

A problem for rationalist responses to skepticism

Sinan Dogramaci

Philosophical Studies

An International Journal for Philosophy
in the Analytic Tradition


ISSN 0031-8116

Volume 168

Number 2

Philos Stud (2014) 168:355-369

DOI 10.1007/s11098-013-0136-4



Your article is protected by copyright and all rights are held exclusively by Springer Science+Business Media Dordrecht (outside the USA) . This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

A problem for rationalist responses to skepticism

Sinan Dogramaci

Published online: 23 April 2013

© Springer Science+Business Media Dordrecht (outside the USA) 2013

Abstract Rationalism, my target, says that in order to have perceptual knowledge, such as that your hand is making a fist, you must “antecedently” (or “independently”) know that skeptical scenarios don’t obtain, such as the skeptical scenario that you are in the Matrix. I motivate the specific form of Rationalism shared by, among others, White (Philos Stud 131:525–557, 2006) and Wright (Proc Aristot Soc Suppl Vol 78:167–212, 2004), which credits us with warrant to believe (or “accept”, in Wright’s terms) that our senses are reliably veridical, where that warrant is one we enjoy by default, that is, without relying on any evidence or engaging in any positive argument. The problem with this form of Rationalism is that, even if you have default knowledge that your senses are reliable, this is not adequate to rule out every kind of skeptical scenario. The problem is created by one-off skeptical scenarios, scenarios that involve a highly reliable perceiver who, by a pure fluke, has a one-off, non-veridical experience. I claim you cannot infer that your present perceptual experience is veridical just on the basis of knowledge of your general reliability. More generally, if you infer that the present F is G , just on the basis of your knowledge that most F s are G s, this is what I call statistical inference, and, as I argue, statistical inference by itself does not generate knowledge. I defend this view of statistical inference against objections, including the objection that radical skepticism about our ordinary inductive knowledge will follow unless statistical inference generates knowledge.

Keywords Rationalism · Mooreanism · Dogmatism · Statistical inference · Knowledge · Lottery

S. Dogramaci (✉)
The University of Texas at Austin, Austin, TX, USA
e-mail: sinan.dogramaci@gmail.com

1 Set up

I'll say you have "considered s " to mean you have an epistemic attitude toward s , either belief, disbelief, or uncertainty. The following closure principle is then highly plausible, and I'll assume it:

Closure: you know that p , only if, for any proposition s that you have considered and know is inconsistent with p , you know that not- s .

Closure doesn't say much about *how* we know this or that; in particular it doesn't comment on the *epistemic dependence* of this or that knowledge on other epistemic states. Here are two competing views of epistemic dependence concerning *perceptual knowledge* that p (with labels roughly following Silins 2007). Note these are alternative and exhaustive possibilities; each negates the other.

Rationalism: necessarily, you perceptually know that p , only if, for any considered proposition s that you know entails [not- p and it appears to you that p], you know that not- s *antecedently* to your perceptually knowing that p .

Mooreanism: it's possible for you to perceptually know that p , even if, for some considered proposition s that you know entails [not- p and it appears to you that p], you *don't* know that not- s *antecedently* to your knowing that p (for example, because you *inferred* not- s from p , like Moore (1939) did).

Antecedently to means *independently of*, *prior to*, or *without presupposing* (see Pryor 2000, p. 525). Probably no definition can be given without recourse to yet other epistemic notions. I mean the same notion that Pryor and other authors on this topic mean. The usual meaning is pre-theoretically clear enough, and my arguments won't assume anything controversial about it. Mooreanism is sometimes also called *dogmatism*, after Pryor (2000).

The restriction to *considered* inconsistent propositions is included here because my interest is in evaluating competing views of how we can *reject* skeptical scenarios—propositions that entail [not- p but it appears to you that p]—when those scenarios are *considered* and *known* to be inconsistent with some item of perceptual knowledge. Views on which skeptical scenarios are diagnosed as propositions we may "properly ignore" (see Lewis 1996, p. 554), for example because they are not salient in a speaker's or a knower's context, are thus not my topic. My interest is in how we, we philosophers, can rule out all the skeptical scenarios that we explicitly consider, including one-off scenarios like *error*. How can we attribute first-personal, present-tense knowledge that these scenarios do not obtain?

Contemporary theorists naturally placed in the Moorean camp include Pollock (1974), Alston (1986), Pryor (2000, 2004), Klein (2000), McDowell (2008), Williamson (2000), Huemer (2001) and Peacocke (2004). Problems for Mooreanism have been raised by Cohen (2002), White (2006), Silins (2007) and Wright (2007): notably, Mooreanism would seem to allow us to acquire certain substantive knowledge implausibly easily, and it can be argued to conflict with modest Bayesian principles. The positive views offered by Cohen (2002), Wright (2002, 2004), White (2006), and Silins (2007) fall in the Rationalist camp.

My aim in this paper is to turn the spotlight on Rationalism. I'll argue that once we see what Rationalists must say in order to reject skepticism, the position turns out to be problematic.

Before continuing, I should flag one point regarding my placement of Wright, White, and Silins into the Rationalist camp, which is that they hold much of their discussions in terms of *justification* rather than *knowledge*. (Replace “know” with “have justification to believe” in the above formulation of Rationalism to get something like the formulations they explicitly discuss; some specific ways of formulating Rationalism in terms of justification will be given and discussed in detail below). My strategy will be to first raise a problem for Rationalism as formulated in terms of knowledge. Then, afterward, I'll argue that if you hold a Rationalist view about justification, then, assuming you are not a skeptic, you must also accept—and inherit the problem for—the Rationalist view of knowledge.

2 A problem created by one-off skeptical scenarios

It will make some formulations more comprehensible if I may use “rule out *s*” to abbreviate “know not-*s*”. So, Closure says that to know *p* you must rule out every considered scenario that you know entails not-*p*. For example, to know that my hand is making a fist, I must rule out *dreaming*, *demon*, *brain-in-a-vat (BIV)*, and *Matrix*. I'll call these the *familiar* skeptical scenarios, and trust their italicized names are understood. The familiar skeptical scenarios all involve massive deception; even *dreaming* is treated as a scenario in which most of your current perceptual beliefs are false.

Rationalism doesn't look its worst when it comes to our ruling out the familiar skeptical scenarios. Since these scenarios are all inconsistent with my reliability, the rationalist can say that we know these scenarios don't obtain because we enjoy “a kind of default justification for assuming the general reliability of our perceptual faculties” (White 2006, p. 552). It's not the most intuitive thing anyone's ever said, and defenders of the view realize that, but anti-skeptical epistemology is a game of finding the least bad cost-benefit balance. Default knowledge of reliability isn't a prohibitive cost if it buys us anti-skepticism, but I will argue there are costly hidden fees that arise for the Rationalist even after paying for the claim that we have default knowledge of our reliability.

To begin, notice that massive deception is inessential to the case for skepticism. To know my hand is making a fist, I must rule out all considered possibilities on which that's false and I'm the unwitting victim of a *one-off* hallucination, hypnosis, hologram, neural short-circuit, or a clever magic trick. Or perhaps there is an evil demon who intervenes in my life only to fool me at this one moment. There are innumerable such one-off possibilities that I know are inconsistent with my making a fist, but consistent with nearly all my other beliefs. Massive deception needn't even be mentioned to make the skeptical argument that you don't know *anything* perceptually; the skeptic need only argue that for *each* purported item of knowledge, there's some such one-off skeptical scenario for the proposition in question which you can't rule out. (Skeptic: “You don't know you're eating an apple, because the

demon could be deceiving you about *eating an apple*. You don't know the window is open *right now*, because the demon could have picked *this* moment to intervene. You don't know..." So, if we're to reject skepticism, Closure requires us to rule out—know the negations of—each and every considered one-off skeptical scenario, and anti-skeptical Rationalists say this knowledge is antecedent to our perceptual knowledge.

Suppose, now, that I know that it appears to me that p , and I consider the following totally generic, one-off skeptical scenario: it appears to me that p , but in fact p is not true. I'm simply having a non-veridical illusion, for whatever unspecified reason. Call this generic skeptical scenario *error*.¹ By Closure, my knowing p requires me to rule out *error*, to know not-*error*. The negation of error can be logically represented as a material conditional: [it appears to me that p] \supset [p is true]. So, anti-skeptical Rationalism says that, *antecedently* to my *perceptually* knowing that p , I know that [it appears to me that p] \supset [p is true].

Now we press this uncomfortable question: *how* can I know not-*error* if, as the Rationalist would have it, this knowledge must be antecedent to my perceptual knowledge that p ? What good options are left when my perceptual knowledge is off-limits? We can partition the possibilities into three, (i), (ii) and (iii). Option (iii) is the only one that's even initially believable and I'm going to say why it's not ultimately believable, but it's worthwhile to quickly state and reject (i) and (ii), if only to motivate (iii).

Option (i) is that I non-inferentially know not-*error*. In other words, for any given item of perceptual knowledge that p , I can *immediately* know the associated material conditional: [it appears to me that p] \supset [p is true]. I take the bizarreness of this view to be apparent. While several philosophers have suggested that we have, by default, immediate knowledge of the negations of a number of the *familiar* skeptical scenarios (*demon*, *BIV*, *Matrix*, etc.), I know of no one who says—and it certainly would be far more of a stretch to say—that we have, *for each* perceptual experience, default immediate knowledge that the experience is, or *will be*, veridical. (Cohen 2002 calls the view absurd; see p. 322.)

Option (ii) is that I know not-*error* via a deductive inference. But if a Moorean deduction from my perceptual knowledge is off limits, this option can only look worse than the previous one. What could I know, *non*-perceptually, that entails that each of my perceptual experiences is or will be veridical? I would require knowledge that entails my perfect perceptual reliability, but that is absurd. Indeed, I know the negation of that: I know that I am *not* perfectly reliable. Even if Rationalism could somehow credit me with a justified belief in my perfect reliability, it would be a

¹ Cross (2010) plausibly observes that skeptical scenarios do not make for *compelling* skeptical arguments unless they include some *explanation* of why the subject is suffering from an illusion. This doesn't affect the argument I'm about to give. My argument just concerns whether Rationalism or Mooreanism is right about structural relations of epistemic dependence between perceptual knowledge and skeptical scenarios. *Error* captures all the features of skeptical scenarios involved in this debate. Consider, for example, the Rationalists' Bayesian argument against Mooreanism. The argument only concerns skeptical scenarios sharing the features of *error*; the explanatory properties Cross is concerned with are irrelevant.

justified belief in a falsehood, and thus if I deduced not-*error* just from this false premise, I would, like the famous cases in Gettier (1963), fail to know not-*error*.²

The last option, option (iii), is that I know not-*error* via a non-deductive inference, that is, an inductive or statistical inference. This would seem to be the option that contemporary Rationalists want to go for. (McDowell (2011, p. 13) still says it is bizarre on its face.) White (2006) proposed we have default justification to believe we're perceptually reliable. Wright (2004), inspired by Reichenbach (1939), explored strategies for crediting us with default knowledge of (or at least "entitlement" to "accept") various so-called "cornerstone" propositions, roughly speaking propositions that the world meets preconditions allowing for our reliable interaction with it. And, Peacocke (2004), though he officially favors Mooreanism (p. 65), provides an argument amenable to Rationalists: his a priori argument that the best explanation of the contents of our perceptual experiences makes them "by and large correct" (p. 98); Russell (1912) and Vogel (1990) make very similar arguments.

However, even though it appears to be the best option for the Rationalist, option (iii) does not ultimately allow for an adequate response to skepticism. Here, now, is the real problem for the Rationalist. The problem is that knowledge of my reliability, although it would allow me to rule out all the *familiar* skeptical scenarios, it is not sufficient to rule out *error*. Knowledge I'm reliable, that my perceptual experiences are by and large correct, could allow me to rule out *dreaming*, *Matrix* or *BIV*, since those scenarios all entail my unreliability, and so I can rule them out by a deduction. But, *error* does not entail my unreliability. *Error* is consistent with the fact that *most* of my perceptual experiences, even the vast majority, are veridical. If I infer the negation of *error*, just on the basis of the premise that most of my experiences are veridical, then this will be a case of *statistical inference*. You draw a *statistical inference* when you infer *Ga* just on the basis of [*Fa* and most *Fs* are *Gs*]. But, I claim, whatever justification the inferred conclusion in such an inference may enjoy, it is not *knowledge*. If statistical inference yielded knowledge, then you could know your lottery ticket will lose, but you don't know this. On the basis of knowledge that most lottery tickets lose, you may be justified in inferring your ticket will lose, but you do not know this conclusion. Likewise, I cannot rule out *error*, I cannot know not-*error*, merely by reasoning as follows: things are mostly as they appear to me, it appears to me that *p*, therefore *p*. For, *error* is just the possibility that this is an occasion when things are not as they appear. Just as I can't use a statistical inference to *know* that my ticket in the million-ticket lottery will lose, I can't use such an inference to know that its appearing to me that *p* isn't that one-in-a-million misfire. *Error* is the possibility that

² It may occur to some to wonder: isn't there one well-known Rationalist, namely Descartes, who would go for option (ii)? I believe that Descartes indicates otherwise in the last two paragraphs of the *Meditations* (Descartes 1641/1984, pp. 60–61), though I admit I am not certain. He does famously think that "clear and distinct" perception is infallible, but it seems he is sensible enough to not think ordinary sensory perception is infallible. Although Descartes thinks he can deduce that he is not dreaming or otherwise massively deceived, it seems he does not think he can deduce the veridicality of individual ordinary perceptual experiences. In any case, I assume no one today believes Descartes's arguments for rejecting skeptical scenarios.

my perhaps *super*-reliable perceptual faculties have, by what may be a one-off fluke, misled me. Thus, option (iii) is not acceptable.

3 Defense

I will reply to two objections against my presentation of a problem for (iii). The first objection is that the lottery has idiosyncrasies, and is thus not a good case for drawing any general epistemological morals. The second objection is that my argument proves too much by entailing radical, skeptical doubts about inductive knowledge.

(1) According to the first objection, I have drawn a bad analogy between Rationalist perceptual knowledge and lottery knowledge. Philosophers have suggested that we may fail to have knowledge in the lottery case for various idiosyncratic reasons. For example, Harman (1986, p. 75) wondered whether our desire to win the money might be responsible for our reluctance to say we know we'll lose.³ Perhaps then, the present objector suggests, even though it is true that I can't know my lottery ticket will lose just on the basis of my knowledge that most tickets lose, nonetheless the Rationalist may know *not-error* just on the basis of her knowledge that most perceptual appearances are veridical.

In reply to any concerns about the lottery's idiosyncrasies, I would turn to any number of simple, clear cut examples of statistical inference. If we roll a pair of dice we know are fair, and they fall under the couch, then we can infer by statistical inference that they did not come up snake eyes, but, without lifting the couch, we intuitively we don't know this. If we spin a roulette wheel, we can infer by statistical inference that it did not come up "00", but intuitively, until we see where it comes to rest, we don't know this. If we shuffle a deck of cards and deal a poker hand face down, we can infer by statistical inference that it is not a straight flush, but intuitively we don't know this. If we pull the arm on a slot machine, we can infer by statistical inference that it won't come up 7-7-7, but intuitively we don't know this. Each of these inferred conclusions may be highly justified; perhaps they are even appropriately assertible; all I'm observing is that we intuitively don't possess knowledge of these conclusions.

I do not know of any case that is as clearly an instance of statistical inference as the above ones and that generates the intuition that we possess knowledge, much less generates the intuition as strongly as the above cases generate the intuition that we don't. However, I should not avoid mentioning one sort of case that I find tricky: many people have the intuition that we can know that if you play the lottery every single day of the coming year then you won't win every day, yet it seems our only way of arriving at this belief is by making a statistical inference. A similar case is Vogel's (1999, p. 165) "Heartbreaker": intuitively, I can know that it's not the case that every golfer in an invitational celebrity tournament will score a hole-in-one on the toughest course, the Heartbreaker; but, how can I infer this except by some

³ Though, earlier in the same chapter (p. 72), Harman also says about lotteries, "I have no idea how to account for our reluctance to attribute knowledge in cases of this sort."

application of statistical inference? But, even in the face of tricky cases like these last ones, I am inclined to hold on to the general view that statistical inference does not yield knowledge. I am inclined to allow that we do apply statistical inference in these cases, but to explain away any temptation to say we acquire knowledge in these last cases. Hawthorne (2004, pp. 16, 19–20) shows how to explain it away; he observes that our temptation to attribute knowledge in these cases seems to evaporate once we realize that the event in question could be repeated, and if repeated long enough, there eventually would be a “winner”. Once we remember that someone would win the lottery daily for a year if enough subjects played the daily lottery over the course of enough years, we lose the temptation to say, for any given subject in a given year, that we *know* he or she will not be the lucky one. And, once we remember that the golfers in some tournament will all score a hole-in-one on the Heartbreaker if enough tournaments were held, we again lose the temptation to say, for any given individual tournament, that we *know* it will not be the very rare occasion where that happens. Perhaps some temptation to attribute knowledge lingers, but it seems a weak basis for arguing that statistical inference in fact can, at least in these odd circumstances, yield knowledge.

Furthermore, whatever lingering temptation some may still have to attribute knowledge in the cases of the daily lottery player and the Heartbreaker even after taking Hawthorne’s considerations into account, these cases do not to seem to provide the most accurate model for the particular statistical inference the Rationalist makes when she infers that her present perceptual experience is veridical. Rather, the Rationalist’s inference seems similar to that of the subject who infers, say, that her unseen poker hand is not a straight flush, just on the basis of the premise that poker hands dealt from shuffled decks are reliably not straight flushes.

(2) The second objection says that, in denying that statistical inference yields knowledge, my view entails an unacceptable, radical skepticism about ordinary inductive knowledge. This objection says that, if we do not have knowledge by statistical inference, then I do not know that the gas tank in my car is running on empty (upon seeing the gauge read ‘E’), that there are penguins in Antarctica, that the sun will rise tomorrow, or that water will come from the faucet when I turn the tap tomorrow morning.

To begin my explanation of why this objection is incorrect, it’s important we distinguish different kinds of inductive inference. I said there are no clear examples of a statistical inference that clearly yield knowledge. However, *inferences to the best explanation* yield knowledge, and they are ubiquitous. So, I’m not (absurdly) denying that we have knowledge by induction; denying that would amount to radical Humean skepticism. But, denying that we have knowledge generated by statistical inference is perfectly fine common sense. We have here two distinct types of non-deductive inference, and, importantly, neither one can be formally reduced to the other. To see that they are not reducible to each other, just notice: the conclusion of a statistical inference is not generally an explanatory hypothesis; and, the conclusion of an inference to the best explanation is not generally an instantiation on a corresponding premise of the form *Most Fs are Gs*.

Inferences to the best explanation provide us with a wealth of ordinary inductive knowledge. The gas tank example is the simplest sort of case. I know the gas tank is

empty because that best explains why the gauge points to 'E'. The penguin example illustrates inductive knowledge by testimony. I know there are penguins in Antarctica because I have heard and read testimony about this from other people, and the best explanation of this testimony is that penguins were actually observed living in the Antarctic by someone earlier in the chain of testifiers. Many epistemologists defend this view that rational trust in testimony reduces to an inference to the best explanation, including Harman (1965), Fricker (1994), Adler (1994; 2002, p. 157), Schiffer (2003, p. 303), Malmgren (2006), and Lipton (2007). Some epistemologists, such as Coady (1992) and Burge (1993) have a different view: they say testimony is trustworthy, or known to be reliable, on some other a priori grounds. Whichever of these views—the two mainstream views—of testimony is right, my position is safe from leading to any skeptical results; knowledge by testimony does not rely on statistical inference.

How does inference to the best explanation yield ordinary knowledge of future events, such as that the sun will rise tomorrow, or that water will come out of my faucet when I turn the tap tomorrow morning? The basic form of the reasoning is simple: if I know that all observed *F*s were *G*s, that will often⁴ serve as an adequate basis for me to come to know, by an inference to the best explanation, the generalization that all *F*s, observed or not, are *G*s. The generalization explains the past observed data. Then, from that generalization, I can know, by deduction, that the next *F* will be a *G*. (So, we also have knowledge by “enumerative inductive inference” as a by-product of knowledge by inference to the best explanation, as Harman (1965) famously noted.)

However, although the basic form of that reasoning is simple, complications arise because it is not always obvious how that basic form can be applied to a realistic case. After all, we all know it is not really true that the sun will *always* rise (since the sun will die), and we know it is not true that water *always* comes from the faucet when the tap is turned (since faucets, on rare occasions, abruptly break). So, we cannot really know, by inference to the best explanation, quite such broad generalizations as these ones to deduce that the sun will rise tomorrow or that my faucet will give water tomorrow.

The basic form can still be applied, however. We just need to select *F* and *G* in ways that suitably restrict the generalizations. There may be many ways to do this, but one simple way is to restrict the cases involved in the inference to all cases that are, to a certain degree, *similar* or *nearby* to the actually observed cases. The idea is to think of *F* as not merely occasions when the tap is turned, but rather occasions when the tap is turned and nearby to the actually observed occasions, nearby temporally and in the modal sense familiar since Stalnaker (1968) and Lewis (1973, 1979). Thinking of *F* (and *G*) this way, my premise, the data that calls for an explanation, doesn't change; it's still just my observations that water came from the faucet every time I recently turned the tap, since these cases are, degenerately,

⁴ I insert 'often' in order to set aside a number of well-known philosophical issues that aren't relevant to our debate here; those issues include (a) how to exclude Goodman's gruesome inferences, (b) how to state the conditions for the absence of defeaters for this sort of inference, (c) how many *F*s must have been observed in order for the data to “call out” for any explanation at all, and (d) how many *F*s must have been observed for the inferred explanation to be strongly enough supported to qualify as known.

nearby to themselves. But, I don't then infer that water *always* comes from the tap; I know, as a matter of independent background knowledge, that's almost certainly not true (any man-made faucet will surely break eventually). Rather, I infer the following generalization, which I have no reason to doubt, and which is the best explanation of the data, among its competitors: water will or would come from the tap every time I recently did turn the tap, were to have turned the tap an extra time, or were to turn the tap in the near future. This explanation is just a generalization over certain nearby modal and temporal cases: water comes from the faucet in all of a class of modally and temporally nearby cases where the tap is turned. Any competitor, logically, must say water wouldn't come from the tap, were I to give, or have given, the tap an extra turn; that's not a better explanation at all. This shows how we can fit the inference into the basic form given earlier: all observed *F*s are *G*s, therefore all *F*s are *G*s.⁵

To infer that water will come from the faucet tomorrow morning, there's no need for the data and the inferred explanation to be broader generalizations. In particular, the data need not be that water *always* came from my faucet when I turned the tap. Although water failed to come from the faucet when there was an abrupt break once last year, the fact that water has consistently been coming from the faucet on recent occasions is by itself data that calls for an explanation. (White (2005, Sect. 1.2.1), provides a clear and useful discussion of when data calls for explanation.) And, while it is always possible to give multiple non-competing explanations, often increasingly richer and more thoroughly satisfying explanations, we may infer a particular hypothesis as long as it is a sufficiently good explanation and best among its competitors. I claim the generalization, "Water comes out when the tap is turned in all cases similar to the recently observed cases", is a good explanation of our data, and best among its competitors. Harman (1965) pointed out that a generalization can serve as an explanation of its sub-cases in this way, and White (2005) agrees, but White also suggests (Sect. 2.2) that we may give a yet better explanation if we can also explain the generalization itself. This extra meta-explanation, if desired, can often be supplied just by adding that there exist laws or mechanisms that are responsible for the generalization in question. For example, we can say: the faucet worked every time I recently turned the tap, and that's because it works in all such similar cases (i.e. it would work, were I to have turned the tap recently or were I to turn it soon), and that in turn is because some mechanism is in place and successfully functions in all those cases to make the faucet give water when the tap is turned. Even without saying anything more about the posited mechanism, this can allow us to enhance our explanation.

⁵ Indeed, given the Stalnaker–Lewis analysis of counterfactual and future conditionals, the inference could be viewed as a simple application of Modus Ponens. On that analysis, I'm in a position to know this conditional, "If I (were to) turn the tap tomorrow morning, water will (would) come from the faucet". For, Stalnaker and Lewis say that conditional comes out true just in case water comes from the faucet in all those worlds where I turn the tap tomorrow morning and that are (otherwise) most similar to ours, and, as argued, this is part of a generalization that best explains my observed data. I don't want to presuppose that the Stalnaker–Lewis analysis of conditionals is correct (in the main text, I only borrowed their talk of similarity or nearness), but it does offer an especially elegant picture of this kind of inductive knowledge.

An inference to the best explanation, on this account, *cannot* provide you with knowledge that your poker hand is not a straight flush, that the slot-machine won't turn up 7-7-7, or that the roulette wheel won't land on "00". Even if I see the roulette wheel land not on "00" on every one of a long string of recent spins, I know, as a matter of background knowledge, that a possible spin landing on "00" is just as similar, just as near in modal space, to these observed cases as a spin landing on any other number. So, I cannot infer that all spins most similar to the recently observed ones will likewise not land "00", and so I cannot infer the conditional, "If the wheel is (were) spun one more time, it will (would) not land "00" again". (Lipton (1992/2004) makes this point that background knowledge defeats the purported inference to the conclusion that the roulette wheel will continue to avoid "00"; and Lewis (1996) points out that cases where your lottery ticket is a winner are as similar to actuality as cases where your ticket is loser.)

So, in response to the objection that I risk radical skepticism by denying that statistical inference yields knowledge, I have responded by elaborating how inference to the best explanation, a distinct and irreducible form of inductive reasoning, provides us with our ordinary inductive knowledge. A critic of Rationalism, and a sympathizer with Mooreanism, should account for ordinary inductive knowledge by appealing to inference to the best explanation.

The Rationalist cannot, by an inference to the best explanation, infer that her present visual experience is veridical; from the Rationalist's premise that most of her perceptual experiences are veridical, that inference is merely a statistical inference, and so not knowledge yielding. However, it may be interesting to notice, at this point, that an inference to the best explanation could explain perceptual knowledge within a Moorean view. The view would say that you can infer p immediately from the premise that it appears to you that p , since the former best explains the latter. This is a view not much discussed, but it's a perfectly fine view. It is not the view of any of the Mooreans cited earlier; they all claim we have non-inferential knowledge of p , the perceptually known proposition. We should also be careful to distinguish it from the above-mentioned Peacocke-inspired Rationalist strategy (Peacocke 2004); that view bases perceptual beliefs on an inference from the premise that our perceptual appearances are *most of the time* accurate, "by and large correct" as Peacocke says. That would of course be a statistical inference. But, knowledge yielding inferences to the best explanation (and enumerative inductive inference) should not, and need not, involve any statistical inferences.

A referee raised the question whether the Rationalist could argue that our perceptual knowledge is analogous to our knowledge as I've characterized it in the faucet case. A first stab at such a strategy would propose that subjects can make use of the following inference: from the (granted) premise that most of my perceptual experiences are veridical, infer the (alleged) explanation that all of my perceptual experiences are veridical. This use of inference to the best explanation is distinct from, but somewhat similar to, the applications I've endorsed, such as in the faucet example. My objection to this style of inference is that it is intuitively too strong, just like the (often) unreasonably strong inference that all F s are G s from the premise that all observed F s are G s. Furthermore, it is false that all my perceptual experiences are veridical, and so, as I noted earlier, any

deduction from this falsehood would, like a Gettier case, fail to yield knowledge. A second stab the Rationalist might take is to instead propose we reason as follows: from the premise that most of my *nearby* perceptual experiences are veridical, infer the (alleged) explanation that all of my *nearby* perceptual experiences are veridical. The problem here is: how can this premise be known? The Rationalist's standard a priori premise, granted here, is that most perceptual experiences are veridical. But, inferring from this that most nearby experiences are veridical requires a statistical inference, which, I of course claim, will not yield knowledge. And the Rationalist cannot say a subject perceptually knows the premise that most of her nearby perceptual experiences are veridical: the Rationalist is giving a response to skepticism, and it would be question-begging to appeal to such perceptual knowledge as a known premise in the explanation of how the subject has perceptual knowledge. By contrast, attributing perceptually known premises is fair game in my discussion of the faucet case, where I was giving an account of inductive knowledge, knowledge that may happily be based on perceptually known premises.

Finally, given the above clarifications about the difference between statistical inference and inference to the best explanation, I can now add one last point concerning my claim that the Rationalist assimilates perceptual knowledge to lottery knowledge. I want to emphasize that the lottery, and statistical inference in general, is the apt analogy or model for Rationalist perceptual knowledge, and not the lottery's well-known cousin case, the preface paradox. I certainly would *not* object to the view that an author *knows* the true claims in her book, even when she also knows some claims in her book are false (hence, I don't object to the rejection of *multi-premise* closure; see Hawthorne 2004). But the author's evidence for the true claims in her book is not her knowledge that most claims in the book are true; rather, each claim has its own special evidence, and the reader can infer that there exists such evidence as its part of the best explanation of the author's testimony. That's not like in the lottery at all. And it's not like the Rationalist view on which we draw upon a broad, default justification that our perceptual experiences are reliable. On the Rationalist view, individual experiences don't *each* have any special evidence that's independent of whatever they *all* share.⁶ A Moorean view, of course, is motivated by just the opposite thought: each experience contributes a justification to the corresponding perceptual judgment antecedently to our rejection of skeptical scenarios, in particular, antecedently to any knowledge that the experience is or will be veridical.

⁶ Cohen (2010) outlines a procedure where, for each potential item of future perceptual knowledge that p , the procedure could now give us a priori justification to believe [it appears to me that p] \supset [p is true]. The procedure is to suppose it appears to you that p , then infer p under the supposition, and then apply conditional proof. But, as I see it, this only shows that a certain justification's availability is *temporally prior*, not prior in the sense of *dependent on* that is at issue here. My knowledge or justification to believe that p has not been shown to be dependent on my knowledge or justification to believe the conditional.

4 Rationalism reformulated?

Wright (2002, 2004), White (2006), and Silins (2007) endorse views in the Rationalist family, but their explicit endorsements are of views formulated not in terms of *knowledge* but *justification for belief* (or for Wright, *warrant to accept*; see Wright (1991, 2004) and substitute as desired below; the differences won't matter here). The principle they've endorsed is this one:

Jp-Rationalism: necessarily, you have a *justified* perceptual *belief* that p , only if, for any considered proposition s that you know entails [not- p and it appears to you that p], your *justification to believe* not- s is antecedent to your perceptual *justification to believe* p .⁷

(Keeping the entailment *known* is harmless; it only makes the principle more defensible.) Why do these authors work with Jp-Rationalism instead of Rationalism? Their reasons are idiosyncratic. Wright (1991), inspired by Russell (1912), thought that it will be good enough if we can save justification (or warrant) from skepticism. And White and Silins were writing in response to Pryor (2000), who endorsed a Moorean view for knowledge but devoted more discussion to the corresponding view about justification. White and Silins also want to give a Bayesian critique of Mooreanism, which is more naturally presented in terms of justification.

Whatever their reasons, it doesn't matter. I'll now argue that any anti-skeptic who accepts Jp-Rationalism faces pressure to accept Rationalism (as originally formulated in terms of knowledge), and thus inherit Rationalism's problem accounting for how we could rule out *error* independently of perceptual knowledge that p .

The first step is just to observe that Jp-Rationalists cannot plausibly deny Jd-Rationalism ("d" for "doxastic" replacing "p" for "propositional"):

Jd-Rationalism: necessarily, you have a justified perceptual belief that p , only if, for any considered proposition s that you know entails [not- p and it appears to you that p], you *have a belief* that not- s that *is justified* antecedently to your perceptual justification to believe p .

If you have considered s , then you either believe it, disbelieve it, or are positively uncertain. Since you know p and s are inconsistent, believing or even being uncertain about s will obstruct your ability to have a justified belief that p . Having a justified belief that p requires you to disbelieve s , which (revisionary solutions to the semantic paradoxes bracketed here) we can equate with believing not- s .

⁷ For instance, here's the first sentence of White (2006)'s abstract: "I argue that its appearing to you that P does not provide justification for believing that P unless you have independent justification for the denial of skeptical alternatives—hypotheses incompatible with P but such that if they were true, it would still appear to you that P." In the final section of his paper, where he sketches his positive view, he presents the view as saying that "in order to gain perceptual justification for believing that P, we must have independent justification for believing that we are not victims of a visual illusion that P" (p. 552).

The next thing to see is that Jp- and Jd-Rationalists are committed to the following principle, at least given the assumption (which I shall help myself to⁸) that any necessary condition on having justification is also a necessary condition on having knowledge:

KJ-Rationalism: necessarily, you perceptually *know* that *p*, only if, for any considered proposition *s* that you know entails [not-*p* and it appears to you that *p*], you have a justified belief that not-*s* antecedently to your perceptual *knowledge* that *p*.

Even though Wright, White and Silins don't usually write in terms of KJ-Rationalism, I expect they would readily accept it (at least until they finish reading this paper!).

Now, KJ-Rationalism itself does not yet allow us to press the uncomfortable question that Rationalism faced regarding how we know one-off skeptical scenarios don't obtain. While you cannot use statistical inference (applied to default knowledge of your own reliability) to *know* that one-off scenarios don't obtain, maybe you can use statistical inference to have a *justified belief* that those scenarios don't obtain. Even if you can't *know* your lottery ticket won't win, maybe you can have a *justified belief* it won't. I'll concessively suppose so.

The problem is that it's not tenable to accept KJ-Rationalism without accepting Rationalism, at least it's not tenable if you're an anti-skeptic about perceptual knowledge. Here's why. The anti-skeptical KJ-rationalist agrees you can perceptually *know* that *p*. By Closure, he'll have to say that you can also *know* the negation of every considered skeptical scenario *s*, including the one-off scenarios, including *error*. Now, while Rationalism says you need to *know* that not-*s* antecedently to your perceptual knowledge that *p*, KJ-Rationalism only says you need a *justified belief* that not-*s* antecedently to your perceptual knowledge that *p*. So, the KJ-Rationalist agrees you *do* have to know not-*s* (because of Closure), but if he *denies* Rationalism, then he says you *don't* have to know it *antecedently* to knowing *p*. The KJ-rationalist can make his usual appeal to default knowledge of your own reliability to explain how you have an antecedently justified belief that not-*s*, because statistical inference can license such a justified belief. But, as the critique of Rationalism showed, statistical inference cannot license *knowledge* that not-*s*. The KJ-Rationalist who denies Rationalism thus faces an explanatory gap: how does a *merely justified* belief that not-*error* become *knowledge* that not-*error* when the only change in the meanwhile is that it perceptually appears to you that *p*? Representing not-*error* as a material conditional, the question is what explains how, just by having a perceptual experience, you can go from having a merely justified belief to having knowledge of this: [it appears to me that *p*] \supset [*p* is true]. A Moorean can happily say that a bit of Moore-style reasoning, premised on your perceptual knowledge that *p*, inferentially leads to knowledge that not-*error*. But the Rationalist, who thinks any such reasoning premised on perceptual knowledge that

⁸ Though I view the assumption as a safe one to rely on, it has, like every philosophical claim, been questioned. A recent example is Kornblith (2008).

p is question-begging, cannot explain how your perceptual experience offers any new epistemic support here.

Many theorists are motivated to reject Mooreanism and move toward views in the Rationalist family because they want to deny that we can inferentially reject skeptical scenarios on the basis of perceptual knowledge. But then their choices will be either full-blooded Rationalism, or skepticism about perceptual knowledge.

A final note: I am aware that some philosophers, following Russell (1912), do not find it very difficult to say that we lack ordinary knowledge, as long as we can hold on to justified beliefs in the same claims. (As mentioned, Wright (1991) expresses sympathy for this, calling it “Russellian retreat” (p. 88).) Perhaps the position is made palatable with the thought that, with mere justified belief, we could still make our way through life in practice; we could make it to the grocery store and back, while skeptics about justified belief might starve. However, our concern as theorists is with what to believe. Is skepticism about knowledge the *true* view? Skepticism about knowledge is a radically revisionary view. Skepticism about knowledge is so dramatic a challenge to our most central common sense beliefs that it should be treated as the position of absolute last resort. Are the problems facing a Moorean anti-skeptical view so severe, do they threaten central common sense beliefs so seriously, that it would be better to give into skepticism about knowledge? I haven't here addressed the problems facing Mooreanism, but I find it hard to believe they are worse than knowledge skepticism. I would even, rather than concede that we have no ordinary perceptual knowledge, make any number of radical revisions to common sense to save a Rationalist view.

Acknowledgments For help improving this paper, thanks to Yuval Avnir, Jim Pryor, Karl Schafer, Roger White and especially Stewart Cohen.

References

- Adler, J. (1994). Testimony, trust, knowing. *The Journal of Philosophy*, 91, 264–275.
- Adler, J. E. (2002). *Beliefs own ethics*. Cambridge, MA: MIT Press.
- Alston, W. (1986). Epistemic circularity. *Philosophy and Phenomenological Research*, 47, 1–30.
- Burge, T. (1993). Content preservation. *The Philosophical Review*, 102, 457–488.
- Coady, C. A. J. (1992). *Testimony: A philosophical study*. New York: Oxford University Press.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65, 309–329.
- Cohen, S. (2010). Bootstrapping, defeasible reasoning, and *a priori* justification. *Philosophical Perspectives*, 24, 141–159.
- Cross, T. (2010). Skeptical success. *Oxford Studies in Epistemology*, 3, 35–62.
- Descartes, R. (1641/1984). *Meditations on first philosophy*. In J. Cottingham, R. Stoothof, & D. Murdoch (Eds.), *Philosophical writings of Descartes* (Vol. II). Cambridge: Cambridge University Press.
- Fricker, E. (1994). Against gullibility. In B. K. Matilal & A. Chakrabarti (Eds.), *Knowing from words*. Dordrecht: Kluwer.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Harman, G. (1965). The inference to the best explanation. *The Philosophical Review*, 74, 88–95.
- Harman, G. (1986). *Change in view*. Cambridge, MA: MIT Press
- Hawthorne, J. (2004). *Knowledge and lotteries*. New York: Oxford University Press.
- Huemer, M. (2001). *Skepticism and the veil of perception*. Lanham, MD: Rowman and Littlefield.
- Klein, P. (2000). Contextualism and the real nature of academic skepticism. *Philosophical Issues*, 10, 108–116.

- Kornblith, H. (2008). Knowledge needs no justification. In Q. Smith (Ed.), *Epistemology: New philosophical essays*. Oxford: Oxford University Press.
- Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
- Lewis, D. (1979). Counterfactual dependence and time's arrow. *Nous*, 13, 455–476.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- Lipton, P. (1992/2004). *Inference to the best explanation*. London: Routledge.
- Lipton, P. (2007). Alien abduction: Inference to the best explanation and the management of testimony. *Episteme*, 4, 238–251.
- Malmgren, A.-S. (2006). Is there a priori knowledge by testimony? *The Philosophical Review*, 115, 199–241.
- McDowell, J. (2008). The disjunctive conception of experience. In A. Haddock & F. MacPherson (Eds.), *Disjunctivism*. Oxford: Oxford University Press.
- McDowell, J. (2011). Perception as a capacity of knowledge. The Aquinas lecture, Marquette University Press.
- Moore, G. E. (1939). Proof of an external world. *Proceedings of the British Academy*, 25, 273–300.
- Peacocke, C. (2004). *The realm of reason*. Oxford: Oxford University Press.
- Pollock, J. (1974). *Knowledge and justification*. Princeton, NJ: Princeton University Press.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, 34, 517–549.
- Pryor, J. (2004). What's wrong with Moore's argument. *Philosophical Issues*, 14, 349–378.
- Reichenbach, H. (1939). *Experience and prediction*. Chicago: University of Chicago Press.
- Russell, B. (1912). *The problems of philosophy*. Oxford: Oxford University Press.
- Schiffer, S. (2003). *The things we mean*. Oxford: Oxford University Press.
- Silins, N. (2007). Basic justification and the Moorean response to the skeptic. *Oxford Studies in Epistemology*, 2, 108–142.
- Stalnaker, R. C. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory. American philosophical quarterly monograph series* (Vol. 2). Oxford: Blackwell.
- Vogel, J. (1990). Cartesian skepticism and the inference to the best explanation. *The Journal of Philosophy*, 87, 658–666.
- Vogel, J. (1999). The new relevant alternatives theory. *Philosophical Perspectives*, 13, 155–180.
- White, R. (2005). Explanation as a guide to induction. *Philosopher's Imprint*, 5, 1–29.
- White, R. (2006). Problems for dogmatism. *Philosophical Studies*, 131, 525–557.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, C. (1991). Scepticism and dreaming. *Mind*, 100, 87–116.
- Wright, C. (2002). (Anti-)sceptics simple and subtle. *Philosophy and Phenomenological Research*, 65, 330–348.
- Wright, C. (2004). Warrant for nothing. *Proceedings of the Aristotelian Society, Supplementary Volume*, 78, 167–212.
- Wright, C. (2007). The perils of dogmatism. In S. Nuccetelli & G. Seay (Eds.), *Themes from G.E. Moore*. Oxford: Oxford University Press.