, 3–41.
- Hills, T., & Dukas, R. (2012). The evolution of cognitive search. In P. Todd, T. Hills, & T. Robbins (Eds.), *Cognitive search: Evolution, algorithms, and the brain* (pp. 11–24). Cambridge, MA: MIT Press.
- Izard, C. (1977). *Human emotions*. New York, NY: Springer.
- Jeannerod, M. (2006). *Motor cognition*. New York, NY: Oxford University Press.
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus and Giroux.
- Knudsen, E. (2007). Fundamental components of attention. *Annual Review of Neuroscience*, 30, 57–78.
- Kok, P., Brouwer, G., van Gerven, M., & de Lange, F. (2013). Prior expectations bias sensory representations in visual cortex. *Journal of Neuroscience*, 33, 16275–16284.
- Kornell, N. (2014). Where is the “meta” in animal metacognition? *Journal of Comparative Psychology*, 128, 143–149.
- Krachun, C., & Call, J. (2009). Chimpanzees (*Pan troglodytes*) know what can be seen from where. *Animal Cognition*, 12, 317–331.
- Levy, D., & Glimcher, P. (2012). The root of all value: A neural common currency for choice. *Current Opinion in Neurobiology*, 22, 1027–1038.
- 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*, 102, 3040–3045.
- Mills, C., Legare, C., Bills, M., & Mejias, C. (2010). Preschoolers use questions as a tool to acquire knowledge from different sources. *Journal of Cognition and Development*, 11, 533–560.
- Milner, D., & Goodale, M. (1995). *The visual brain in action*. New York, NY: Oxford University Press.
- Mischiati, M., Lin, H. T., Herold, P., Imler, E., Olberg, R., & Leonardo, A. (2015). Internal models direct dragonfly interception steering. *Nature*, 517, 333–338.
- Mysore, S., & Knudsen, E. (2013). A shared inhibitory circuit for both exogenous and endogenous control of stimulus selection. *Nature Neuroscience*, 16, 473–478.
- Nagel, J. (2014). The meanings of metacognition. *Philosophy and Phenomenological Research*, 89, 710–718.
- Nelson, T., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. Bower (Ed.), *The psychology of learning and information* (vol. 26, pp. 125–174). San Diego, CA: Academic Press.
- Panksepp, J. (1998). *Affective neuroscience*. New York, NY: Oxford University Press.
- Perner, J. (2012). MiniMeta: In search of minimal criteria for metacognition. In M. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of metacognition* (pp. 94–116). New York, NY: Oxford University Press.

- Perry, C., & Barron, A. (2013). Honey bees selectively avoid difficult choices. *Proceedings of the National Academy of Sciences*, 110, 19155–19159.
- Proust, J. (2012). Metacognition and mindreading: One or two functions? In M. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of metacognition* (pp. 234–251). New York, NY: Oxford University Press.
- Proust, J. (2014). *The Philosophy of metacognition*. New York, NY: Oxford University Press.
- Seligman, M., Railton, P., Baumeister, R., & Sripada, C. (2013). Navigating into the future or driven by the past. *Perspectives on Psychological Science*, 8, 119–141.
- Smith, J. D. (2005). Studies of uncertainty monitoring and metacognition in animals and humans. In H. Terrace & J. Metcalfe (Eds.), *The missing link in cognition* (pp. 242–271). New York, NY: Oxford University Press.
- Smith, J. D., Beran, M., Couchman, J., & Coutinho, M. (2008). The comparative study of metacognition: Sharper paradigms, safer inferences. *Psychonomic Bulletin & Review*, 15, 679–691.
- Smith, J. D., Shields, W., & Washburn, D. (2003). The comparative psychology of uncertainty monitoring and meta-cognition. *Behavioral and Brain Sciences*, 26, 317–373.
- Sober, E., & Wilson, D. (1999). *Unto others: The evolution and psychology of unselfish behavior*. Cambridge, MA: Harvard University Press.
- Topolinski, S., & Strack, F. (2009). The architecture of intuition: Fluency and affect determine intuitive judgments of semantic and visual coherence and judgments of grammaticality in artificial grammar learning. *Journal of Experimental Psychology: General*, 138, 39–63.