10-1-2015

How You Know You are Not a Brain In a Vat

Alexander Jackson
Boise State University

How You Know You are Not a Brain In a Vat

Alexander Jackson
Boise State University

Abstract

A sensible epistemologist may not see *how* she could know that she is not a Brain In a Vat (BIV); but she doesn't panic. She sticks with her empirical beliefs, and as that requires, believes that she is not a BIV. (She does not inferentially base her belief that she is not a BIV on her empirical knowledge—she rejects that 'Moorean' response to skepticism.) Drawing on the psychological literature on metacognition, I describe a mechanism that's plausibly responsible for a sensible epistemologist coming to believe she is not a BIV. I propose she *thereby knows* that she is not a BIV. The particular belief-forming mechanism employed explains why she overlooks this account of how she knows she is not a BIV, making it seem that there is no way for her to know it. I argue this proposal satisfactorily resolves the skeptical puzzle.

Keywords: skepticism; metacognition; brain in a vat; BIV; Moore's argument

1. The Puzzle Facing Sensible Epistemologists

C.S. Peirce wisely commanded, "Let us not pretend to doubt in our philosophy what we do not doubt in our hearts" (Peirce 1992 p. 29). With that in mind, please answer the following question.

> Are you a massively deceived Brain In a Vat (BIV), whose misleading experiences are produced by a computer simulation run by an evil scientist?[1]

Hopefully you did not suspend judgement on whether you are a BIV, and so are now suspending judgement on all empirical matters. Freaking out in that manner is not a sensible response to considering the skeptical scenario. Hopefully you did not suspend judgement on whether you are a BIV, incoherently combining that attitude with judging that you have hands, it is a sunny day, etc.[2] Hopefully you kept your empirical beliefs, and added the belief that you are not a BIV.

Every sensible person who considers the issue believes herself not to be a massively deceived BIV; I assume the reader is in the club. A philosopher who suspends judgement on whether she is a BIV is like a patient suffering from Obsessive Compulsive Disorder who cannot leave the house because she keeps returning to check she really did turn off the stove. The patient with OCD doubts her memories of checking the stove; her doubts are pathological, not epistemically virtuous.[3] Similarly, it is pathological to really suspend judgement on whether the skeptical scenario obtains and hence on whether one interacts with real people who should be treated with moral respect, and with some of whom one has a loving relationship. Suspending judgment on those matters has crazy consequences for decision-making. (Consider what one is willing to do in a computer game such as *Grand Theft Auto*.) Plausibly, if one believes that one might be a BIV then one suspends judgement on whether one is.[4] So if a contextualist about "knows" and related words holds that there are special contexts in which one should accept that one "might be a

[1] Equivalently: are you the plaything of Descartes' evil demon, systematically misled by the experiences the demon induces in you?
[2] Vogel (1990) and Hawthorne (2004) defend the view that it is incoherent to suspend judgement on whether one is a handless BIV while judging that one has hands (for example). Dretske (2005) and Nozick (1981, ch. 3) dissent.
[3] The example of OCD comes from de Sousa (2008).
[4] This is the majority view in the literature on epistemic modals; see Yalcin (2011) and the other papers in that volume. The dissenting minority includes Dougherty & Rysiew (2009).

BIV", she is recommending temporary insanity.[5] Some people find it easy to *say* they might be BIVs; their attention should be directed towards the consequences for their attitudes towards what appear to be other people, and then to Peirce's maxim with which we began. (Sometimes one needs to pull oneself together, and not let a feeling of anxiety overwhelm one. On the way to the airport, I have to tell myself that I don't need to check again that I have my passport. The same goes if one starts to worry that one might be a BIV.)

Sensible epistemologists have already settled that they are not BIVs. Yet the skeptical puzzle continues to tax them. So the puzzle does not concern whether one is a BIV. I suggest the puzzle is to say *how one can know* one is not a BIV. It can seem that there is *no way* to know that one is not a BIV—all the suggestions in the literature seem wrong. But that result is unacceptable: believing that one cannot know that one is not a BIV rationally commits one to suspending judgement on whether one is a BIV, and hence to skepticism. I will focus on the puzzle about knowledge; but the same problem arises as to how one could rationally believe that one is not a BIV. (The following footnote sharply separates the puzzle of BIV skepticism from the so-called 'lottery puzzle'.[6])

This paper aims to solve the puzzle here *facing sensible epistemologists*. To do so, it is enough to give an account *we* can contentedly endorse of *how* one knows that one is not a BIV. Such an account can assume that one is not a BIV, that one has reliable perceptual faculties, and so on. The problem with existing accounts is that they seem implausible *to we sensible epistemologists*; the problem is not that they will be rejected by someone who thinks she might be a BIV.[7] This paper does not aim to change the mind of someone who believes she might be a BIV (and is not resisting a strong inclination to believe she is not). Such a person is not to be persuaded by presenting an argument; she is to be treated in the same way as patients with OCD—medically. The paper should change the mind of someone who is strongly inclined to believe she is not a BIV, but feels forced not to because she sees no way she could know such a thing.

The skeptical puzzle is hard because it seems illegitimate to use one's empirical knowledge as evidence that one is not a BIV—which does not leave much to appeal to. In particular, the 'Moorean' solution seems unacceptable. Sensible epistemologists agree that one has perceptual knowledge that it is a sunny day, for example. For all p, p logically entails it is not the case that: p is false and one is deceived on that score because one is a BIV. The Moorean tells one to *infer* that logical consequence from what one knows by perception.[8] But most of us find the recommended inference repugnant. We are told that one should first settle by perception that it is a sunny day, and that matter having been settled satisfactorily, then infer it is not the case that: one is a BIV and it is not a sunny day. But it seems that one has not settled satisfactorily that it is a sunny day if one leaves it open at that point in inquiry that one is a BIV deceived on that score. It is not appropriate to use one's perceptual beliefs as an inferential basis for other beliefs until one has answered that skeptical worry. In other words, the Moorean says that one can know one is not a BIV *epistemically posterior* to one's perceptual knowledge; but that seems false to most of us

[5] Note that contextualists can say that even in the sense operative when considering the skeptical worry, one "knows" one is not a BIV. (See Cohen 1999 esp. p. 85 n. 27, and DeRose 2004.) They would then need to say how that can be so—they need a theory of the kind I propose in this paper. Lewis' (1996) contextualist theory is naturally taken to imply that one should sometimes suspend judgement on whether one is a BIV.

[6] We don't react to the BIV scenario by dropping our empirical beliefs. By contrast, the so-called 'lottery puzzle' concerns cases in which considering a certain worry *does* cause us to drop a belief. For example, I stop (flat-out) believing that I will teach next Tuesday if I consider the possibility that I will sick that day and have to cancel. The puzzle arises because most of our beliefs are vulnerable in this way to some carefully chosen worry. Does this show that people really know very little? Or do people know a lot, and they make a mistake in reversing their verdict? Is there a third diagnosis, according to which the initial belief and its subsequent dropping are both 'appropriate'? Contextualism, subject-sensitive invariantism and relativism are theories of the third, 'shifty' kind. (For more on the lottery puzzle, see Vogel 1990 and Hawthorne 2004.)

As observed at the start of this footnote, the institutions that generate the lottery puzzle are different in kind from those that generate the puzzle of BIV skepticism. There is no reason to presume the two puzzles have a common solution. I am sympathetic to contextualist or relativist solutions to the lottery puzzle, but reject such solutions to BIV skepticism. (The lottery puzzle comes up again in footnote 38.)

[7] Pryor (2000 pp. 517-8) puts the point by saying that we should abandon the 'ambitious anti-skeptical project' of using only premises a skeptic would grant, and focus on the 'modest' one, using any premises *we* accept. Amongst others, Nozick made this point too (1981, pp. 197-8; see p. 262 of the reprint in Sosa et al. 2008).

[8] E.g. Pryor (2004, 2012); Davies (2004); Tucker (2010).

sensible epistemologists. (Saying the Moorean inference 'begs the question' strikes me as ambiguous between the objection just made, and the irrelevant point that *skeptics* will claim the inference is illegitimate because one doesn't know the premise is true.) Some philosophers have tried to explain away the apparent illegitimacy of the Moorean inference, but I am not convinced.[9] This means Mooreanism is unable to explain the pull of skepticism, as well being implausible.

A second problem for Mooreanism is that its being a sunny day *does not* entail that one is not a BIV. That it is a sunny day doesn't even make it likely that one is not a BIV. So the Moorean inference does not even purport to be a way to know that one is not a BIV (as opposed to knowing it is not the case that: one is a BIV and it is not a sunny day). But one needs a way to know one is not a BIV. It would be incoherent to form a perceptual belief that it is a sunny day while thinking, "Maybe I'm a BIV, but at least I'm not a BIV who is wrong about whether it is a sunny day!" So Mooreanism is, at best, a radically incomplete theory of how to answer skeptical worries.[10] This is not the place to explore possible responses to this objection, though I think Matthew McGrath (2013) shows the prospects are bleak.[11] I just want the objection on the table, as it won't even be a prima facie problem for my proposal (§3.2).

I've described two problems for the Moorean account of how one can know one is not a BIV. There are other accounts in the literature, including 'abductivism'[12] and Crispin Wright's 'entitlement' theory (2004); I don't find them plausible either. I worry that all the existing accounts are unsatisfying. That is, they bite a bullet: they rest on some claim that still seems false after the theorist has had a go at explaining our mistake. This paper presents a new account—one that makes the initial mistaken intuition go away. The rest of this section outlines the proposal, and indicates how I will go about fleshing it out.

Every sensible person who considers the skeptical scenario comes to believe—in the same kind of way—that she is not a BIV. I suggest that everyone who has considered and dismissed the skeptical scenario in the normal way thereby *knows* that she is not a BIV. Some ways of coming to believe that one is not a BIV, such as on the basis of a coin toss, are not legitimate. The way sensible epistemologists come to believe that they are not BIVs *is* legitimate—it produces knowledge.

§2 describes a plausible mechanism by which sensible epistemologists come to believe they are not BIVs, drawing on empirical psychology. For now, let me just convey the flavour of the view. A sensible epistemologist does not infer she is not a BIV, settling that matter epistemically posterior to settling by sense perception that it is a sunny

[9] Pryor (2004 §§5-7, 2012 §V) and Davies (2004) try to explain the mistake we allegedly make in rejecting the Moorean inference. Pryor observes that if one is unreasonably suspicious that one is a BIV, or that certain epistemologies of perception are correct, then one won't be rationally able to form perceptual knowledge, and so the Moorean inference won't be available. But sensible epistemologists harbor no such unreasonable suspicions; so the source of our repulsion from the Moorean inference has not been identified. Davies observes that one can't make the Moorean inference when one is trying to 'settle the question' of whether one is deceived. But if one assumes that one knows one is not a deceived BIV, and wonders merely *how* one knows it, then Davies endorses the Moorean inference. But sensible epistemologists do assume they know they aren't BIVs; so Davies has not identified the source of our typical anti-Moorean reaction.

[10] I don't think Moorean can respond that one must inspect one's body, and infer that one is not a BIV, before one can form any other perceptual knowledge. A principled Mooreanism must be grounded in a 'liberal' epistemology of perception, i.e. one which holds that perceptual knowledge does not automatically rest on any independent knowledge (Pryor 2000, 2004, 2012).

[11] Pryor (2004 p. 374 n. 33) suggests that having a perceptual experience that p makes one justified in believing *that one knows that* p. Michael Huemer (2011) suggests that humans are so built that when it appears to one that p, it will usually appear to one *that one knows that* p. On Huemer's view, the latter appearance makes one justified in believing that one knows that p. So Pryor and Huemer could respond to the objection to Mooreanism by saying that one can infer from one's knowledge *that one knows it is a sunny day* that one is not a BIV. I argue in §2.2 that this response does not yield a psychologically plausible account of how people actually come to believe that they are not BIVs. Nor does it allow us to satisfactorily explain away the pull of skepticism (cf. footnote 9 and §4). McGrath (2013 §5) gives other potent objections.

[12] Beebe (2009) includes an overview and references.

day (for example). Nor does she settle that she is not a BIV epistemically prior to settling it is a sunny day (as conservatives such as Wright demand). Rather, she settles that she is right that it is a sunny day, and that she is not a BIV, in one go—at the same point epistemically speaking.

Having described a belief-forming mechanism with this flavour, I refine the following account of when it produces knowledge. Suppose that—setting aside how she actually deals with the skeptical worry—a sensible epistemologist meets the conditions for knowing it is a sunny day. Those conditions also make it the case that, by judging that she is right it is a sunny day and that she is not a BIV 'in one go', she knows she is not a BIV. Her belief that she is not a BIV is a good one *because* the relevant perceptual belief has certain good-making features. Knowledge that she is not a BIV 'comes for free' with her perceptual knowledge that it is a sunny day.

I embed my account of how one knows one is not a BIV in a general account of how one knows one is not going wrong in a specific way (such as by misremembering, or mis- inferring). §2 presents a general psychological account of how one typically judges one is not going wrong. I draw on empirical work on metacognition (the monitoring and control of one's thinking); a unified picture emerges from work on memory, inference, and perception. My account of how a sensible epistemologist judges she is not a BIV is an instance of the more general account I give. §3 gives a general epistemological theory of when it is rational to judge in the relevant way that one is not mistaken, and when one thereby knows one is not mistaken. Thus the account of how one knows one is not a BIV is psychologically and epistemologically principled, not ad hoc.

I develop the psychological account in §2, layering the epistemological account on top in §3.1. §3.2 distinguishes the proposal from some others in the literature. I indicate why it implies that one settles that one is not a BIV neither epistemically prior nor epistemically posterior to settling matters by perception. §4 argues the proposal should resolve sensible epistemologists' philosophical puzzlement. Sensible epistemologists become puzzled because they overlook the idea that they came to believe they are not BIVs in exactly the right way, and thereby came to know it. My psychological account explains why the oversight is so natural and common. So my account explains away the pull of skepticism.

I am confident that the key to resolving the skeptical puzzle is to explain how one comes to believe that one is not a BIV without being inclined to judge that one thereby forms knowledge. The psychological explanation I propose is speculative, however. I trust you won't let doubts about the details obscure the appeal of the strategy.

2. The Psychology of Answering Worries

This section gives a psychological theory of how people generally come to judge that they are not mistaken about a particular matter. I summarize a popular picture of how people generally come to judge that they are right about a given matter (§2.1), then speculatively extend the picture to cover judgements that one is not making a given kind of mistake (§2.2). I apply this general theory to say how a sensible epistemologist comes to judge she is not a BIV (§2.3).

Let's put some background in place. One need not believe things to be the way they seem to one to be. How this happens is one of the topics psychologists address under the heading of metacognition. The phenomenon is widespread: one can distrust the deliverances of perception, memory, and inference.[13] We can say that in such cases, one has a perceptual, mnemonic or inferential *seeming*, without forming a *judgement* with the same content. (I come back to the nature of seemings in §2.1.)

Psychologists think that humans—and some other animals—can distrust a seeming without judging that they are going wrong: such a conceptual representation of oneself, or the seeming, is not necessary.[14] But our concern is with how one comes to judge that one is not mistaken, and hence with cases in which the subject *does* represent that possibility. The subject considers whether she might be going wrong in a particular way, and appreciates that answering that she might be so mistaken commits her to suspending judgement on the content of the relevant

[13] Michaelian (2012) explains how this kind of metacognition increases one's reliability.
[14] See Proust (2010; 2013 esp. chapter 5), and Carruthers (2011 chs. 9.4-9.5).

seeming (assuming there is no other reason to believe it). Either the subject judges that she might be mistaken and hence suspends judgement on the content of the seeming, or she judges that she is not mistaken and affirms the content of the seeming in judgement. Call this process "self-conscious metacognition".[15]

My working example will be a case of 'propositional' or 'semantic' memory. Suppose the pub quizmaster asks Amy what the capital of Mongolia is. It then seems to Amy that Ulan Bator is the capital of Mongolia, though she doesn't recall how she learnt that fact, nor any supporting evidence. Prompted by her team-mate, Amy asks herself whether she is right, or whether she might be misremembering on that score. The process of asking and answering that question is an instance of self-conscious metacognition. If she answers that she might be misremembering, she is committed to suspending judgement on the content of the mnemonic seeming. In fact, Amy judges that she is not misremembering. (Misremembering, i.e. the degradation of stored information, is one kind of worry about memory. Another worry is that things went wrong in the initial acquisition of the information, such as that one acquired it from a misprint in a magazine, or on the say-so of an unreliable friend. I intend my theory to apply to both kinds of worry.)

2.1 Metacognition and FORs

How does Amy come to judge that she is right that Ulan Bator is the capital of Mongolia, and that she is not misremembering? I start my answer by sketching some popular claims from the psychology literature on metacognition. This literature tends to concentrate on the case of memory; Asher Koriat (2007) provides an excellent survey. Valerie Thompson argues persuasively that the metacognitive framework developed for the case of memory applies to the case of inference too (2009; also Thompson et al. 2011, 2013a). Koriat (2011) applies the framework to certain visual judgements, such as which of two objects is larger. Like Joelle Proust (2013) and Jennifer Nagel (2012), I will suggest that the framework should be applied to all the standard ways of forming a judgement.

According to the psychologists we'll look at, answers retrieved from memory are accompanied by a Feeling Of Rightness (FOR) of a certain strength. 'FOR' is Thompson's terminology; Koriat calls it a "sheer subjective feeling" of confidence (2007 pp. 301, 314). The idea is introspectively compelling. My memory that London is the capital of the UK is accompanied by a very strong Feeling of Rightness. My memory that Lima is the capital of Peru is accompanied by a weaker Feeling of Rightness—weak enough to be called a Feeling of Doubt.[16]

Officially, let's say that Amy tokens a mental representation with the content that Ulan Bator is the capital of Mongolia, and a strong FOR attaches to that token representation. (This will seem natural when describing how FORs are generated, below.) I will sometimes abbreviate this by saying the relevant *content* Feels Right. Plausibly, a seeming just is a token mental representation that has a FOR attached. So we can say that it seems to Amy that Ulan Bator is the capital of Mongolia.[17]

This view of seemings supports Jack Lyons' account of the relation between seemings and judgements (Lyons 2009, pp. 70–2, 90–2). Plausibly, a judgement just is a token mental representation that plays a certain functional role in guiding action and belief-formation. So a token mental representation can count as both a seeming and a

[15] Hilary Kornblith convincingly attacks over-intellectualized theories of this process, which he calls 'reflection' (2012 esp. chapter 5).

[16] One might worry that cognitive life is not always so full of feelings. The following reply is inspired by Whittlesea & Williams' (1998) work on feelings of familiarity. They point out that we only have the feeling that someone is familiar or unfamiliar when that fact is surprising. Analogously, we could say that people only have strong Feelings of Rightness when there is a need for them, such as when they consider worries about their judgements.

[17] Tucker (2013 §1) surveys the disagreement in the epistemology literature about the nature of seemings. In Tucker's taxonomy, my proposal is an 'experience view' of seemings. It is good to make seemings a mental state psychologists talk about.

judgement.[18] On this account, judging that p when it seems to one that p does not clog one's head with two token mental representations that p; one such representation is sufficient.[19] In such a case, I see no harm in saying that the *judgement* Feels Right.

Typically, the subject has no introspective access to what's responsible for the strength of a FOR. Psychologists come up with clever experiments that confirm theories about what's responsible. The standard view is that features of the retrieval process causally determine the strength of the FOR attached to a retrieved memory.[20] For example, the ease of retrieval from memory affects the strength of the accompanying FOR. The details don't matter for our purposes; what matters is that subjects typically can't say why a certain content Feels Right.

A strong FOR typically causes a subject to affirm the relevant content in judgement, treating the question as settled. A weak FOR causes a subject to suspend judgement, which can prompt further investigation. For example, I might ask a friend with a smart phone to google the capital of Peru to check I'm right; I'm not going to ask someone to check that London is the capital of the UK—that matter is settled.[21]

A strong FOR is typically produced, and secures a judgement, in a way that does not causally depend on beliefs about one's competence, such as that one's memory is good. Koriat distinguishes metacognition involving FORs, which he calls "experience- based", from that involving prior beliefs about competence, which he calls "information- based" or "theory-based" metacognition (Koriat 2007 pp. 295–301, 313–4). Judging usually involves experience-based, and not information-based, metacognition. (It is controversial whether work surveyed by Alter & Oppenheimer (2009) shows that beliefs about the truth-conduciveness of 'feelings of fluency' always play a role in judgement. As I explain in the following footnote, I am skeptical.[22] Regardless, what matters for what follows is that all that's open to introspection is that one has a strong FOR. The role (if any) of beliefs about feelings of fluency is not introspectible, and hence won't be scrutinized as part of self-conscious metacognition.)

---

[18] In my view, nonconceptual perceptual representations can play the functional role needed for them to count as judgements, i.e. occurrent beliefs. (One can stipulate another technical sense in which only conceptual representations count as 'beliefs'.) A full defence of this claim would take us too far afield. Note however that the claim is needed to give a unified account of the metacognitive regulation of perception, memory and inference (see below).

[19] The terminology can get confusing. Its seeming to one that p, and one's judging that p, are two different token 'mental states'. A person tokens those two mental states by tokening a single representation that has two features (namely: it has a FOR attached, and it has plays a certain functional role).

[20] The empirical literature refers to the non-introspectible processing which produces a FOR as 'metacognition'; that processing must be distinguished from what I've called self-conscious metacognition.

[21] Consider again patients with OCD who cannot leave the house because they keep returning to check they really did turn the stove off. Plausibly, the problem is that these patients have weak FORs accompanying their mnemonic representations, i.e. feelings of doubt rather than strong FORs (de Sousa 2008). This characterization, according to which OCD is a matter of inappropriate *feelings*, is supported by the fact that the standard treatment is an anti-depressant, Prozac (de Sousa 2008 pp. 197-8).

[22] Many kinds of processing go into reading and answering a question, and each kind can be more of less fluent. The ease of reading the question, which is affected by the font used, is a case of 'perceptual fluency'. 'Answer fluency' is the case with which an answer is retrieved from memory or inferred. Alter & Oppenheimer (2009) think that both kinds of fluency affect a single feeling. They say feelings of fluency affect confidence via one's 'naïve theory' about what such feelings indicate. For example, Unkelbach (2007) examines how people answer 'hard' questions like: "The first Olympic gold medalist (new games) was James Connolly—true or false?", and 'easy' questions like: "Aristotle was a Japanese philosopher—true or false?" Perceptual fluency plays a role in answering 'hard' questions: people are usually more likely to say that a statement is true if it is written in a font that is easy to read. However, appropriate explicit beliefs cause people to take fluency to count against a statement's truth, rather than for it (as can special training). So beliefs about what's indicated by fluency play a role in educated guesses. But Unkelbach also found that manipulating the legibility of statements has no effect on how people answer 'easy' questions, to which they presumably retrieve the answer. Thompson et al. (2013a, 2013b) found that perceptual disfluency causes subjects to think longer about answers they infer, but does not reduce their confidence. Answer disfluency does reduce confidence. So Thompson et al. deny that the varieties of processing fluency result in a single feeling (e.g. 2013b p. 258). Rather, perceptual fluency affects a "global sense of (un)ease", while answer fluency affects the FOR produced. In their view, FORs do not need to be 'interpreted' by a naïve theory to affect confidence.

Thompson (2009) argues that 'dual-process' theories of reasoning should posit FORs in the case of inference. According to dual-process theories, we should distinguish fast, intuitive and automatic kinds of reasoning from slow, analytic and controlled kinds; these are called 'type 1' and 'type 2' processes respectively.[23] For example, consider the following questions from Shane Frederick's (2005) Cognitive Reflection Test (CRT):

> A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost? _____ cents
> If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 widgets? _____ minutes

Most people have the answers '10 cents' and '100 minutes' leap to mind when addressing these questions; those answers are produced by paradigmatic 'type 1' processes. Most people—including most undergraduates at elite universities (Frederick 2005)—accept those answers. However, people sometimes make the effort to slowly and sequentially calculate the right answers, namely 5 cents and 5 minutes. Such calculation is the paradigm of 'type 2' reasoning.

Thompson wants to explain why so many people are taken in by the cognitive illusions in the Cognitive Reflection Test. More generally, why does a given person sometimes go along with the response produced by a type 1 process, and sometimes override such a response in favour of engaging in type 2 reasoning?[24] Part of Thompson's answer is that type 1 reasoning produces a representation with a FOR of a certain strength attached. Usually, a strong FOR causes the subject to judge the relevant content to be true, omitting type 2 reasoning into the matter. She says (2009 p. 175): "It is this … Feeling of Rightness (FOR), that is the reasoner's cue to look no further afield for the answer." Evans & Stanovich (2013a, 2013b) endorse Thompson's postulation of FORs. Thompson and collaborators (2011, 2013a) corroborate the role of FORs experimentally. They also show that the fluency with which an answer is inferred affects the strength of the attached FOR.

A strong FOR typically causes the subject to believe the content of the representation to which the FOR attaches. Further, if the subject considers whether she is right about the relevant matter, a strong FOR will typically cause her to believe she is right that p. Thompson calls the result a Judgement of Rightness (JOR) (2009 pp. 181, 186). A JOR is a conceptual representation, with the content that (for example) one is right that Ulan Bator is the capital of Mongolia.[25] A FOR is not a judgement, but rather "an affective response" (Thompson 2009 p. 181). If one asks oneself the question, a strong FOR accompanying one's seeming will typically cause one to judge that one is right, i.e. form a JOR. So says Thompson:[26]

> In most cases, it is assumed that the FOR will be a sufficient basis for judgement, such that one's JOR is completely determined by the strength of the FOR with little, if any, conscious effort. (Thompson 2009 p. 181)

[23] Evans (2008) surveys the domains in which the 'type 1/type 2' terminology has been invoked. On pp. 207-1, he warns against running different distinctions together under those headings, as do Evans & Stanovich (2013a, 2013b). In particular, a process that uses working memory need not be intentionally controlled.

[24] Stanovich emphasizes that some people are more likely to engage in type 2 reasoning than others (2011 esp. ch. 2). People have different 'thinking dispositions', which are not well correlated with scores on standard IQ tests, and hence are not themselves a matter for 'algorithmic' type 2 processing (pp. 36-7). This important point does not undermine the need to posit FORs to explain *of a particular person* why she sometimes relies on a type 1 response and sometimes doesn't.

[25] Maybe JORs represent that a particular seeming is right, rather than representing oneself as right about a particular matter. As a reviewer pointed out, this might make a difference to the mental state concepts a creature needs to form JORs. I don't think this issue makes a difference to our main topic, which concerns how creatures that can represent scenarios according to which they suffer misleading seemings come to judge that such scenarios do not obtain.

[26] Koriat (2007) claims a 'subjective feeling' of confidence is often casually sufficient for a judgement about how likely it is that one's answer is right. Presumably people often think their answer is definitely right (else they wouldn't actually believe their answer). So Koriat's claim supports Thompson's.

Admittedly, a strong FOR does not always cause one to believe the content of the relevant seeming, or form a JOR if one is considering whether one is right: beliefs about one's competence can also cause a subject to make or withhold the relevant judgements. We saw that Koriat allows for both 'experience-based' and 'information-' or 'theory- based' metacognition. Thompson draws a similar distinction between 'implict' and 'explicit' metacognitive cues:

> In sum, it is proposed that the JOR, like other memory-based metacognitive judgments, is multiply determined by both implicit and explicit cues. Implicit cues [i.e. FORs] are based on properties of the retrieval experience, such as its fluency; explicit cues are derived from beliefs that are accessible to conscious introspection. Note that, as is the case with decisions based on heuristic outputs, metacognitive decisions may be based on implicit cues, even when a more accurate judgment could be derived from explicit sources. (Thompson 2009 p. 183)

The important point for our concerns is that it is perfectly normal for someone to judge that they are right about some matter, based solely on a strong FOR.[27]

This subsection has summarized an attractive and empirically well-supported picture of memory and reasoning. It is attractive to extend the claims about FORs to the case of perceptual appearances too. Daniel Kahneman (2003) emphasizes the similarities between type 1 reasoning and perception (they both produce 'impressions', i.e. seemings). Koriat applies a common metacognitive framework to visual judgments that one picture covers more area than another (2011), to memory (2012), answers to moral questions (Koriat & Aviv 2011), and to personal preferences (Koriat 2013). We see the same kind of metacognitive phenomena in the cases of memory, reasoning and perception, which suggests that a uniform account is appropriate. Proust (2013) and Nagel (2012) take such a view. For the rest of this paper, I will assume that the account sketched here of FORs and JORs applies to perceptual seemings as well as mnemonic and inferential seemings.

2.2 A Two-Alternative Forced-Choice Task

§2.1 gave an account of how people come to judge that they are right about a given matter. This subsection extends the story to say how people judge that they are not misremembering (or mistaken in some other specified way). Unfortunately, the psychologists don't address this question[28], so we will have to go off-piste. My proposal is very much in the spirit of §2.1. Plausibly, the job of strong FORs is to secure assent, subjects treating the relevant matters as settled without further investigation. It seems part of this job that strong FORs attached to an answer cause subjects to dismiss unsupported worries, such as that they are misremembering, if they are considered.

Koriat and Thompson do not talk explicitly about people who consciously ask themselves whether they are misremembering. I will assume that considering the possibility that one is making a certain kind of mistake does not always rob a strong FOR of its power to produce assent. Amy can consider whether she is misremembering, and judge she is right that Ulan Bator is the capital of Mongolia, dismissing the worry, when all that's consciously present is the strong Feeling of Rightness accompanying her answer. Looking ahead: considering the BIV scenario does not reduce one's confidence either. Of course, if it appears to Amy that her seeming is in danger of being misleading, she won't judge the relevant content true. But in the case we are talking about, she merely considers the proposition that her seeming is misleading, without it appearing to her that there is a danger it is. (It is not relevant that Amy would be less confident were she to think about things differently; see notes 6 and 38 on the lottery puzzle). (The following footnote addresses the objection that considering a proposition containing a negation, such as that one is *not mis*remembering, causes one to engage in type 2 processing, and thus the original FOR loses its potency.[29])

---

[27] Ernest Sosa says that 'reflective endorsement' of a belief must draw on a web of beliefs about the world and one's place in it (2009 pp. 151-2, 239-43). Assuming that 'reflective endorsement' is the same thing as forming a JOR, Sosa overlooks that JORs can be experience-based rather than information-based.

[28] Koriat concurs (p.c.).

[29] This objection trades on conflating characteristics that used to be associated under the banner of 'type 2 processing' (cf. note 23). Deutsch et al. (2009 p. 445) conclude that processing negations uses working memory, but "may occur very quickly and unintentionally even without extended and consistent practice of the negated expressions." Hence, they say, "negation processing is probably best characterized as a semi-automatic process." There is no support here

What causes Amy to judge that she is not misremembering, given that a strong accompanying FOR is what causes her to judge that she is right that Ulan Bator is the capital of Mongolia? Let me start by arguing against the view that she *infers* that she is not misremembering, from the fact that she is right (i.e. from a JOR). Set aside the worry that people *should not* reason in that way (just as the Moorean inference seemed illegitimate in §1). The view under consideration is psychologically implausible. It predicts that a strong FOR will accompany the inferentially produced representation that she is not mistaken. Hence Amy would feel she has proved her conclusion, just as she would if she infers that Ulan Bator is not in Africa, from the fact that it is the capital of Mongolia. But, I suggest, that's not how it feels to judge that one is not mistaken. To see this, let's consider a variant of our working example.

Suppose someone asks Amy whether she was deceived by an unreliable testifier when she originally formed the belief that Ulan Bator is the capital of Mongolia. Amy judges that she is right about the matter, and so judges she was not originally lead astray; but it will not feel as if she has proved she wasn't. She will not have the strong FOR predicted by the view that she infers she wasn't deceived. What will Amy say if challenged as to how she knows that she was not originally deceived? The view under consideration predicts that she will say that she is right that Ulan Bator is the capital of Mongolia, from which she can infer that she was not deceived when she formed the belief. But that's not what she will say. She will be baffled as to how she knows she was not originally deceived (she can't remember anything about how she acquired the belief!). She will try to ignore the question, insisting that she's right that Ulan Bator is the capital, and that she was not deceived when she acquired that information. (If she tries to give reasons for thinking she wasn't originally deceived, they will not be her real reasons for believing; they will be confabulated, mere rationalizations. Compare Nisbett & Wilson 1977; Evans 2008 pp. 258–9; Haidt 2001; Carruthers 2011 p. 39–40 and ch. 11.)

It is not simply obvious to Amy that she was not deceived when she originally formed her belief, nor can she say how she knows she was not. I think the explanation is that when Amy judges that she was not deceived, no FOR attaches itself to *that* representation. (A seeming is a representation that is accompanied by a FOR of a certain strength; so Amy does not have a *seeming* that she was not originally deceived.) The FOR that attaches to the mnemonic seeming causes Amy to treat it as settled that Ulan Bator is the capital of Mongolia. In particular, that FOR causes her to judge she was not originally deceived. But because no FOR attaches to the representation that she was not deceived, she will not automatically judge that she just knows it, and that it's obvious. Rather, she will cast around—without success—for an account of how she knows it that appeals to familiar elements such as memory and inference. The rest of this subsection spells out how Amy comes to judge that she is not mistaken, without a FOR attaching to that representation. I come back to the effects of such judgements in §4, where I argue one becomes puzzled as to how one could know one is not a BIV because that judgement does not itself Feel Right.

The key lies in how Amy sets up the problem she addresses. In considering the worry that she might be misremembering, Amy frames a Two-Alternative Force-Choice (TAFC). (As I explain below, TAFCs are a pervasive and well-studied part of cognition.) One of Amy's options is to judge that she is right that Ulan Bator is the capital of Mongolia, i.e. to make a JOR. Her other option is to refuse to judge that she isn't misremembering, i.e. to allow that she might be misremembering. The choice is 'forced', because Amy takes the two options to be exhaustive and mutually exclusive. The options are exhaustive, because she will not reject both of them: it makes no sense for her to judge she is not misremembering, but suspend judgement on whether she is right that Ulan Bator is the capital of Mongolia (assuming misremembering is the only kind of error she is considering). She rightly takes the two options to be mutually exclusive; setting up the decision-task in this way constitutes her respect for the prohibition on judging she is right while suspending judgement on whether she is mistaken.

If Amy's task is a Two-Alternative Force-Choice, then we can see how Amy comes to judge that she is not misremembering without a FOR attaching to that judgement. The only thing that affects Amy's decision is the strong FOR attached to the mnemonic representation. That FOR tilts the scales decisively towards making a JOR and away from refusing to judge that she is not misremembering. The output of Amy's decision task is both the picking of the former option and the rejection of the latter. That is, the output is *both* a JOR *and* a judgement that she is not misremembering. No FOR is generated that attaches to the judgement that she is not misremembering.[30]

for the claim that considering a negation generally shifts one away from responding 'intuitively'. This should be obvious: a paradigmatic 'intuition' is that the subject in a Gettier case does *not* know.

[30] FORs earn their keep here. The reader would be puzzled if I merely said that Amy judges she is not misremembering without it seeming true to her.

The proposal looks plausible in the light of other kinds of Two-Alternative Force- Choice. There is reason to think similar mechanisms produce one's answer, and one's level of confidence in it, in perceptual TAFCs (such as saying which of two objects is larger), general knowledge tasks, and value-based choices (such as choosing one of two possible free gifts).[31] So it is reasonable to draw comfort from the following analogies between Amy's choice and other kinds of TAFC.[32]

Value-based decision-making is a matter of choosing between options, and at least some of the time, one picks one option and rejects the others *in one go*. Suppose one has difficulty picking between two job offers. Eventually one comes to prefer the first job to the second. It would be strange to think the preference can only directly cause one to intend to accept the first job, which then causes one to intend to reject the second. Rather, the output of the decision-task can be both the intention to accept the first job and the intention to reject the second. The two intentions are outputted in one go.[33] Analogously, I suggest that the task of answering Amy's self-conscious metacognitive question outputs two judgements in one go: that she is right, and that she is not misremembering.

The suggestion sits happily with the biologically-motivated models of perceptual and value-based TAFCs described by Bogacz et al. (2006 esp. pp. 702–8; 2007). In these models, two neural populations constitute 'decision units' that represent the options to be chosen between. The mean level of firing of neurons across such a population constitutes the 'level of activity' of the decision unit. Both decision units starts at a baseline level of activity. Inputs to the decision increase the activity of one decision unit *and* inhibit the activity of the other. (On some models, inputs directly inhibit one decision unit; on other models, they do so indirectly, because the activity of one decision unit inhibits that of the other.) When the activity of one of the decision units reaches a certain threshold, a decision is reached. When one decision unit reaches the threshold, the other will have been inhibited down to the baseline level of activity. It is biologically realistic to talk of the scales tipping away from one option and towards the other. The end-state of the decision- making process represents both the picking of one option and the rejection of the other. There's no impediment to the direct output of the decision-task including both the acceptance of one option and the rejection of the other.

To summarize: Amy frames a Two-Alternative Forced-Choice, where the options are judging she is right that Ulan Bator is the capital of Mongolia, and not judging she is not misremembering. The strong FOR attached to the memory settles the matter: it causes the decision-task to be resolved by forming a JOR *and* judging that she is not misremembering. Let's introduce some terminology for this cognitive procedure: she *judges that* she is not misremembering *as an accompaniment to judging that she is right that* Ulan Bator is the capital of Mongolia. When

[31] Bogacz et al. (2006) review and compare models of the mechanism by which people perform perceptual choices. Bogacz et al. (2007 §5) extend one of those models to value-based choice. They assert that perceptual and value-based choices "involve a common selection mechanism" (2007 p. 1664), though that kind of mechanism is implemented in different regions of the brain in the two cases (p. 1669). Koriat gives a unified treatment of degrees of confidence in perceptually based answers (2011), answers to TAFC general knowledge questions (2012), answers to moral questions (Koriat & Aviv 2011), and in matters of personal preference (Koriat 2013). Koriat (2011) directly addresses arguments that different mechanisms are employed in perceptual and general knowledge TAFCs.

[32] The analogy has limits. Amy's choice differs in some ways from a general knowledge TAFC, such as whether Lincoln was Republican or a Democrat. Firstly, one could suspend judgement on both options for Lincoln's affiliation, even if one is forced to write down an answer; but one cannot stably suspend judgement on whether one might be mistaken—either one leaves the possibility open or one doesn't. That is, a skeptical worry presents one with a forced choice in judgement, not just in what one writes down. Secondly, one can resolve whether Lincoln was a Democrat or a Republican by remembering that he was a Republican, and can then infer that he was not a Democrat. By contrast, one cannot first judge that one is right, and then infer that one is not mistaken (assuming one lacks independent information about one's success). Such an inference is forbidden—see §3.2.

[33] A value-based decision can be affected by whether one focuses on picking one option, or on rejecting one (Wedell 1997). One might think that in these cases, the two intentions are not outputted in one go. That's not a problem for the analogy given in the text, as long as there are value-based choices in which one focuses equally on picking one option and rejecting the other.

it is obvious which JOR is being referred to, I will abbreviate by saying that Amy judges *as an accompaniment* that she is not misremembering. I conjecture that for any belief-forming faculty, such as perception and inference[34], humans will usually judge that they are not mistaken in this way.

This psychological theory makes a prediction. We should not expect someone to be able to articulate much of a defence of a judgement that she is not mistaken, when she judges it as an accompaniment to the relevant JOR. She might believe introspectively that the relevant claim Feels Right to her; but such a belief cannot do the causal (or epistemic) work of the feeling itself. We can only expect the following kind of insistence from Amy: "I *am* right that Ulan Bator is the capital of Mongolia; I am not misremembering." This prediction seems correct. (It seems correct for perception and inference too.)

2.3 Application to the BIV worry

I claim that a sensible epistemologist *judges that* she is not a BIV *as an accompaniment to judging she is right that* she has two hands, it is a sunny day, etc. More carefully, let me distinguish three cases.

Firstly, someone considering the BIV worry could fixate on one empirical claim, such as that it is a sunny day. Such as case is strongly analogous to Amy's. The subject frames a forced choice, in which forming a JOR, and allowing she might be a BIV, are the two options. Nothing significantly speaks in favour of her being in a danger of being a BIV. The strong FOR attached to the perceptual representation causes the decision task to be resolved in favour of forming a JOR and judging she is not a BIV. She thus judges that she is not a BIV without a FOR attaching to that representation.

Secondly, someone might take the BIV worry to present indefinitely many forced choices of the kind just described—one for each empirical belief she has. Such a subject can't consider every such decision task, but she can consider some examples of empirical contents that Feel Right, understanding that such examples can be multiplied indefinitely. (If a particular content doesn't strongly Feel Right, she moves on to other examples.) In each case, the subject judges that she is not a BIV as an accompaniment to the relevant JOR. She thus appreciates there are indefinitely many ways for her to judge she is not a BIV. She might express the resulting mental state by saying: "I am right that I have two hands, I am right that it is a sunny day, *and so on*; I am not a BIV."

Thirdly, I want to tentatively allow that the sense that one's empirical beliefs form a highly explanatory system to be a partial cause of one's judging that one is a BIV. In such a case, the subject frames a forced choice between judging that she is right about a plurality of matters, which form an explanatory web, versus allowing that she might be a BIV. The FORs attaching to the individual judgements are partial causes of the decision task being resolved in favour of a JOR for the explanatory web; a Feeling of Explanation (a 'sense of understanding') adds to the attraction of the system. (On the 'sense of understanding', see Gopnik 1998; Trout 2007; Lombrozo 2012.)

In summary: when sensible epistemologists dismiss the skeptical worry, they judge they are not BIVs as an accompaniment to certain JORs. Given an appropriately framed forced choice, the strong FOR of the relevant empirical content causes the decision task to be resolved in favour of forming a JOR and against allowing that one might be a BIV. In Koriat's terms, the belief that one is not a BIV is an 'experience-based' metacognitive judgement, not an 'information-based' one. As predicted by the general theory, a sensible epistemologist will not be able to say much in defence of her belief that she is not a BIV. She will only be able to insist, "I am right about these empirical matters; I am not a BIV."

[34] The case of inference requires a tweak. One can be totally confident that one is not inferring fallaciously, without being totally confident in one's premises and hence one's conclusion. So the FOR attached to one's belief in the conclusion does not explain the confidence that one is not mis-inferring. Rather, one addresses the possibility that one infers fallaciously by reasoning from the *supposition* that the premises are true, rather than the *belief* that they are. The forced choice task is then between judging under the scope of the supposition that one is right about the conclusion, versus judging that one might be inferring fallaciously. The| FOR attached to the conclusion under the scope of the supposition will be strong enough to explain the confidence that one is not inferring fallaciously.

3. The Epistemology of Answering Worries

§2 gave a psychological theory of how humans generally come to believe that they are not mistaken about a certain matter, and applied that theory to the case of believing that one is not a BIV. §3.1 says when such beliefs constitute knowledge, and when they are rational. §3.2 distinguishes the resulting proposal from some others in the literature. §4 explains why the resulting view should dissolve sensible epistemologists' philosophical puzzlement.

3.1 The Proposal: Knowledge by Accompaniment

I propose that when one judges that one is not mistaken as an accompaniment to a JOR, one can thereby *know* that that one is not mistaken. Specifically, suppose that—setting aside how she actually deals with the worry she is misremembering—Amy meets the conditions for knowing Ulan Bator is the capital of Mongolia. Those conditions also make it the case that, by judging she is not misremembering as an accompaniment to judging that she is right that Ulan Bator is the capital, she *knows* she is not a misremembering.

Consider the following underlying facts: Amy originally acquired the belief that Ulan Bator is the capital of Mongolia in a good way; she stored it normally; she now lacks significant reason to think she is in danger of misremembering (i.e. she lacks 'defeating' beliefs); and she judges that she is not misremembering as an accompaniment to judging that she is right that Ulan Bator is the capital of Mongolia. In virtue of those facts (and maybe some others), two epistemic facts hold: her JOR constitutes knowledge; and her judgement that she is not misremembering constitutes knowledge. Call this *knowledge by accompaniment* that she is not misremembering. It comes for free with her knowledge that Ulan Bator is the capital of Mongolia. Of course, Amy can know in this way that she is not misremembering, without knowing how she knows it.

I endorse a related account of what makes the relevant judgements rational. Having a seeming with a strong FOR, and lacking significant reason to think one is in danger of going wrong, makes it rational to judge that one is not mistaken as an accompaniment to forming the relevant JOR. (The case of inference requires more subtlety.[35]) It obscures things to suppress reference to the approved *way* of judging, simply saying that a seeming with a strong FOR makes it rational to judge that one is not mistaken. Such a formulation invites the question of *how* its seeming that p could make it rational to believe some other content. By contrast, there is no pressing question as to *how* a strong FOR makes it rational to resolve the forced choice by judging that one is not mistaken as an accompaniment to a JOR. The normative principle gives the thumbs-up to a kind of psychological operation that intuitively deserves it.[36]

Arguably, the view I am sketching is incompatible with a certain version of evidentialism. Evidentialism is the view that rational judgements are all based on good evidence. The particular version to be considered holds that (when the FOR is strong) a seeming is good evidence for the content of the seeming (e.g. Feldman & Conee 2001). Arguably, the analogous explanation for why it is rational to believe one is not mistaken is not plausible: it is not plausible that a mnemonic seeming is *evidence* that one is not misremembering. Some might take this to be a problem for my proposal, but I take it to be a problem for this version of evidentialism. I hope that by the end of the paper, my proposal is very attractive; whereas I argue in Jackson (2011) that this version of evidentialism can't be coherently motivated.[37]

---

[35] My (2011) argues that inferential seemings are evaluable for rationality, unlike perceptual and mnemonic seemings. So when one infers fallaciously, the inferred belief, and the belief that one is not inferring fallaciously, are *not* all-things-considered rational. As I explain in that paper, it is still true that having an inferential seeming with a strong FOR, and no defeating beliefs, 'rationally requires' believing the fallacious conclusion.

[36] One might worry about taking it as bedrock that a strong FOR has a variety of normative consequences. Expressivism about rationality helps one relax about such matters. I approve of the normal ways of responding to a strong FOR, and so give them the thumbs up. (I favour the way Fine 2001 formulates expressivism.)

[37] Teaser: the ethics literature distinguishes *having a reason* to act in a certain way from being in a mental state that *rationally requires* so acting (e.g. Broome 1999). I think that having a mnemonic seeming with a strong FOR, and no defeating beliefs, rationally requires believing the content of the seeming. There is no reason to go further and say that the seeming is *a reason* to believe the content, or *is good evidence* for such a judgement.

(The reader might be distracted by the following thought. One knows one misremembers from time to time; so it seems one cannot know that this is not such an occasion. So, allegedly, we should not give a theory explaining how one *does* know one is not misremembering. I respond in the following footnote.[38])

Let's apply the general epistemological theory to the BIV worry. I claim a sensible epistemologist has knowledge by accompaniment that she is not a BIV. She lacks significant reason to think she is in danger of being a BIV; she now judges that she is not a BIV as an accompaniment to forming JORs corresponding to some of her empirical beliefs (in one of the ways described in §2.3)[39]; and those empirical beliefs are well-formed (setting aside the proper ruling out of the skeptical scenario). In virtue of those underlying facts, two epistemic facts hold: the JORs constitute knowledge that she is right about the relevant empirical matter; and her judgement that she is not a BIV constitutes knowledge. (If at least one of the relevant empirical beliefs constitutes knowledge, then she has knowledge by accompaniment that she is not a BIV. The following footnote explores how to be generous towards someone who happens to consider only empirical contents she doesn't know, when judging that she is not a BIV.[40]) Judging that one is not a BIV is just a side-effect of holding the empirical beliefs one should; the epistemic standing of doing so derives from its connection to those empirical beliefs. I also endorse the corresponding account of what makes it rational to judge that one is not a BIV as an accompaniment to a JOR.[41]

One might object that strong FORs cannot be what make it rational to believe that one is not a BIV, because a BIV would also have strong FORs. This objection is clearly mistaken. What makes it rational to believe one is not a BIV *must* be something one has in common with a BIV, because it is also rational for a BIV to believe it isn't one. (Further, this objection smacks of the evidentialism I was happy to reject earlier in this subsection.)

---

[38] This objection presses the lottery puzzle (explained in footnote 6). Let me make three points in response. Firstly, there is no counterexample here to the proposal for knowledge by accompaniment. To press the lottery puzzle is to argue that normal adults *do* have significant reason to think they are in danger of misremembering, i.e. they have 'defeating' beliefs. If that claim is granted, then the account of knowledge by accompaniment correctly predicts that one does not know one is not misremembering. Secondly, the lottery puzzle simply distracts us from what's at issue with BIV skepticism; for no such argument can be made that adults have significant reason to think they are in danger of being BIVs. Thirdly, suppose that widespread skepticism is not the right response to the lottery puzzle, but rather we need a contextualist or relativist account of knowledge. Then pressing the lottery puzzle changes the context in a way that prevents us from addressing legitimate and interesting questions about how one knows various things, including that one is not misremembering.

[39] If one only notices the skeptical challenge to a single belief one fixates on, one can still know one is not a BIV. However, it is epistemically better to *understand* the full scope of the skeptical worry that one might be a BIV. It is epistemically better to appreciate that each of one's empirical beliefs can play the same role in rejecting the possibility that one is a BIV.

[40] We can distinguish occurrent from dispositional knowledge. For example, someone who never considers the question still has dispositional knowledge that she is not a lemur. If someone happens to use an empirical belief that doesn't constitute knowledge to dismiss the skeptical worry, then maybe she lacks occurrent knowledge that she is not a BIV but has dispositional knowledge she isn't. This is especially plausible in the case of a subject who is aware she has indefinitely many routes to judging she is not a BIV, but only considers empirical beliefs that don't constitute knowledge. Suppose someone thinks, "I am right that Sydney is the capital of Australia, I am right that Toronto is the capital of Canada, *and so on*; I am not a BIV." If the subject has some *knowledge* she is disposed to take as an instance of the "and so on", then she dispositionally knows she is not a BIV. Maybe she even knows it occurrently: the case might be analogous to Warfield's (2005) examples of coming to know something by deducing it from a falsehood. (One of his cases: one miscounts and judges there are 53 people in the room instead of 52, and deduces there are fewer than 100 people present.)

[41] One might wonder whether a sensible epistemologist can judge she is not a BIV as an accompaniment to JORs corresponding to *any* of her empirical beliefs, or only to a subclass of them. Candidates for the relevant subclass include: her perceptual beliefs; her non-inferential empirical beliefs (including beliefs stored in memory); her non-inferential empirical beliefs plus those that constitute an explanatory understanding of some empirical phenomenon; and all of them. I remain officially neutral on this matter.

3.2. Distinguishing the Proposal From Other Views

Knowledge by accompaniment avoids the two objections to Mooreanism described in §1. Firstly, it does not involve *inferring* that one is not a BIV from any empirical knowledge. Secondly, knowledge by accompaniment is still available when the content of the relevant empirical knowledge (such as that it is a sunny day) does not make it likely that one is not a BIV. Knowledge by accompaniment is available precisely because being a BIV is a skeptical worry, i.e. because it would be incoherent to believe that it is a sunny day while allowing that one might be a BIV. Further, §4 argues that my proposal can explain away the pull of skepticism—unlike the Moorean proposal.

I have not avoided Mooreanism by adopting the standard 'conservative' alternative (e.g. Wright 2004). On one way of glossing the issue, Mooreans say that one can settle that one is not a BIV *epistemically posterior* to one's perceptual knowledge, while conservatives object that one must settle that one is not a BIV *epistemically prior*. Knowledge by accompaniment seems to provide a third alternative: one can settle that one is not a BIV, and that one is right about a certain empirical matter, *at the same point epistemically speaking*. Let's say that the two judgements are *epistemically parallel*.[42]

Let me briefly explain how I regiment this terminology. Some occurrent mental states and processes count as 'settling that p epistemically prior to settling that q', and some as failing to do so. That makes the mental state or process rationally permissible or not, in a particular respect. For example, the Moorean inference counts as settling that one is not a BIV epistemically posterior to settling that one is right about a given empirical matter; doing so is rationally impermissible.[43] This prohibition is not a matter of the connection between premise and conclusion being too weak: one should believe that if one has hands then one is not a handless BIV, but the corresponding inference is forbidden. It is similarly forbidden to form a JOR while suspending judgement on whether one is a BIV; it is not forbidden to form a JOR while ignoring the question of whether one is a BIV. Judging that one is not a BIV as an accompaniment to judging that one is right about a given empirical matter counts as settling those questions epistemically in parallel, not as settling one matter prior to the other.[44] Settling those matters epistemically in parallel is permissible (assuming one lacks defeating beliefs). The rational way to address a skeptical worry about a belief is to formulate an appropriate TAFC (§2.2); the Moorean inference is not an acceptable alternative.

My proposal may remind the reader of Christopher Hookway's work (2003, 2008). However, Hookway argues that strong FORs indicate that skeptical scenarios "can be ignored" (2008 p. 63). I say that if they are not ignored, they can easily be known not to obtain.

The existence of knowledge by accompaniment is inconsistent with traditional foundationalism, according to which the foundational beliefs do not derive their status from that of other beliefs, while the rest of our beliefs are inferentially supported by the foundations. On my proposal, one knows that one is not a BIV non-inferentially (else the proposal would be Moorean). Yet the belief that one is not a BIV derives its status from its relation to one's empirical beliefs, so it would be misleading to say it is foundationally justified. It is because some empirical beliefs are in all other respects fit to count as knowledge (or rational) that the belief that one is not a BIV gets the relevant positive status.

---

[42] Klein (2004) says two beliefs are "epistemically on a par" iff they are both supported by some third belief. There is no such third belief in cases of knowledge by accompaniment.

[43] Nevertheless, someone who makes the Moorean inference still knows she has hands.

[44] A more detailed treatment will also distinguish epistemic independence. The official ideology is: at t, S settles that p *in way Wp* epistemically prior (/in parallel) to settling that q *in way Wq*. To illustrate, suppose a sensible epistemologist considers two empirical matters, that she has hands and that she has feet, in judging she is not a BIV. She settles she is not a BIV *as an accompaniment to judging that she is right that he has hands* epistemically in parallel to settling that she is right that she has hands. She settles she is not a BIV *as an accompaniment to judging that she is right that she has feet* epistemically in parallel to settling that she is right that she has feet. This does not imply that she settles that she is right she has feet epistemically in parallel to settling she is right she has hands (i.e. 'at the same point'). (If they were epistemically parallel, then it would be impermissible for her to reason from one judgement to the other with the aid of another premise, which is not the case.)

Nor does accepting knowledge by accompaniment have much in common with traditional coherentism. According to coherentism, inferential support flows in both directions between mutually supporting beliefs. My proposal suggests that inferential support flows in neither direction between the belief that one is not mistaken and the relevant JOR. (I already forbade inferring in the Moorean direction.) According to traditional coherentism, membership of the coherent set is what makes every belief in the set justified; there is no relevant asymmetry between the beliefs that cohere. On my proposal, the belief that one is not mistaken constitutes knowledge partly in virtue of the fact that certain other empirical beliefs are in all other respects fit to do so, and not vice versa.

4. A Satisfying Solution to the Puzzle of BIV Skepticism

This section explains why accepting that one has knowledge by accompaniment that one is not a BIV is a satisfying solution to skeptical puzzle. The puzzle is to say how one knows one is not a BIV (§1). A satisfying solution identifies the mistaken intuition, makes it go away, and explains why we had it in the first place. I think the solutions in the literature, including Mooreanism, clearly fail the second and third of these tasks. They still seem wrong when all is said and done; they bite a bullet. My proposal does not. A sensible epistemologist might initially assume that the way she came to believe she is not a BIV does not render it knowledge; but there isn't a residual intuition that there is no knowledge by accompaniment. Reflection removes any intuition that there is no such knowledge. This allows us to decide the question on theoretical grounds. The psychological claims made in §2.2 are speculative but plausible. Knowledge by accompaniment is the only plausible story on the table about how one knows one is not a BIV, it seems to me; and it is part of a plausible general story about how humans know they are not mistaken on a particular occasion. So it is reasonable to conjecture that there is knowledge by accompaniment. All sensible epistemologists have knowledge by accompaniment that they are not BIVs; accepting that's how they know it should end their puzzlement.

Our initial tacit rejection of knowledge by accompaniment is easily explained away. When a judgement Feels Right, it is easy to dismiss the question of how one knows; one insists that one simply knows, and one does not doubt that one knows. For example, one might not be able to say *how* one knows that nothing could be red all over and blue all over at the same time; but one will insist it one knows it, and that it is obvious; one does not start to worry that there is no way one could know such a thing. But suppose one asks oneself how one knows one is not a BIV. The judgement that one is not a BIV does not itself Feel Right (§2). So one does not immediately judge that one simply knows one is not a BIV, and that it's obvious. Without an attached FOR and a consequent judgement that one simply knows one is not a BIV, one brings one's beliefs about one's knowledge- forming capacities to bear on the question—one does some 'information-based' metacognition. There are a number of ways one believes one can know things, including perception, memory, and inference; but those ways do not provide a way to know one is not a BIV. So it will seem that there is no way for one to know one is not a BIV. If the judgement were accompanied by a strong FOR, one would add the process by which one judges that one is not a BIV to one's list of ways one knows things; but in this case there is no FOR. There is no mystery as to why, in the absence of a FOR, people considering the skeptical puzzle don't spot the possibility that the belief-forming method they are actually employing produces knowledge. So there is no mystery as to why it initially seems that one cannot know that one is not a BIV, and hence that one cannot know anything about the external world. Hopefully, the non-obvious argument of this paper will cause you (in a spirit of bold conjecture) to add knowledge by accompaniment to the list of ways you believe you can know things.

A philosopher might worry that my response to skepticism is unsatisfying, because it does not respect some 'internalist' constraint on a solution. I feel in a good dialectical position to reject any alleged constraint that really is incompatible with my proposal.[45] One should not think it is easy to give an 'externalist' response to skepticism. In fact, I think mine is the first plausible one to be presented. It sure doesn't feel like I've cheated, ducking the hard question about skepticism by 'going externalist'.

[45] My solution is compatible with the 'internalist' claim that what it is rational to believe supervenes on one's current mental state. (Compare Feldman & Conee's (2001) 'mentalism' about justification.)

5. Conclusion

I've explained that as a sensible epistemologist, you knew that you are not a BIV before you read this paper. Now you have a speculative but plausible account of *how* you know it. The account provides a satisfactory resolution to the skeptical puzzle.[46]

References

Alter, A. L., & Oppenheimer, D. M. (2009). Uniting the tribes of Fluency to Form a Metacognitive Nation. *Personality and Social Psychology Review,* 13(3), 219–235.

Beebe, J. R. (2009). The Abductivist Reply to Skepticism. *Philosophy and Phenomenological Research,* 79(3), 605–636.

Bogacz, R., Brown, E., Moehlis, .J, Holmes, P., & Cohen, J. D. (2006). The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two- Alternative Forced-Choice Tasks. *Psychological Review,* 113(4), 700–765.

Bogacz, R., Usher, M., Zhang, J., & McClelland, J. L. (2007). Extending a Biologically Inspired Model of Choice: Multi-Alternatives, Nonlinearity and Value-Based Multidimensional Choice. *Philosophical Transactions of the Royal Society B,* 362, 1655– 70.

Broome, J. (1999). Normative Requirements. *Ratio,* 12(4), 398–419.

Carruthers, P. (2011). *The Opacity of Mind*. Oxford, UK: Oxford University Press. Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives,* 13, 57–89.

Davies, M. (2004). Epistemic Entitlement, Warrant Transmission and Easy Knowledge. *Proceedings of the Aristotelian Society Supplemental,* 78, 213–245.

DeRose, K. (2004). Single Scoreboard Semantics. *Philosophical Studies,* 119(1–2), 1–21. Dretske, F. I. (2005). The Case Against Closure. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology*. Oxford, UK: Blackwell.

de Sousa, R. B. (2008). Epistemic Feelings. In G. Brun, U. Doguoglu, & D. Kuenzle (Eds.), *Epistemology and Emotions*. Aldershot, UK: Ashgate.

Deutsch, R., Kordts-Freudinger, R., Gawronski, B., & Strack, F. (2009). Fast and Fragile: A New Look at the Automaticity of Negation Processing. *Experimental Psychology,* 56(6), 434–446.

Dougherty, T., & Rysiew, P. (2009). Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions. *Philosophy and Phenomenological Research,* 78(1), 123–132.

Evans, J. St. B. T. (2008). Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition. *American Review of Psychology,* 59, 255–78.

Evans, J. St. B. T., & Stanovich, K. E. (2013a). Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science,* 8(3), 223–241.

Evans, J. St. B. T., & Stanovich, K. E. (2013b). Theory and Metatheory in the Study of Dual Processing: Reply to Comments. *Perspectives on Psychological Science,* 8(3), 263– 271.

Feldman, R., & Conee, E. (2001). Internalism Defended. *American Philosophical Quarterly,* 38(1), 1–18.

Fine, K. (2001). The Question of Realism. *Philosophers' Imprint*, 1:1.

Frederick, S. (2005). Cognitive Reflection and Decision Making. *The Journal of Economic Perspectives,* 19(4), 25–42.

Gopnik, A. (1998). Explanation as Orgasm. *Minds and Machines* 8(1), 101–118. Haidt, J. (2001). The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review,* 108(4), 814–834.

Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford, UK: Oxford University Press. Hookway, C. (2003). Affective States and Epistemic Immediacy. *Metaphilosophy*, 34(1–2), 78–96.

Hookway, C. (2008). Epistemic Immediacy, Doubt and Anxiety: On a Role for Affective States in Epistemic Evaluation. In G. Brun, U. Doguoglu, & D. Kuenzle (Eds.), *Epistemology and Emotions*. Aldershot, UK: Ashgate.

[46] I would like to thank the following for helpful comments and conversations: Andrew Cortens, Stephen Crowley, John Hawthorne, Brian Kierland, Peter Klein, Asher Koriat, Tania Lombrozo, Matthew McGrath, Ram Neta, Ernie Sosa, Jim Stockton, Jennifer Nagel, Nina Strohminger, Jonathan Weinberg, Jonathan Weisberg, an anonymous reviewer for another journal, and the organizers and participants at Southern Methodist University's 2014 epistemology workshop.

Huemer, M. (2011). The Puzzle of Metacoherence. *Philosophy and Phenomenological Research,* 82(1), 1–21.
Jackson, A. (2011). Appearances, Rationality, and Justified Belief. *Philosophy and Phenomenological Research,* 82(3), 564–593.
Kahneman, D. (2003). A Perspective on Judgment and Choice: Mapping Bounded Rationality. *American Psychologist,* 58, 697–720.
Koriat, A. (2007). Metacognition and Consciousness. In P. D. Zelazo, M. Moscovitch & E. Thompson (Eds.), *The Cambridge Handbook of Consciousness*. Cambridge, UK: Cambridge University Press.
Koriat, A. (2011). Subjective Confidence in Perceptual Judgments: A Test of the Self- Consistency Model. *Journal of Experimental Psychology: General,* 140(1), 117–139.
Koriat, A. (2012). The Self-Consistency Model of Subjective Confidence. *Psychological Review,* 119(1), 80–113.
Koriat, A. (2013). Confidence in Personal Preferences. *Journal of Behavioral Decision Making,* 26(3), 247–259.
Koriat, A., & Adiv, S. (2011). The Construction of Attitudinal Judgments: Evidence From Attitude Certainty and Response Latency. *Social Cognition,* 29(5), 577–611.
Kornblith, H. (2012). *On Reflection*. Oxford, UK: Oxford University Press.
Lewis, D. (1996). Elusive Knowledge. *The Australasian Journal of Philosophy,* 74(4), 549–567. Lombrozo, T. (2012). Explanation and Abductive Inference. In K.J. Holyoak & R.G.
Morrison (Eds.), *Oxford Handbook of Thinking and Reasoning*. Oxford, UK: Oxford University Press.
Lyons, J. C. (2009). *Perception and Basic Beliefs*. Oxford, UK: Oxford University Press. McGrath, M. (2013). Dogmatism, Underminers and Skepticism. *Philosophy and Phenomenological Research,* 86(3), 533–562.
Michaelian, K. (2012). Metacognition and Endorsement. *Mind and Language,* 27(3), 284– 307.
Nagel, J. (2012). Intuitions and Experiments: A Defense of the Case Method in Epistemology. *Philosophy and Phenomenological Research,* 85(3), 495–527.
Nisbett, R. E., & Wilson, T. D. (1977). Telling More That We Can Know: Verbal Reports on Mental Processes. *Psychological Review,* 84(3), 231–259.
Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press. A relevant selection is reprinted in Sosa et al. (2008).
Peirce, C. S. (1992). *The Essential Peirce Volume 1*. N. Hauser & C. Kloesel (Eds.). Bloomington, IN: Indiana University Press.
Proust, J. (2010). Metacognition. *Philosophy Compass,* 5(11), 989–998.
Proust, J. (2013). *The Philosophy of Metacognition: Mental Agency and Self-Awareness*. Oxford, UK: Oxford University Press.
Pryor, J. (2000). The Skeptic and the Dogmatist. *Nous* 34(3), 517–549.
Pryor, J. (2004). What's Wrong with Moore's Argument? *Philosophical Issues,* 14, 349–378. Pryor, J. (2012). When Warrant Transmits. In A. Coliva (Ed.), *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*. Oxford, UK: Oxford University Press.
Sosa, E. (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge Volume II*. Oxford, UK: Oxford University Press.
Sosa, E., Kim, J., Fantl, J., & McGrath, M. (Eds). (2008). *Epistemology: An Anthology, Second Edition.* Oxford, UK: Blackwell.
Stanovich, K. E. (2011). *Rationality and the Reflective Mind*. Oxford, UK: Oxford University Press.
Thompson, V. A. (2009). Dual-Process Theories: A Metacognitive Perspective. In J. St. B.
T. Evans & K. Frankish (Eds.), *In Two Minds: Dual Processes and Beyond*. Oxford, UK: Oxford University Press.
Thompson, V. A., Prowse Turner, J. A., & Pennycook, G. (2011). Intuition, Reason, and Metacognition. *Cognitive Psychology,* 63(3), 107–140.
Thompson, V. A., Prowse Turner, J. A., Pennycook, G., Ball, L. J., Brack, H., Ophir, Y., & Ackerman, R. (2013a.) The Role of Answer Fluency and Perceptual Fluency as Metacognitive Cues for Initiating Analytic Thinking. *Cognition,* 128(2), 237–251.
Thompson, V. A., Prowse Turner, J. A., Pennycook, G., Ball, L. J., Brack, H., Ophir, Y., & Ackerman, R. (2013b). The Role of Answer Fluency and Perceptual Fluency in the Monitoring and Control of Reasoning: Reply to Alter, Oppenheimer, and Epley (2013). *Cognition,* 128(2), 256–8.
Trout, J. D. (2007). The Psychology of Scientific Explanation. *Philosophy Compass,* 2(3), 564–591.
Tucker, C. (2010). When Transmission Fails. *Philosophical Review,* 119(4), 497–529.

Tucker, C. (2013). Seemings and Justification: An Introduction. In C. Tucker (Ed.), *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford, UK: Oxford University Press.

Unkelbach, C. (2007). Reversing the Truth Effect: Learning the Interpretation of Processing Fluency in Judgments of Truth. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(1), 219–30.

Vogel, J. M. (1990). Are There Counterexamples to the Closure Principle? In M. D. Roth & G. Ross (Eds.), *Doubting*. Dordrecht, Netherlands: Kluwer. Reprinted in Sosa et al. (2008).

Warfield, T. A. (2005). Knowledge From Falsehood. *Philosophical Persepectives*, 19, 405–16. Wedell, D. H. (1997). Another Look at Reasons for Choosing and Rejecting. *Memory and Cognition*, 25(6), 873–887.

Whittlesea, B. W., & Williams, L. D. (1998). Why Do Strangers Feel Familiar, But Friends Don't? A Discrepancy-Attribution Account of Feelings of Familiarity. *Acta Psychologica*, 98(2–3), 141–165.

Wright, C. (2004). Warrant For Nothing (And Foundations For Free)? *Proceedings of the Aristotelian Society Supplemental*, 78, 167–212.

Yalcin, S. (2011). Nonfactualism About Epistemic Modality. In A. Egan & B. Weatherson (Eds.), *Epistemic Modality*. Oxford, UK: Oxford University Press.