

Comment on Susanna Siegel, *The Rationality of Perception*

Peter Railton

University of Michigan

Correspondence

Peter Railton, University of Michigan.

Email: prailton@umich.edu

Abstract

In Susanna Siegel's compelling presentation of the case for the rationality of perception, a "significant part of the constructive defense" is played by the idea that there are "inferential routes to perceptual experience" (Siegel 2017, p. 94). Inferences, after all, are epistemically evaluable and bear on the rational standing of their conclusions. She argues that an obstacle to accepting this idea is a "Reckoning Model" of inference, and shows by example that we recognize as inferences various familiar kinds of responses to information that do not fit this model. She offers a more general approach to the nature of inference that fits these examples and accommodates inferential routes to perceptual experience. I argue that Siegel needs to say more about the mental processes involved in such inferences, and how it can be more than merely associative and yet still distinct from Reckoning. Fortunately, a psychologically- and conceptually-grounded distinction between model-free vs. model-based learning and guidance processes can provide a characterization that plays the role Siegel needs.

*In this comment I will be focusing on Chapter 5 of Susanna Siegel, *The Rationality of Perception* (Siegel 2017), as well as a more recent essay of hers that expands upon Chapter 5. Since I expect that one of the key challenges Siegel will face in making her case for the rationality of perception concerns the nature of inference and its relation to rationality, I hope that this focus will prove helpful.

1 | INTRODUCTION—THE PROJECT

Inference—when thought of as an actual mental process, *inferring*,¹ in which the mind transitions in a directed way from some information-bearing states to others—is of special interest in epistemology in part because it has one foot in the psychological realm and one foot in the normative. Inferring, along with perceiving, promises to be a process that is incontestably causal, yet nonetheless can realize normative properties, e.g., by being more or less *aptly responsive to reasons*. And thanks to this duality of perceiving and inferring, we, too, can qualify as more or less aptly responsive to reasons.

However, it is inferring in particular that has been thought to have an especially close connection to the normative property of *rationality*. Perceptual experience, some will say, enables us to be aptly responsive to reasons, but it is essentially passive—the content of immediate experience is, on this view, simply “given”—and so deficits in perceptual acuity, for example, my poor vision, are not deficits in rationality. By contrast, how I *respond* to the fact of my poor vision—whether I take steps to check my vision against my other senses or against the testimony of others, whether I take my poor vision into account in claiming authority on perceptual matters, etc.—*does* bear on my rationality. In this case, then, it is not the sensory information I simply happen to *receive* that bears on my rational standing, but what I *do with* this information—how I compare it with other sources, or connect it with other information, or project it into the future, or use it to guide the search for more information—that qualifies for rational assessment. These doings are exercises of *epistemic agency* for which I can be held responsible or given credit as a more or less rational epistemic agent. And they are, or at least mostly are, *inferred*s. Thanks to inferring, I can play an active role in determining the extent to which my beliefs are *supported* in epistemically appropriate ways by the information I receive, even if that information itself is epistemically defective. In a paper extending the account of inference in *The Rationality of Perception*, Susanna Siegel writes:

- (i) Inference is a paradigm of person-level reasoning that redounds well or badly on the subject. [p. 1]²
- (ii) The hallmark of inference is that the conclusions drawn by inferers epistemically depend on the premises from which they are drawn. [p. 1]³

Inferring comes in many forms, corresponding to different kinds of *support relations*—inductive inference supported by evidential, causal, or statistical relations, practical inference supported by past experience and means-ends relations, deductive inference supported by logical or algebraic relations, etc., which form the ground upon which epistemic dependence supervenes.

¹Annoying as this will probably prove to be, I will often use the term ‘inferring’ or ‘inferreds’ rather than ‘inference’ or ‘inferences’ when talking about psychological processes, to try to avoid conflation with the abstract argument forms we also call ‘inferences’.

²Siegel (2019). Unattributed page references in the text refer to this article. In what follows, I will be taking ‘reasoning’ here in a generic sense that does not simply *entail* that reasoning is an inference.

³I will be identifying four “marks of inference” in Siegel’s (2019) discussion, (i)-(iv), but should stress that these are not broken out and numbered as such in Siegel’s text, so I am taking a certain liberty in doing so. Most of what follows will not depend upon whether all four are, for her, “marks” or “hallmarks” of inference, whether they embody redundancy, or whether she would add others.

A good account of inferring should help us see why inferring is important for epistemic agents, and one plausible answer fits well with Siegel's approach: inferring is of special value to inferers since it is a way of *leveraging* the information they already have to reach a new, or more qualified, or better-understood information-state, even when they cannot obtain new information. If they leverage their existing information by *following* or *being responsive to* the kinds of evidential, causal, probabilistic, etc., connections upon which epistemic support supervenes, then the new or changed information-state has some chance of being an improvement—at least by the individual's own lights. So we can think of inferring as a *response* to one's information-state, explicitly or implicitly intended to reach a conclusion that is supported by that state. Siegel thus advances a generic "Response Hypothesis" concerning what is common to all the different forms of inference. According to this hypothesis, inferring is:

- (iii) ... a distinctive kind of *response* to an informational state, or to a combination of such states, that produces a conclusion [p. 10, emphasis added]

Let's unpack this a bit. To call the final state in a causally-connected sequence of thoughts produced in a mind the "conclusion" of that sequence is to say something more than that it is the last member of the sequence. To count as a conclusion, the *content* of the final member of the sequence must somehow, at least from the individual's standpoint, *follow from*, or *be drawn from*, or *be supported by*, or *be reached via*, etc., the *contents* of the earlier steps. A conclusion thus reflects the way the individual "takes" the contents of the previous members of the sequence to bear upon the last member: they are "taken" to epistemically support the final member, and this is arguably characteristic of deductive, inductive, abductive, analogical, means-end, etc. inference. Of course, we do not require that individuals possess such concepts as *<epistemic support>* or *<epistemic dependence>*, but inferers will nonetheless be responsive to features that can bear (or seem to bear from the standpoint of the individual) the relations that correspond to these concepts. For example, if I conclude that *q* given that I take things to be *p*, I will understand that if *p* is called into question, this will also to some degree call into question my conclusion that things are *q*. Features (i)-(iii) thus are plausibly thought to be marks of a very wide range of inferential practices.⁴

The idea that inferring involves not only arriving at a conclusion-state, but *producing* this state in response to one's antecedent information-state, brings to light the causal as well as the epistemic aspect of inferring and their close relationship. In inferring, the individual reaches her final thought in part *because* of the earlier thoughts, where this *because* relation is agential as well as epistemic. The inferring must be something she *does*, and so be a "person-level" activity on her part.

Does this "person-level", agential aspect of the sort of inferring we are here concerned with imply that a mental transition will count as an inference just in case the individual is explicitly aware of the full contents of the states involved and expressly controls each step along the way? That would hardly be plausible—inferential processes are often very complex, and, as any student of logic knows, to spell out all the assumptions and premises even of perfectly ordinary inferences would be a very demanding task. Consider this example of an everyday inference, provided by Paul Boghossian:

⁴Even in means-end reasoning, the contribution a means is thought to make toward attaining one's end is reason to believe—epistemic support for—taking the means to be justified.

On waking up one morning I recall that:

(1) It rained last night.

I combine this with my knowledge that

(2) If it rained last night, then the streets are wet.

to conclude:

(3) The streets are wet.

This belief then affects my choice of footwear. [Boghossian 2014, p. 2]

Boghossian's knowledge is no doubt richer than (2) suggests—for example, he only has certain streets in mind, and if he were in an arid part of the world, or if the night's rain had been brief, or if the weather were sunny and streets were gravel, etc., he would not expect the streets still to be wet. It is because of this richer background, and his assumptions about these factors, most of which are no doubt implicit, that he is willing to rely upon (2) in concluding (3) from (1) and making a decision about footwear. Recall that *inferring* is a mental action not a logical schema, so we need to explain how he “took” his information-state such that (3) was a conclusion he'd reach in the circumstances, and (1) and (2) give us only a glimpse of this—most of the complex causal knowledge in play remains implicit. Indeed, even the conclusions of typical “person-level” inferences can be implicit, as when one realizes, while talking with someone in a conditional mode about future planning, that one has already reached a conclusion about what one will do, without having explicitly noticed how or when.

However, as Siegel notes, we must be careful to distinguish *implicit* elements in person-level inferences from entirely *sub-personal* processes, lest we run afoul of mark (i). Vision scientists, for example, often use the word ‘inference’ to characterize computations that take place early in the stream of visual processing, e.g., the “inference” from a certain pattern of firing of discrete retinal cells to the positing of a continuous boundary between objects in the visual field. Siegel agrees that this use of ‘inference’ is legitimate, insisting only that individuals are not rationally accountable for purely sub-personal inferences:

Some phenomena aptly labelled ‘inference’ don't redound on the subject's rational standing at all. For instance, inferences in which the premise-states are states of early vision with no epistemic power to justify beliefs fall into this category. Here I set those phenomena aside.
[p. 2]

An important distinction between implicit person-level inferences and sub-personal inferences is that person-level inferences, as actions, have an intentional character, giving them *success conditions*, that is, conditions for what will count as performing that particular action well or completely. Given the epistemic purport of inferences, their success conditions include an epistemic component. Thus Siegel writes:

(iv) Inferences can be epistemically better or worse, depending on the epistemic status of the premises and the relationship between the premises and the conclusions. [p. 1]

2 | FROM IDLING TO RECKONING

We can perhaps get a clearer picture of the combined force of (i)-(iv) by considering a form of person-level mental transition that does not bear these marks, mental idling:

Idling: Thoughts are occurring in an individual's mind one after another in a causal sequence, such that each thought plays a causal role in the genesis of the next, but the individual herself is not imparting any direction to the sequence. As far as the agent is concerned, the sequence of thought simply happens.

Mental idling is a form of mental activity, which might be more or less implicit, but it has no success conditions and issues in no conclusion, though it might at any moment cease. If we were to happen upon someone just emerging from a bout of Idling, and ask him, "So, then, what do you conclude?", he would likely reply, "I didn't conclude anything at all. I wasn't trying to figure anything out." This bout of mental idling therefore would not *itself* be successful or unsuccessful, though one could be successful or unsuccessful at entering or leaving the state of mental Idling. For example, if by chance the last thought in the idling sequence was a non-obvious logical consequence of the contents of the thoughts that had gone before, since there was no guiding inferential intention, the relation of epistemic support between the contents of earlier and later thoughts would be happenstance from the standpoint of the individual—not something he could claim rational for, or would be likely to. Perhaps he'd be just as surprised as we are. "Fancy that! I had no idea." Mental Idling, then, does not bear marks (i)-(iv)—it is not bad inferring, but *non-inferring*.⁵

An orthodox theory of inferring does not, however, stop at identifying certain "marks" of inference, such as (i)-(iv). Rather, it seeks to give an account of the nature of inferring in terms of necessary and sufficient "structure and components" (p. 2) that will *explain* why inferring, but not, for example, mental idling, bears such marks as epistemic evaluability and rational accountability. The orthodox theory of inferring upon which Siegel focuses has at its base what she characterizes as the *Reckoning Model*:

Reckoning Model (RM): When you infer q from p , you have in mind premise states with p (or its conjuncts) as its content. You also have in mind a *reckoning state* with the content: p supports q . You draw the conclusion that q *because* you have in mind both the premise states and the reckoning state. [adapted from Siegel 2019, introduction]

The Reckoning Model is the target of Siegel's criticisms, but not because she is proposing an alternative account of the "structures and components" essential for a mental process to be an inferring. Indeed, her project is one of "illumination without analysis":

... the nature of inference may be illuminated even without positing any structure beyond what's posited by the hypothesis that inferring is a distinctive kind of response to an informational state, or to a combination of such states, that produces a conclusion. [p. 10]

She calls this "the response hypothesis", and notes that, while "[t]he reckoning model entails the response hypothesis, the response hypothesis does not entail the reckoning model" (p. 10).⁶

So, rather than building up from more elementary mental mechanisms to acts of mind that could bear marks (i)-(iv), Siegel proposes to start with marks (i)-(iv), and canvas a wide array of mental

⁵Siegel uses the term 'jogging', drawn from John Broome (2014), for something like this. But actual jogging is an action that might be well or badly done. Even if the brain "idles well", the way a car's engine can "idle well" when it doesn't run too fast or stall out, this is not an action on the part of the agent.

⁶Indeed, a recent survey of attempts to provide such an analysis concludes that we must either give up the idea of ourselves as rational epistemic agents or accept something like rule-following as an unanalyzed primitive (Boghossian 2014).

processes to which we normally do—or do not—apply the kinds of epistemic evaluation and rational responsibility reflected in (i)-(iv).⁷ Once we've embarked on this project, she argues, we'll quickly see that inference is a much broader phenomenon than proponents of the Reckoning Model and other traditional models recognize. By canvassing this diversity guided by marks (i)-(iv), we can open the way for a more adequate understanding of inferring in its various guises. The result in Siegel's hands is indeed illuminating: truer to actual inferential practice and a welcome fresh start in an area of philosophical inquiry where attempts to give outright analyses so often lead to circularity or dead ends.

3 | THE CANONICAL RECKONING MODEL OF INFERENCE

According to the Reckoning Model (RM), an inferring differs from (for example) a bout of mental idling in two primary ways: (1) in inferring, the inferrer "takes" her antecedent thoughts to support her consequent thoughts, where this involves representing to herself the (supposed) support relationship via a "reckoning-state";⁸ (2) the inferrer reaches the consequent thought she does at least in part *because* of this reckoning-state: it plays a causal role, along with the antecedent thoughts (premise-states), in generating the consequent thought (conclusion-state), where this role depends upon that "taking". The reckoning-state thus figures in the agent's *rationale* in reaching the conclusion she does, and inferences can be more or less well done depending upon the epistemic quality of this rationale.

Must the agent be *aware* of this rationale in order to infer? As noted above, requiring explicit awareness would be excessively restrictive. So a weaker self-awareness condition is called for, which Siegel (2019, p. 2) formulates, following Boghossian (2014), as:

Self-Awareness Condition (SAC): "[Inference is] mental action that a person performs, in which he is either aware, or can become aware, of why he is moving from some beliefs to others."

In any given instance of inferring, the contents of the reckoning-state, or of the premise-states, or of the conclusion-state, or of any combination of these, might be implicit. But (SAC) places a limit on this—the rationale of the inferring, the "why", must, at least under conditions favorable to reflection, be such that the agent could be aware of it. Combining (SAC) with (RM) yields what Siegel calls the "Canonical Reckoning Model", which she takes to be the best-motivated and most stable form of that model, and can be formulated as:

Canonical Reckoning Model (CRM):

The reckoning state (see RM) is in principle accessible by reflection (the inferrer "can become aware of why he is moving from some beliefs to others").

⁷Here as elsewhere it is important to remind ourselves that a theory of *inferring*, as a psychological process or action, is not a theory of *inference* as a logical operation.

⁸I will hyphenate 'reckoning-state' in order to make clear that it is a term of art. Lots of different kinds of mental states might deserve to be called "reckoning states" in the ordinary sense. Similarly for 'premise-state', 'conclusion-state', etc. And we will capitalize 'Reckoning' when used to refer to the model or its advocates, to distinguish this special use from ordinary uses of the term.

The becausal condition is not merely causal, but includes the inferrer's first-person rationalization of why the conclusion is drawn.

Siegel's dispute is not with the idea that thought processes satisfying the Canonical Reckoning Model constitute inferences—perhaps even be the “pinnacle of intelligent response” to one's informational-situation (2019, p. 17). She disputes instead whether the Canonical Reckoning Model, even with the weaker Self-Awareness Condition, formulates anything like a necessary condition for a succession of thoughts to constitute an inferring.

4 | A DILEMMA?

Siegel uses a series of well-crafted examples to persuade us that the Canonical Reckoning Model is too restrictive: there exist, she argues, many familiar thought processes that do not satisfy the conditions of the Canonical Reckoning Model and yet constitute responses to one's informational-situation that bear the marks (i)-(iv) and deserve to be called inferences. The Response Hypothesis acknowledges this, hence the slogan, “inference without reckoning”.⁹

Siegel is aware, however, that her appeal to examples, even if intuitively plausible, faces a fundamental objection from Canonical Reckoners:

Perhaps the most principled challenge to the Response [H]ypothesis is a dilemma. Either in inference one appreciates or registers the rational relationship between inferential inputs and conclusions (or purports to do so) in the form of a *reckoning state*, or else one's mind is merely caused to move from one state to another. If there is no such reckoning state, then the information state can make a causal impact on the thinker, but cannot make a rational impact. [p. 16]

Siegel wants to argue that this dilemma is a false one:

The picture of inference without reckoning allows a third option. It is possible to respond rationally to an information state without a reckoning state that represents what makes that response rational. [p. 16]

How is this possible? If:

One's acknowledgment of rational support *consists in the response*, rather than taking the form of a state that represents the support relation. [p. 17-18; emphasis added]

I agree that this is possible. However, I suspect that defending the possibility of “inference without reckoning” may require more by way of characterizing the processes involved than (i)-(iv) afford. In particular, we will need to be able to see how, at least in principle, a response innocent of any higher-order representation of the support relation between the antecedent and consequent thoughts—what we will call a *first-order mental transition* from one information-state to another—could *itself* embody an “acknowledgment of rational support”. This will, I claim, require advocates of the Response Hypothesis to venture

⁹It might be more accurate to say “inference without a reckoning-state”—we'll see why this matters, below.

more deeply into questions about mental “structure and components” than Siegel proposes to do here, but in a way that I think will help ground her overall project.

We can see this most clearly, I think, if we press on the question of how to distinguish Siegel’s “inference without reckoning” from nearby *non-inferential* non-reckoning responses to one’s information-state. Consider what might be Siegel’s simplest example.

5 | PUDDLE

Puddle, Siegel tells us, is an example of “inference without reckoning” that is simple enough to “not involve any more cognitive sophistication than what’s needed to ... keep one’s feet dry” (p. 5):

Puddle: ... while walking along a rainy street, Sin-yee might come to a puddle and think that it is too big to hop across, so she will have to go around it. She need not think to herself that *she has to walk around the puddle if she wants to keep her feet dry*. [p. 5; emphasis added]¹⁰

In Puddle, the italicized phrase would constitute the relevant “reckoning-state”—a state whose content *represents* the support relation between Sin-yee’s ends (getting home and keeping her feet dry) and the means she takes (walking around this particular puddle). Siegel is claiming, in effect, that Sin-yee’s grasp of her situation can be such that it makes walking around the puddle reflect an intelligible, rational response on her part, an *inference*, even though Sin-yee never formulates—implicitly or explicitly—a reckoning-state.

However, consider the classical Behaviorist, who will dispute the idea that Sin-yee’s behavior is to be explained inferentially. He elaborates: “Sin-yee has encountered puddles in the past and hopped. Because she likes having dry feet, she was negatively reinforced for jumping when the puddles were too large to clear—i.e., three or more feet in width, as it happened. It took her a while, but she learned from foot-soaking hops to be averse to jumping whenever she encounters a puddle wider than three feet. This does not require her to mentally imagine some remoter goal ‘in light of which’ trying to jump a puddle of this size is a poor idea—the aversion to jumping is *conditioned* directly to the stimulus condition of encountering a puddle of this size. Moreover, on those occasions in which, faced with a puddle wider than three feet, she instead walked around, Sin-yee was rewarded by being able to continue on her way home with dry feet, only slightly ‘punished’ by the extra steps and the lost fun of hopping. As with the aversion to jumping, this conditioned motivation to walk around is conditioned directly to the stimulus situation, without need for any thinking about the future. So now Sin-yee has acquired a well-calibrated stimulus-response disposition. She is happier, and her parents are happier. All this suffices to *establish* in Sin-yee the behavior of walking around puddles larger than three feet.”

On this interpretation, Sin-yee would clearly have enough “cognitive sophistication” to keep her feet dry, but would we say that she manages this by an *inference*, or simply by *conditioned association*?¹¹ After all, the point of the Behaviorist revolution was to purge psychology of the supposedly “pernicious mentalism” that led to positing unobservable “internal actions” such as inferences

¹⁰For expository convenience, I have substituted a proper name, ‘Sin-yee’, for ‘you’ in Siegel’s original framing of the example. And to keep with the idea that a high level of cognitive sophistication is not being demanded, let us imagine that Sin-yee is a kindergartener walking home from school.

¹¹Siegel herself *contrasts* purely associative transitions in thought with inferences, though the examples she considers constitute a subset of potential associative responses to one’s information-state (Siegel 2017, section 5.2.).

performed on private mental representations of non-existent future states. A scientific psychology, the Behaviorists argued, must stick to quantifiable observables and the Law of Effect, explaining Sinyee's behavior in terms of a history of reinforcement that yields well-calibrated stimulus-response dispositions that operate more like direct associative connections than mental inferences. Behaviorists were ambitious—they would seek to give a similar kind of explanation even of the more complex behavior in Siegel's other examples, eschewing any "internal actions" like inferring. Contemporary psychologists call this kind of reinforcement learning and behavioral control *model-free*, since it depends only upon acquired "cached values" (e.g., aversion or attraction) for responses in situations, with no reference to past experience, possible future outcomes, or general knowledge such as causal models of situations and actions. While contemporary psychologists have largely abandoned the full ambitions of classical Behaviorism, they recognize a place in animal and human behavior for such model-free learning and control processes.

Siegel concludes with a brief discussion of emotional and aesthetic responses which she sees as being "analogous" to "inference without reckoning" in that one can be ignorant of the features of situations or objects to which one is responding (and thus have no available "reckoning-state" to guide the response) while still making a response that is sensitive to relations of epistemic dependence and that renders oneself epistemically evaluable.

These kinds of emotional and aesthetic responses are arguably intelligent yet partly self-ignorant responses. In this respect, they are directly analogous to self-ignorant inference without reckoning. ... [I]nference without reckoning allows that inference can tolerate the kinds of self-ignorance described here. Whatever epistemic improvements might result from being able to pinpoint what one is responding to and why, in aesthetic, emotional, or rational domains, the initial responses one makes prior to any such attempt can still reflect the intelligence of the responder. [p. 18]

However, the ability to discern levels of *intelligence* in responders is not enough to make a case that *inferring*, understood as an exercise of potentially rational epistemic agency, is taking place. Model-free, associative responses still count as learning, and animals and humans can be better or worse at this kind of learning, e.g., in their sensitivity to, facility in acquiring, reliability in acting on, or proportioning of effort to, the reward-contingencies of given acts in given situations. Such responses are not "merely causal", in the way that the succession of thoughts in mental idling was hypothesized to be merely causal. We know from recent work in machine learning just how powerful model-free reinforcement learning can be—given enough data and processing power, it can yield highly sophisticated and successful behaviors in complex situations (Mnih et al., 2015). And animals, humans included, have a continuous stream of rich data about their internal states and the world around them, and a large number of neurons and synapses, so sophisticated and successful behaviors in complex situations would not be beyond their reach, either. Moreover, behaviors guided by model-free learning and control can be said to be successful in part *because* such systems track epistemically-relevant phenomena, e.g., learning to discriminate signal from noise, to identify the most informative cues in a given context, to calibrate responses to relative frequencies and magnitudes, etc. These discriminative abilities non-accidentally operate in ways that pattern on evidential relations, enabling model-free learners to solve a wide array of problems, and general problem-solving ability is one common way of thinking about intelligence.

The fact that model-free reinforcement learning can yield behavior that is responsive to evidentially-relevant features shows that the original dilemma presented on behalf of the Canonical Reckoner, above, is a false one, as Siegel argues. Model-free reinforcement learning represents a third option

distinct from self-conscious reckoning, on the one hand, and from “merely causal”, non-intelligent responses to information, on the other.

However, the Canonical Reckoner can reply that this could *weaken* rather than strengthen Siegel’s case, since it suggests that examples like Puddle can’t be used to “illuminate without analysis” the phenomenon of “inference without reckoning” on the ground that Sin-yea shows intelligent, adaptive behavior “without reckoning”, since, if the Behaviorist were right, the behavior might well be “without inference” as well. It would seem that we need to inquire more fully into the underlying “structure and components” of individuals’ responses to the situations in order to say whether these are, in fact, cases of “inference without reckoning” rather than the upshot of well-calibrated stimulus-response dispositions.

This strikes me as a legitimate rejoinder to Siegel—if we are inclined to see cases like Sin-yea’s as inferences, the Reckoner will say, that is probably because we are *already* assuming a good deal about the inferential character of the underlying “structure and components”. Indeed, Siegel’s descriptions of her two most extended examples contain mentalistic idioms that already assume that inference, or something close to it, is going on. In **Kindness** she speaks of “what leads you to [the] thought” or enables you to “reach the judgment” that the clerk is kind (p. 4). And in **Pepperoni** she speaks not only of responding to features, but of how “the conclusion of your inference would weaken” if those features were to change (p. 4).

Is there a way to proceed here that could clearly distinguish “inference without reckoning” from well-calibrated stimulus-response dispositions, while at the same time not presupposing that implicit higher-order representational states are at work underneath? This way of posing the question suggests a path forward: we might try to establish the possibility of *first-order* processes at work underneath, effectuating the connections needed to render implicit mental transitions of the kinds that figure in Siegel’s examples genuinely dependent upon epistemic support relations. If so, we might be able to make good on Siegel’s idea that, in “inference without reckoning”:

One’s acknowledgment of rational support *consists in the response*, rather than taking the form of a state that represents the support relation. [p. 18, emphasis added]

Of course, what would be involved is *implicit* acknowledgment, since we are looking for a form of inference in which explicit representations of rational support relations do not play a role. But for such acknowledgment to exist, it would seem that at least two conditions must be met: (a) we must be able to see how a first-order response could manifest or embody normatively-appropriate sensitivity to relations of epistemic support, even in the absence of higher-order representation of those normative support relations as such, however implicit; and (b) this first-order response would have to be more than “automatic” sub-personal calculation or mere association, and involve person-level activity of a kind that opens up the possibility of evaluating the process and the individual as more or less responsive to reasons. Finding *bona fide* mental processes that meet these two conditions would add a fourth category to the original (supposed) dilemma: between Canonical Reckoning and mental idling, there would be, in addition to well-calibrated stimulus-response dispositions, another form or forms of response that would, thanks to meeting (a) and (b), pass a certain threshold of eligibility for the label “inference without reckoning”.

One way to make headway on finding or characterizing mental processes meeting conditions (a) and (b) is to look for instances where there isn’t much temptation to posit implicit higher-order presentations of epistemic support relations, where first-order thought holds sway. A plausible candidate would be the psychology of those innocent of normative epistemic concepts, such as very young infants or non-human animals. So let’s ask: Why have practicing psychologists largely abandoned classical Behaviorism—whether about the behavior of the Sin-yees of the world, or of very young infants, or of the intelligent animals in their laboratories?

6 | *TU QUOQUE*: STATES, PROCESSES, AND SKILLS

But before we do that, it might be worth pausing to ask what is at stake in the question whether we can find first-order, person-level mental processes in which support relations play the kind of *because* role required if these processes are to count as inferences in an epistemically-evaluable sense. This might seem to be a narrow question, arising as a problem specifically for Siegel's project of vindicating "inference without reckoning". Instead it is, I believe, an instance of a very general problem of explaining how inference is possible at all.

To see why, and why this problem arises equally for the defender of the Canonical Reckoning Model, let's say that I am in an information-state *I* comprising antecedent-states with the content *p*, which as a matter of fact would epistemically support moving to an information-state *J* which includes as well the content *q*.¹² And suppose that I indeed make this mental transition from *I* to *J*, where the fact that I am in state *I* partially causes and explains the fact that I come to be in state *J*. With the *because* relation thus thinly described, however, this transition might be no more than mental idling, and so not inferential. Suppose, however, we add the constraint that the move to *J* must also depend in some way upon the *information content* of *I* and *J*. Still, this constraint on the *because* relation could be met by a purely associational transition from *I* to *J*, and, as Siegel points out, purely associational transitions are not inferential in the sense that interests us here.

The Canonical Reckoning Model comes in at this point to say that, in order to have the right *because* relation, we must add to my antecedent information-state a potentially self-aware reckoning-state *R* with the content *p supports q*, and this state and its content, must play a role in the transition to *J*. However, adding another *state* might not be to the point, since this would still be compatible with the move from the antecedent information-state *I + R* to the subsequent information-state *J + R* being merely associational—sensitive to information-content, but non-inferential. One solution that will not work would be for the Canonical Reckoner to answer by saying that the move from *I + R* to *J + R* becomes inferential because it involves a *because* role for a potentially self-aware recognition of the epistemic support relation between *I + R* and *J + R*. For then we would need to add a reckoning-state *S* with the content (*p & (p supports q) supports q*), playing a *because* role. And down this path lies regress.¹³

Now, the Canonical Reckoner will surely object. He might say: "That's the wrong way to think about it. Don't think of the role of a reckoning-state in inference as that of an additional *premise-state*—we all know where *that* leads. Think instead of the reckoning-state as playing a *shaping role* in the movement of mind, *guiding* the transition from the antecedent information-state to the consequent information-state, without entering as an additional *step* in the inference. That kind of *guided response* is what it *is* to be sensitive to the reason-making consideration provided by the reckoning-state's content."

That's a reasonable response on his part. But why isn't it a response Siegel can make with respect to the original movement of mind from *I* to *J*? That is, why couldn't Siegel say that *what it is* to be sensitive to the support relations between the contents of *I* and *J* is for the evidential, causal, probabilistic, means-end, etc. relations contained in *I* to *shape* and *guide* the mental transition from *I* to *J*? If the Canonical Reckoner is willing to entertain the idea that a mental transition—in this case, from *I + R* to

¹²For ease of exposition, we will be using as examples only cases where the content of the antecedent information-state actually supports the content of the consequent information-state.

¹³Siegel herself does not wish to rest her case for "inference without reckoning" on regress problems (Siegel 2017, p. 94; Siegel 2019, p. 9), but it's worth noticing that the problem of regress arises for the Canonical Reckoning Model as well in order that Siegel not bear an unfair burden of proof.

$J + R$ —can be an inferring if it manifests a sensitivity to the support relations contained in the contents of $I + R$, without a higher-order reckoning-state to stand behind and guide this transition, then the same should be true for the information about relations that sustain epistemic relevance in the contents of I in relation to J . Recall Boghossian’s inference about the streets. What guides the transition from (1) “It rained last night” to (3) “The streets are wet” is not a higher-order reckoning state mentioning epistemic support relations, but a causal generalization, (2) “If it rained last night, then the streets are wet”, that would, if true, make it the case that an epistemic support relation obtains between (1) and (3). If the worry of the Canonical Reckoner is that the inferrer might not understand *why* he is making the mental transition he does, then that worry would seem to be met by (1)–(3). Indeed, one might think that (2) does a better job of *explaining why*, or manifesting the inferrer’s *take* on the situation, or *understanding* of it, than would a reckoning-state. Typically, the antecedent information-state I in a given ordinary inference will contain a large number of relations—predictive, causal, means-end, etc.—upon which epistemic support relations supervene, and so be capable of providing the inferrer an appropriate *rationale* for drawing the conclusion he does in the circumstances. No commentary on the existence of epistemic support relations need be added—though, of course, making the inference attests to an implicit competence in, for example, using causal relations predictively (or retrodictively, or hypothetically, etc.) to guide thought and action, such as making a rational choice of shoes for the day.

To make a start in answering, a distinction is needed. We are sometimes tempted to think of inferences in terms of inputs and outputs—an antecedent information-state containing the premise-content goes in, and a consequent information-state containing the conclusion-content comes out. This makes it look as if the *response* to the antecedent information-state in which we are interested, and which we will be assessing for intelligence or rationality, is the consequent information-state. For example, when Siegel talks of “epistemic dependence”, it seems in some places that she has in mind epistemic dependence of the *conclusion*-content upon the *premise*-content. She offers, as

... the main diagnostic of inference: epistemic dependence. You could have better or worse reasons for the conclusion ... , and that would make the conclusion better or worse.
[p. 4]

But it might be more in the spirit of her approach to say that “you could have better or worse reasons for *your concluding*, and that would make *your concluding* better or worse”. That is, the response she ultimately is interested in is not the content of the conclusion independent of the process that generated it, but the whole consisting of the process and content. For example, after mentioning a number of *non-inferential* mental transitions—mental jogging, rhyming, association, etc.—she writes,

On the face of it, what’s lacking from these cases is a *distinctive way of responding* to the [antecedent]-state that produces the [consequent]-state. These transitions fail to be inferences, because they lack this kind of response. [p. 15; emphasis added]

Rhyming or association could deliver a consequent-state with a content that is in fact a legitimate conclusion from the content of the antecedent-state, but this would not constitute inferring, on her view. And the conclusion-producing processes she’s interested in are of a particular kind—they must be *actions* on the part of the individual responding that make epistemic evaluation of her and her thinking fitting. So that, if we are to make good on Siegel’s claim that, in “inference without reckoning”,

One’s acknowledgment of rational support *consists in the response*, rather than taking the form of a state that represents the support relation. [p. 18, emphasis added]

we need to be able to say what it would be for a mental action to be a process of a kind that constitutes an appropriate acknowledgment of support without needing to say this of itself. While that would be mysterious if we focused only on the content of the consequent-state, as if a consequence-state with content q that was produced by association would in itself be different from a consequence-state with the same content produced by inference. However, once we think of the individual's response as incorporating the process producing that content, the picture becomes clearer. To ask what it might look like for a process of thought transition to constitute "acknowledgement" of the relevant support relations, we will begin by considering how many psychologists distinguish inferential from associative thought processes.

7 | A PROTO-INFERENTIAL INFRASTRUCTURE

Within psychology, a pivotal factor in the rejection of classical Behaviorism as a research program was that even laboratory rats, to say nothing of humans, turned out to comport themselves in ways that did not fit the Behaviorist's associationist account.

Rats showed themselves capable of learning spontaneously by exploration without external reinforcement—relations were learned, and used to pursue goals effectively in novel situations, without conditioning. For example, rats allowed to explore a maze without reinforcement, and then trained by reinforcement to take a particular path through the maze, were able, when they encountered a newly-introduced barrier along their trained path, to spontaneously improvise an alternative path that took them efficiently to the location where they had previously found food (Schmidt & Redish 2013). Doing so required that they integrate information learned in separate contexts (e.g., free exploration and training) and be able to project this information into novel contexts and synthesize novel actions that effectively meet their goals. They were *leveraging* the information they acquired, not just accumulating it. The pioneering critic of classical Behaviorism, Edward Tolman, identified such learning and behavior control as "purposive" or "molar", to draw attention both to the fact that it is organized around goals and not merely <situation-response> pairs and the fact that it occurs at the level of the whole animal—drawing widely upon the rat's varied sources of information and capacities to act, in contrast to "local", behaviorally-stereotyped stimulus-response dispositions. In another example, rats trained in a maze, when they encountered the maze flooded with water, promptly swam to the food tray, a complex behavior that had never been previously deployed or conditioned (Tolman 1948). Using behavioral observation, Tolman hypothesized that rats form "cognitive maps" that integrate the results of spatial exploration and reward history into an intelligent guidance system that gives them considerable "autonomy" in relation to whatever particular training regimes they have undergone, and permitting them to "infer" relative location independent of the particular paths run (1948). Since Tolman, psychologists working at the neural level have assembled an impressive body of evidence that the expression "cognitive map" is not an anthropocentric projection. Patterns of neural activation in the hippocampus and entorhinal cortex as the rat explores a space register two kinds of spatial information, relational and grid-like. Systematic activation of these neural patterns takes place not only when the rat is moving in the maze, but when it is resting or sleeping, where repeated, stimulus-independent activations correspond to simulated trajectories in these "maps", including "short-cut" trajectories the rats have never previously traveled. Moreover, when rats are at choice points in the maze, activation in these "maps" spreads forward, ahead of the rats' actual location, sweeping the two branching paths and connecting this information with previous learning about reward magnitudes and frequencies—indeed, neural firing rates are found to correspond to the "expected value" of the available paths in light of the rat's previous experience. Once this expected value information is computed, it projects into motor-control circuits where comparison occurs and the rat chooses a path accordingly (Johnson et al.

2007). Moreover, “maps” are not merely passive in real-time behavioral guidance—the outcome of a given choice is continuously compared with prior expectation, and discrepancies function to update the rat’s “cognitive map” and adjust out-going motor commands (Redish 2016).

Such “map-like” learning, choice, and motor control is called *model-based*, in contrast to the model-free associative learning discussed above. Behavioral and neural evidence suggests that model-based learning and control in rats includes not only spatial relations and the distribution, magnitude, and frequency of rewards, but also degrees of uncertainty, causal connections, and social relations. Such a model-based architecture makes possible both *forward*, expectation-based guidance of choice and behavior and *inverse*, discrepancy-based guidance of updating. While rats themselves, as far as we know, do not have concepts like *<utility>*, *<probability>*, *<epistemic support>*, etc., their behavioral choice and motor control can be mathematically modeled to a good approximation by Bayesian inference and rational choice theory (Körding & Wolpert 2006).

Model-based architecture makes good sense for foraging animals, but especially for those who have complex needs and social relations and are likely to inhabit an environment that does not simply afford stable “reward contingencies” for behaviors, and who face strong metabolic limitations on computation. It is no fluke of evolution that highly-intelligent animals have this kind of mental architecture. Contemporary control theory suggests that, if we think of the brain as the “regulator” of the organism’s interactions with the environment in light of its goals, then achieving robust effectiveness and efficiency in this regulatory task requires the building of models (Conant & Ashby 1970). Whether or not we philosophers are happy to call the ways in which rats construct and utilize multi-dimensional mental models *inferential*, it seems to me less controversial that model-based regulation of cognition and behavior affords an example of how an information-state could encode spatial, causal, probabilistic, means-end, etc. relations in such a way that they can literally *guide* the animal’s responses in learning and behavior.

Model-based learning and control thus provide an entry-point for assessment not only in terms of *intelligent behavior*, but in terms such as *accuracy* of spatial, causal, and predictive representations, and *effectiveness* and *efficiency* in representing and evaluating alternatives in order to make a choice in terms of expected value rather than acquired habit. This is evidenced in the psychologists’ ability to ask and answer serious questions about how well the patterns of neural firings of rats (or other intelligent animals) approximate normative models of learning or decision-making, and how this could *explain* flexible, intelligent behavior in circumstances where associative accounts cannot. To be sure, the claim is not that rats represent such normative features as evidential or rational support *as* such, which would require normative concepts, but that they construct mental models that encode the relations upon which evidential and rational support relations supervene. This makes it the case that rat mental and behavioral responses are, for example, attributable at the whole-animal level to responsiveness to considerations of kinds we can identify as *evidentially-relevant* or *reason-giving*.

Such model-based architectures are helpful in thinking about how inference is possible without regress. Recall that what we needed, whether we are Siegelians or Canonical Reckoners, is a credible account of what it would be for an information-state that encodes the grounds of epistemic support relations to operate to *guide* or *shape* mental transitions in ways that are appropriately responsive to support relations without introducing a “premise” or “rule”—or “reckoning-state”—that would require a further inferential step. Mental models are constituted by *networks* of weighted, directed connections, and information travels the paths made by these connections in proportion to these weights. In a simple example, the weights and directions of connections in the model arise in response to frequencies and sequences of perceptual elements in past experience,

and new sensory inputs can enter such a network to yield a particular pattern of activation pattern that terminates in a perceptual identification (or a degree of confidence in a particular perceptual identification) of a persisting object. A “premise” or “rule” connecting these input features with the identification of an object of that kind does not play a role, so we need not *add* an inference to the process that, if we were psychologists, we might call “perceptual inference”. Further, this object identification can then enter a model of causal relations to yield expectations about the object’s subsequent behavior, which in turn can guide the individual’s subsequent motor control and learning from discrepancy. All of this is a continuous flow of information in which the antecedent information-state of the rat, prior to the original sensory input, provided a model-based structure that shaped and guided the rats’ responses throughout in line with how the rat “understands” relevant spatial, causal, predictive, reward, and control relations—how she “took” her situation to be and what she “took into account” in acting as she did. Even if no higher-order reckoning-state representing the epistemic support relations on view here was involved, we still can provide an answer to such questions as “Why did the rat make the choice it did?”, “What did she see in taking the right-hand path?”, “What is guiding her on-going behavior?”, etc. So, the rat’s response as a whole is not a mere relation of input to output, but a complex “take” on her situation and its prospects grounded in a large “knowledge structure” and set of competencies, and once we understand this take, we can see why it acts as it does.

Call such structured, information-sensitive, projective, thought- and action-guiding responses *proto-inferential*. Proto-inference gives us an idea of how a first-order mental process could, even in the absence of higher-order representations of support relations, nonetheless constitute a response to new information that non-accidentally embodies apt sensitivity to relations of evidential or means-end relevance “in light of” the individual’s antecedent information-state, fulfilling desideratum (a), above. Moreover, this response is not activity in a special-purpose, “automatic” sub-agential calculation or a “mere” well-calibrated disposition, but an “organism-level” activity of the rat “in light of” its overall “take” on the situation as bearing upon its goals, fulfilling desideratum (b). The rat has thus become evaluable in such epistemic terms as *intelligence*, *accuracy*, *sensitivity to relevant evidence*, *proportionality of expectation to evidence*, and so on. This does not make the rat’s mind the mind of a rational individual, since rational individuals can do more by way of representing, evaluating, and guiding the work of their own minds than is made possible by the first-order sensitivity to reason-making features embodied in a rat’s mental models. But it does show how an information-state *I* with a content *p* that epistemically supports a transition to an information-state *J* with content *q* could, in itself, *guide* the individual’s thought to *q* by those considerations contained in *I* that constitute the grounds of the epistemic support relation to *J*, and in the ways they are relevant to the content of *J*.

We now face the next step in understanding “inference without reckoning”, which requires us to see how this kind of proto-inferential structure could be embedded within a psychology capable of forming and using self-representations in ways that could constitute “inferences without reckoning” that qualify for, and qualify the individual for, evaluation as more or less rational. Richard Feynman’s autobiography recounts his days as a youth repairing radios (1985, p. 20). Once, when confronted with a puzzling form of noise, he did not open the radio but paced back and forth asking himself what might cause such a noise. It came to him that such a noise could result from tubes heating up in the wrong order, then he promptly opened the radio, reversed the tubes, and solved the problem. This made his reputation, as the owner of the radio reported to everyone, “He fixes radios by *thinking!*” How would an inherited proto-inferential system help Sin-yeek keep her feet and shoes dry inferentially, that is, by *thinking*?

8 | PUDDLE REVISITED

So let's return to Puddle. How would a model-based account of Sin-ye'e differ from the model-free account provided by the classical Behaviorist, and what might this say about Sin-ye'e as a potentially rational inferrer?¹⁴

On a model-based account, Sin-ye'e leaves school heading in a particular direction because she has the goal of being home and mentally maps the available pathways connecting school and home and their relative advantages or disadvantages. While she might take pretty much the same route every afternoon, so that the matter seldom receives much thought, she is quite capable of taking alternate routes if her goal changes to include stopping by a friend's or a store on the way or staying out of the hot sun, which could require mentally piecing together paths she's taken for different purposes in the past to synthesize an overall route she has never traveled. She can mentally simulate possible combined pathways before she has physically taken them, enabling her to notice possible advantages or disadvantages, or to compare her confidence that one or the other might work.

Let's suppose that Sin-ye'e inhabits a fairly arid part of the world, and so she has encountered only small, shallow puddles in the past, easily stepped over or tip-toed across. This morning, however, there was an unusually heavy rain and, as she heads home, Sin-ye'e encounters her first really large, deep puddle. Sin-ye'e has never had the problem of wet shoes and feet owing to failed puddle-hopping, but she once got her shoes and feet soaked on a visit to the shore, and did not at all like walking back to where her family was staying, squishing in her shoes, or being scolded by her parents for the damage caused to her shoes. So her capacity to model situations causally includes such possibilities. Playing hop-scotch has given Sin-ye'e an idea of how far she can jump, and walking around obstacles has given her an idea of what that involves. Though she need not explicitly make any calculation, as she stares at the puddle her proto-inferential capacities quickly simulate and evaluate the possibilities of hopping vs. walking around, and she opts to walk around. Should we call this an inference on her part, given that she does not "think to herself that *she has to walk around the puddle if she wants to keep her feet dry*", as Siegel puts it?

So far, one might say, Sin-ye'e's mental activity would seem to just what one would observe in her furry proto-inferential animal relatives in like circumstances—leveraging previous learning to simulate, evaluate, and compare options at a choice point, such that motivation to act shifts accordingly. We don't want to be speciesist—is there a *principled* difference between the two cases that would warrant our saying that Sin-ye'e infers, or manifests rationality, while the rat does not? There is, I believe, and it has nothing to do with Sin-ye'e being human in particular, and nothing to do with forming a reckoning-state. It has to do with her overall standing as an agent, her capacity for acting intentionally.¹⁵

Suppose we come across Sin-ye'e, stopped staring at the edge of the puddle, and ask her the Anscombean question, "What are you doing, stopping there like that?" She might reply, "I don't know, I'm just standing here" or "Have I stopped? I hadn't noticed." But that is very unlikely. More likely she'll say, "I'm thinking." Sin-ye'e, in other words, is not just host to a bout of proto-inference—she is thinking *intentionally*, under an idea of what she's doing, not fully articulated but still there, as the Anscombean question is thought to reveal. The guiding idea is something like *figuring out what to do*

¹⁴Debate continues over the relative proportion and importance of model-free vs. model-based mental processes in intelligent animals and humans (Dayan & Berridge 2014). So our contrast is more precisely put as between an agent capable of both model-based and model-free thought, vs. one capable only of model-free thought, like the Behaviorist's Sin-ye'e.

¹⁵I am assuming for the sake of argument that rats, even though their behavior is guided by expected-value assessments, are not "acting under an idea" in the manner of intentional action in humans. This could be unduly uncharitable—rats keep surprising us.

about the puddle given that she wants to get home with dry feet. This idea gives structure and a success condition to what she is doing, and converts her thinking from unself-conscious proto-inferential *activity* into potentially self-conscious inferential *action* in which she is trying to utilize the information she has to solve a problem posed by the situation, and how well she does this will “redound” to whether she can do this better or worse. For contrast: Imagine that Sin-yee has been told by her parents and teachers not to *stoop* so much when she walks and stands. If we were instead to have asked Sin-yee, standing staring at the puddle, “What are you doing there, stooping like that?”, she’d very likely reply, “Am I stooping? I hadn’t noticed.”

We can further ask Sin-yee, “Thinking about what?”, and she’ll likely have a ready reply, “The puddle.” “And what are you thinking about it?” “It looks too big.” “Too big for what?” “Too big to jump over, of course.” “So what will you do about that?” “Well that’s what I’m *thinking about*, silly.” Sin-yee has learned how to focus her attention to a situation in a way that engages her proto-inferential capacities, and how to sense when they’ve been exercised enough, or how confident she should be of their working. In this way her inferential capacities are like other skills she’s acquired. She is able to deploy her acquired competence in “figuring out” in much the same way as she is able to deploy her acquired competence in language. When she speaks, she speaks intentionally and grammatically, a form of person-level action even though she would be at a loss to say how she uses her implicit model of the grammar of the language so effectively, or what it consists in. Still, the relations in that model can suitably guide her speech, and she can be aware of whether a given sentence “sounds right” or “sounds wrong”, even as she is saying it. Let’s say that she speaks a different language at home from the language she speaks at school. Sin-yee then can intentionally employ one language competence rather than another, not simply by responding to the immediate context—she can enjoy using her school language at home to share secrets with her siblings that her parents can’t understand—but by adapting choice of language, language register, etc., to her goals. Sin-yee’s intentional control of her speech, and the credit she deserves for speaking well, or slyly, or aptly, does not require that she be able to say why one vocabulary or sentence-form was chosen rather than another. None of these processes are purely sub-personal or opaque—we can, by asking the right questions, unearth a number of important elements of her thinking and linguistic and social competence. Her thinking skill likewise. She often will be able to supply appropriate answers in the form of considerations that *did* play a role in her thinking, because her thinking was model-based, and these considerations were integral parts of those models. “Why do you care whether the puddle is too big to jump?”—“Because I don’t want to get my shoes wet.” “Why care about whether you get your shoes wet?”—“Because it’s yucky walking with wet shoes.” “I thought you liked hopping?”—“I do, but it’s worse to have wet shoes.” Moreover if Sin-yee were halfway around the puddle, only to discover that the ground there is very muddy and the puddle past that point much shallower than she thought, this discrepancy will catch her attention, and she will be able to focus on her new situation in a way that mobilizes other relevant experiences and abilities, and delicately tiptoe across. And if we were to ask, “Why are you tip-toeing like that?”, she’d have an answer.

Sin-yee is able to see what is being demanded by our questions because of her capacity to model herself as well as the world around her, allowing her to enter into critical dialogue on questions of about why she is doing what she’s doing, and join the realm of self-reflective and potentially rational agents, much to the delight of those around her. But, as before, her apt responsiveness has to do with her ability to supply conscious direction to her model-based *figuring out* capacities, and their ability to deliver intelligent, reliable, evidence-sensitive responses owing to her proto-inferential capacities. What she needs if she is to be skilled in figuring out is that her information-state itself have enough structure and content to encode the spatial, causal, probabilistic, means-ends, etc. relations upon which the support for, and guidance of, her *figuring out* depends. And that is what a model-based

infrastructure provides, in which she can use her higher-order representational capacities to ask herself the questions we are asking her.

Sin-ye'e's capacity for inferring that bears features (i)-(iv), I've been arguing, depends upon her capacity to impart intentional direction to her proto-inferential capacities—like Feynman, she can act under the idea of *figuring something out*, and do this with increasing skill as she matures. This intention, like any intention, binds together the various elements of her *figuring out* behavior, and explains why *figuring out what to do about the puddle* is *her* response to the situation before her, rather than something that simply happens to her.

Recently, however, Paul Boghossian has argued that we cannot invoke intention to explain what makes a mental transition an inference. After all, doesn't acting on an intention already involve inferring? Boghossian makes this point forcefully against an intentional reading of following a rule:

Intentional Construal (IC): On this Intentional construal of rule-following, ... my actively applying a rule can only be understood as a matter of my grasping what the rule requires, forming a view to the effect that its trigger conditions are satisfied, and drawing the conclusion that I must now perform the act required by its consequent. ... [Thus] on the Intentional view of rule-following, rule-following requires inference. [Boghossian (2014, p. 13)]

Translating to the case at hand, an (IC) account of what it would take for Sin-ye'e to be *figuring out what to do about the puddle* would have to be understood in terms of her grasping what is to be figured out, reaching a view to the effect that a given course of thought would be instrumental to carrying out that task, forming the intention to do this, and applying that intention in action. Plainly, we have just inserted further inferential processes into the purported explanation of inferential processes. This is worse than no headway—we are regressing. Boghossian writes, “I am now inclined to think [that this difficulty] is at the heart of Wittgenstein's discussion of following a rule” (2014, p. 12). Rather than give up altogether on the idea that we are rational animals capable of genuine inference, Boghossian suggests that we take rule-following to be an unanalyzable primitive (2014, p. 17).

However, the picture of acting intentionally offered by (IC) cannot constitute the *basic* understanding of what it is to act intentionally. To be sure, sometimes we form express intentions, or ask and determine what they require, or decide to set them in motion, etc. But all of these are themselves intentional actions, so we clearly haven't gotten to the bottom of things. If acting intentionally is *ever* to be possible, then it must be possible to *act intentionally* without having to *form and execute an intention* to do so. And Sin-ye'e can intentionally think about what to do about the puddle without forming an intention to do so—she stopped to think abruptly, but not accidentally or unwittingly or against her will. As her answers to our questions reveal, her thinking was guided by the idea of getting home and her preference for dry shoes and feet. No formation of an intention to stop, or to deliberate, or to conclude her deliberating, need occur.

The moral of the model-based approach to cognition and control is that one can be intelligently guided by one's “take” on a situation—one's spatial and causal modeling, one's preferences or aims, the feedback one receives in the course of acting, etc.—in a continuous process, without engaging in distinguished component actions like “forming an intention”, “applying an intention”, etc. An idea of what one is seeking and why can guide one's mental and physical activity through control and feedback in such a way as to constitute purposive, intentional action, without positing a string of internal actions to carry this out. And that had better be the case, if acting intentionally is to be possible at all.

Reviewing: Sin-ye'e's *figuring out what to do about the puddle* bears the relevant marks, (i)-(iv). It is person-level mental action that results in reaching a conclusion in response to the individual's

information state, and that redounds well or badly on the individual in part according to whether the contents of the antecedent states epistemically support the contents of the conclusion. The “right sort” of *because* relation for inferring, if her case is typical, involves intentional and not merely mechanical causality—it exists when the movements of the mind are organized under and guided by the individual’s implicit or explicit aim of *trying to figure something out*, even if her knowledge of the antecedent information that will be brought to bear or the support relations linking this information to her conclusion is limited. We might not be good at introspecting the inner workings of most of our figurings out—ordinary causal inference, for example, draws upon a vast amount of acquired information and multiple simulations. But, at least by Sin-yee’s age, we are pretty good at directing our minds to a particular task of figuring something out. Just as we can mobilize the rich resources of our perceptual system for interpreting something in our visual field by directing our attention to it and focusing in on it, we can mobilize the rich resources of our inferential infrastructure and current information-state for figuring something out by directing our mental attention to it and focusing in on it. In neither case do we need to know—nor could we know, in general—just *why* we reach the conclusion we do, so the Self-Awareness condition is not fully met. But the intentional character of the *figuring out* gives it an epistemic success condition, with the result that it becomes liable to epistemic assessment and makes us liable to assessment for rationality.

9 | BEYOND THE PUDDLE

We have spent this long on Puddle, since Siegel needs only to show that there are some examples of “inference without reckoning” that count as inferences. The other examples she offers (both in Siegel 2019 and in chapter 5 of *The Rationality of Perception*, Siegel 2017) raise other, interesting issues. Hide-and-Seek finds the searcher intentionally trying to *figure out where the hider is hiding* by bringing to bear her proto-inferential capacities to model not only the situation and its possibilities, but the hider’s mind as well. Kindness finds the person waiting in line forming implicit conclusions about the character of the clerk simply by focusing her eyes and attention upon the clerk—and smooth eye-movement and attention are, the vision scientists tell us, person-level voluntary actions (Lisberger 2010). And Pepperoni finds the pizza-eater in a quandary because her expectations, based upon her self-model of her taste, appetite, etc., have been violated, bringing her to focus her mind upon the problem of *figuring out what’s wrong with this pizza*, which then intentionally deploys her multi-dimensional proto-inferential causal modeling capacities *inversely*, to find possible causes for an observed effect, of which the “yuckiness” of the pepperoni seems initially most plausible.

In these examples, we see epistemic agents acting epistemically, though not under that description and not through reckoning-states. Instead, they are deploying their proto-inferential competencies to leverage the information they have, forward or inversely, under the idea of *figuring something out*. *Figuring out*, like any intentional action, can be done implicitly or explicitly, well or badly, and rationally or irrationally. Without the underlying proto-inferential architecture of the mind we inherit (and which is itself explained by very general facts about successful learning and control), there would be no competence to answer to these intentions to figure out. Perhaps an underlying model-free or associationist architecture could yield the appearance to ourselves or others of figuring out by inferring, but it would be an appearance only. And without the competence in intentional action—including such mental actions as inferring—that characterizes the mind of potentially-rational agents (and which is itself explained by very general facts about the nature of rational choice and action), we would be limited in our deployment of our model-based architecture to the kinds of proto-inferences made by our

furry relatives. Either way, we'd be intelligent animals. But only with the two factors combined do we have a chance of being rational animals. Like Sin-yeec.

10 | CONCLUSION

We have been trying to make a case that “inference without reckoning” can be genuine inference. Whether Siegel will welcome this defense is another matter. It saves “inference without reckoning”—indeed, it shows how “inference *with* reckoning” might be possible without regress—but at the cost of saying more about underlying mechanisms that Siegel would appear to want to say. At the same time, the possibility of alternative, e.g., associative, explanations of how examples like Puddle, Hide-and-Seek, Kindness, and Pepperoni work leaves the defender of “inference without reckoning” little choice but to wade into questions about process. Not *ad hoc*, but on the basis of highly general considerations about the possibility and nature of cognition, choice, control, and action.

REFERENCES

- Boghossian, P. (2014). What is inference? *Philosophical Studies*, 169, 1–18.
- Broome, J. (2014). Comments on Boghossian. *Philosophical Studies*, 169, 19–25.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Sciences*, 1, 189–197.
- Dayan, P., & Berridge, K. C. (2014). Model-free and model-based Pavlovian reward learning: Revaluation, revision, and revelation. *Cognitive, Affective, and Behavioral Neuroscience*, published online.
- Feynman, R. (1985). *Surely You're Joking, Mr. Feynman!*. New York: W.W. Norton.
- Johnson, A., van der Meer, M. A., & Redish, A. D. (2007). Integrating hippocampus and striatum in decision-making. *Current Opinion in Neurobiology*, 17, 692–697.
- Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10, 319–326.
- Lisberger, S. G. (2010). Visual guidance of smooth pursuit eye movements: Sensation, action, and what happens in between. *Neuron*, 66, 477–491.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518, 529–533.
- Redish, A. D. (2016). Vicarious trial and error. *Nature Reviews: Neuroscience*, 17, 147–159.
- Schmidt, B., & Redish, A. D. (2013). Navigation with a cognitive map. *Nature*, 497, 42–43.
- Siegel, S. (2017). *The rationality of perception*. New York: Oxford University Press.
- Siegel, S. (2019). Inference without Reckoning. In M. Balcerak-Jackson & B. Balcerak-Jacksoni (Eds.), *Reasoning: Essays on theoretical and practical thinking*. New York: Oxford University Press.
- Tolman, E. (1948). Cognitive maps in mice and men. *Psychological Review*, 55, 189–208.

How to cite this article: Railton P. Inference, Inferring, and the Rationality of Perception. *Philos Phenomenol Res.* 2020;101:735–754. <https://doi.org/10.1111/phpr.12735>