

SKEPTICISM ABOUT CONTEXTUALISM

A Thesis

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Master of Arts

in

The Department of Philosophy and Religious Studies

by

William G. Osborne III
B.A. Millsaps College, 1997
M.F.A., Louisiana State University, 2000
December 2005

Acknowledgements

Numerous professors in the LSU philosophy department have helped to bring this project to completion. Over and above his peers, Dr. Jeff Roland has shown with patience, exactitude, and good humor the model of professor working with student. Above the many merits of Professor Roland and his colleagues, however, I must cite my wife Annie as the greatest aid and influence I've known—despite, or perhaps on account of, her lack of patience for brains-in-vats.

Table of Contents

Acknowledgements.....	ii
Abstract.....	iv
Introduction.....	1
External World Skepticism.....	3
Stewart Cohen’s Contextualist Response to EWS.....	7
The Problem of Induction.....	15
A Contextualist Solution to the Problem of Induction.....	22
Conclusion.....	26
References.....	27
Vita.....	29

Abstract

Contextualism is the epistemological thesis that holds context to significantly affect the truth value of claims such as “S knows that p .” A shift in context can lead to a shift in the standards by which we evaluate propositional knowledge claims, and thus a shift in the truth values of these claims: a statement “S knows that p ” may be true when evaluated in one context while simultaneously false when evaluated in another context. A contextualist says it is by playing on these shifting standards that the skeptic manages to destroy knowledge with her skeptical arguments. Once this is understood, the contextualist holds, we can see that skeptical arguments like those of external world skepticism do not threaten our ordinary knowledge.

In this thesis I argue that if contextualism is as successful as it claims, then it should be possible to use the techniques of contextualism to solve skeptical problems other than those with which the contextualist is usually concerned; namely it should be possible to solve the problem of induction. Taking Stewart Cohen’s *relevant alternatives* derived contextualism as my representative, I show that it is possible to construct such a solution consistent with the tenets of contextualism, but that in doing so certain troubling question-begging assumptions are shown in sharp relief, invalidating the proposed solution. In the end we are left with the question: Can a relevant alternatives contextualism successfully solve any skeptical argument without a proper solution to the problem of induction?

Introduction

Contextualism in epistemology attempts to defuse skeptical challenges to knowledge claims, usually those challenges concerning knowledge of the external world. There are many different contextualist approaches to this task, but we can conceive of contextualism in general as preserving our intuitions about having knowledge by arguing that a skeptical challenge only *appears* to endanger everyday knowledge claims because of shifts in epistemic standards that accompany shifts in types of conversational contexts.¹ What “to know” means in one situation is different, the contextualist holds, than what “to know” means in another situation: knowledge for a contextualist is thus completely relative to contexts of knowledge attribution. Once these shifting standards are recognized, the contextualist argues, we can see that the truthfulness of a knowledge claim depends on the standards determined by different contexts. The skeptic may well be correct in denying a knowledge claim in the stringent environment created by skeptical issues, but the contextualist counters that when we are in ordinary, everyday, non-skeptical contexts, the truthfulness of the same claim is upheld. Contextualists conclude that our normal knowledge claims are in the end safeguarded from the threat of skepticism.

The skeptical problem about the external world is a familiar fixture in epistemology. Modern philosophy is traced to Descartes’ meditations where he explored the possibility of being deceived by an evil god-like genius into holding beliefs about an external world that do not correlate with the way things actually are. Our senses are on occasion deceived in rather trivial ways—a stick appearing bent in water, hallucinations, dreams, *etc.* Descartes recognized that from these examples it follows that it is possible to be deceived all of the time. And not being able to rule out this possibility, it appears one cannot be fully justified in asserting the contrary; *i.e.*, that things *are* as they seem. Descartes appealed to a benevolent creator to escape potential solipsism. This argument, though, is widely seen as folly. Numerous subsequent attempts to justify beliefs about an external world have followed Descartes. None, however, has won widespread acclaim.

The contextualist attempt to disarm external world skepticism appeals to the supposed relative nature of “to know” in order to secure our knowledge claims about the external world. In a very skeptical context the standards for knowledge are very high—impossibly high, says the contextualist—rendering any claim based on the existence of the external world false. On the other hand, the same claim—all things being equal—is seen to be true in a context with much lower standards, for it reaches these lower standards.

Note, though, the following premise in the contextualist argument: our everyday claims to knowledge are safe when they are examined in the light of a non-skeptical context. Aside from appeals to common sense and ordinary language use, this claim is offered without any support. “No doubt some will feel it is ad hoc,” admits the contextualist Stewart Cohen. But, Cohen maintains, all one “can hope for is a workable option that enables us to avoid skepticism. We avail ourselves of that option for no other reason than we want to avoid skepticism.”² Cohen has stated elsewhere that contextualism only has purchase “if we are antecedently convinced of the falsity of skepticism.”³ By this account, the falsity of skepticism (the “ad hoc” option) is

¹ James Pryor characterizes contextualism as “the doctrine that the truth of a knowledge-ascription can vary with the context in which that ascription is made or assessed.” James Pryor, “Highlights of Recent Epistemology,” *British Journal for the Philosophy of Science* 52 (2001): 100.

² Stewart Cohen, “Knowledge and Context,” *Journal of Philosophy* 83 (1986): 582.

³ Stewart Cohen “Contextualism, Skepticism, and the Structure of Reasons,” *Philosophical Perspectives* 13 (1999): 83.

used as an explicit premise in the contextualist attempt to safeguard knowledge from the threat of skepticism. Again, Cohen: “But doesn’t this just beg the question against the skeptic? In a sense it does, but no more than the skeptic begs the question against common sense.”⁴ Leaving aside whether or not one can justifiably beg the question, I’d like to examine how it is that one can view contextualism as question-begging and look at what this might show us about their project.⁵

Many philosophers have taken issue with that which the contextualist approach helps itself in order to accomplish its task. Contextualism’s detractors have examined various points of contention—contextualism’s irrelevance to epistemological problems (Kornblith),⁶ its unsupportable semantics (Schiffer),⁷ its concern with meta-linguistics as opposed to epistemology (Sosa),⁸ and its presumption of semantic support for knowing as an indexical (Stanley)⁹ among other concerns. Here I’d like to demonstrate how the contextualist question-begging approach becomes acutely evident—and philosophically interesting—when their techniques and analyses, developed as a response to the skeptical problem of the external world, are applied to another important skeptical argument, the problem of induction. I will take a particular form of contextualism, as offered by Stewart Cohen, and see how it might be taken to provide a solution to the problem of induction. We will see that his contextualist approach explicitly begs the question in regards to the problem of induction. This question-begging will show in sharp relief the question-begging inherent in Cohen’s approach to external world skepticism and suggest that question-begging is a greater problem than Cohen or any other contextualist has heretofore recognized.

The structure of my project is as follows: I begin with the problem of the external world (in “External World Skepticism”) and present in detail one contextualist approach to solving external world skepticism (“Stewart Cohen’s Contextualist Response to EWS”); I then lay out the problem of induction (in “The Problem of Induction”); following this I detail a contextualist solution to that problem based on Cohen’s work and then weigh the efficacy of such a solution (in “A Contextualist Solution to the Problem of Induction”). Here we see the problems generated by this approach, at which point I will examine the implications of my project.

⁴ Stewart Cohen, “How to be a Fallibilist,” *Philosophical Perspectives* 2 (1988): 113.

⁵ For philosophers who embrace contextualism, what I am suggesting is question-begging is decidedly uncontroversial. It’s as matter-of-fact as the common sense and ordinary language that is contextualism’s starting point. The larger question is ultimately whether the linguistic concerns of contextualism have any bearing on epistemological concerns.

⁶ Hilary Kornblith, “The Contextualist Evasion of Epistemology,” *Philosophical Issues (Nous)* 10 (2000).

⁷ Stephen Schiffer, “Contextualist Solutions to Skepticism,” *Proceedings of the Aristotelian Society* 96 (1996): 317-333.

⁸ Ernest Sosa, “Knowledge in Context, Skepticism in Doubt,” *Philosophical Perspectives* 2 (1988): 139-155.

⁹ Jason Stanley, “Context and Logical Form,” *Linguistics and Philosophy* 23 (2000): 391-434.

External World Skepticism

External world skepticism (EWS) denies one can have knowledge of an external world, where knowledge of the external world consists of having a justified, true belief¹⁰ with regard to some proposition concerning a world existing outside of the experience of the mind (*e.g.*, that there are people; that in the United States fire trucks are often red; that my office is cold today; or even, that there exists an external world). More precisely, EWS is the position according to which, if *p* is any proposition concerning an external world, no subject *S* does not know *p*.

The EWS skeptic presses for sufficient reason to grant that a subject has knowledge of an external world—reason the skeptic maintains lies well beyond the reach of the putative knower. In general, a would-be knower cites sensation-based evidence as grounds for knowledge of an external world—“I know something about some object or attribute because I have experienced it through sight, smell, taste, touch, *etc*”: a subject relies upon the subjective phenomena of sensory experience to justify her beliefs about an objective external world. The skeptic asks for justification for the leap from this mental experience to claims about a world external to the mind: seeing flames, feeling the warmth of burning wood, smelling charred logs, and hearing the crackle associated with a fire are insufficient, says the EWS skeptic, for claiming knowledge that there’s a fire nearby.

In his first *Meditations*, Descartes presents what has become the traditional argument for EWS.¹¹ Descartes sought to secure knowledge in light of skeptical challenge. To do so, he systematically examined what was taken to be evidence for knowledge claims and evaluated the threats levied against that evidence. In the text, Descartes is first troubled with the possibility of illusions—if one has been mistaken before, why not now?—and then raises the stakes with the potential of a subject dreaming; as Descartes continues his project, the stakes are continuously raised and security of knowledge correspondingly lessened—more and more the relationship between what we think we know and what is actually the case is called into question.

Descartes’ evil genius emerges at the apex of these threats: instead of individual cases of perceptual knowledge, the evil genius threatens knowledge of the external world taken as a whole—all instances of knowledge of the external world—conceptual as well as perceptual—whether or not one can have knowledge of an external world in *any* capacity. To perform such a task, the evil genius orchestrates events so that a subject’s phenomenal life appears just as it would were the subject actually experiencing an external world. This of course results in an

¹⁰ I am here ignoring Gettier cases, situations of “accidental” knowledge that satisfy the traditional tripartite conditions for knowledge (*i.e.*, justified true belief). In a Gettier case, I may believe something that is true and have justification for that true belief, but the justification is based on reasons that don’t justify the belief in the way we think justification ought to do so: our justification seems to work by accident by accident. Say I have the true belief that someone in my office plays the cello and has recently purchased one of these instruments (belief). My belief is based on the following reasons: I notice a coworker has begun to carry cello scores around; I see that she has a cello in the back of her car; and I overhear her mention the going rate for a cello to another colleague. Now, it is actually the case that this colleague has *not* recently purchased a cello—she merely rents one—and does not even play one: the one in her car is for her child, she checks out scores for her child on her lunch break, and she is aware of the prices for a decent cello because she cannot afford one for her child. But, as it turns out, another colleague, one I wouldn’t even suspect has any interest in music, has purchased and begun to play the cello. The problem that arises is that I have the true belief that there’s a cello player in my office who recently acquired a new cello, and I have justification for this true belief, but it seems counter-intuitive to call this knowledge, as the justification seems the wrong sort, and the true belief accidental. I recognize the importance of Gettier cases, but they need not factor in to my project.

¹¹ Rene Descartes. *Selected Philosophical Writings* (New York: Cambridge University Press, 1998).

apparent—albeit mistaken—justification for beliefs about that apparent external world. If the world appears to us as though it has the features we normally take the world to have, we readily respond by thinking it legitimate to hold beliefs about the world and make knowledge claims to the effect that things are as they appear. Seeing that we would still maintain claims about a world *even* when that world did not even exist (outside of the inspirations of the evil genius), it becomes clear that we lack justification for our beliefs and any claims we might make about an external world. And so, in demonstrating the *insecurity* of a subject’s sensory-based knowledge, the evil genius scenario raises the threat of EWS.

Hilary Putnam provides an updated version of EWS with a “brain-in-a-vat” (BIV) scenario. The BIV scenario has become the subject of and platform for much philosophical debate concerning EWS and as such the argument based on the BIV scenario is *the* argument to disarm in the contemporary landscape of epistemological skepticism. Quite simply, BIV is the Cartesian evil genius scenario *sans* evil genius—as we’ll see, technology plays the role here. Putnam invites us to entertain the scenario that—independently of an evil-genius-like agent—brains-in-vats have arisen through certain cosmological events to become the status quo: brains-in-vats that have always, for as long as they’ve been brains at all, been brains-in-vats, i.e. brains that have never been embodied. In this scenario, the brains’

nerve endings [are] connected to a super-scientific computer which causes the [brains] to have the illusion that everything is perfectly normal. There seem to be people, objects, the sky, etc.; but really, all the [brains are] experiencing is the result of electronic impulses traveling from the computer to the nerve endings.¹²

The argument that results from this scenario is a non-agential argument with the same driving point as the Cartesian scenario: a subject S cannot rule out that she is a brain suspended in a vat of nutrients electrochemically stimulated so as to produce a simulation of an external world—all of her sensory experiences, memories, etc. are the result of this stimulation, so she cannot get “outside” of this information to verify its source, and as a result she has no grounds to rule out a scenario such as being a BIV; but, if S doesn’t know she’s not a BIV, then S doesn’t know the things we would take her to know were she not a BIV (such as whether she has hands or not); one can therefore conclude, it seems, that S does *not* know she has hands. Let’s unpack this a little and see how the BIV scenario and its related argument give us EWS.

I rely upon the sum of the sensations I am currently experiencing—the sound of the air conditioner, the color of the flowers outside my window, the feel of my back against my chair, a soreness in my *hands*, even a general sense of hunger—to situate myself bodily within an external world. This world I take to consist of things that reasonably correspond with these sensations; in this case, I take it that I am in my office late Sunday morning. Colors, sounds, bodily sensations—these all suggest to me that there are things external to my mind that have the very attributes of color, timbre, etc. that I take to cause my mental sensations—I feel a keyboard with my hands, I see that my typing brings words to the screen, I can see fleshy protuberances dangling in front of me—with all of this, I feel it reasonable to assert that it is by virtue of having hands that I experience this sensory information. But, all the same, it seems I cannot rule out the possibility that I am actually *only* a brain-in-a-vat which is being fed information in a way other than by contact with the world as I experience it. And if I were *only* a brain, then it is clear I

¹² Hilary Putnam, *Reason, Truth, and History* (Cambridge: Cambridge University Press, 1981), 6-7.

would have no body, and with no body I would have no hands. So, given I can't rule out that I'm merely a brain, it seems to follow that I do not know I have hands. Also, I therefore do not have any justification for my belief that I am in my office on Sunday. Indeed, I have no justification for *any* claims to knowledge concerning the external world. In fact, it appears I should be skeptical of anyone having justified knowledge of an external world at all.

These two EWS arguments, *viz.* the Cartesian evil genius argument and the argument based on the BIV scenario, are obviously very similar—so similar in fact these sorts of arguments share a structural form:

- P1. S does not know that not-H.
- P2. If S does not know that not-H, then S does not know that O.
- C. S does not know that O.

Following Keith DeRose, we can call these types of arguments “Arguments from Ignorance” (AI), a name that notes a defining feature of these arguments, *i.e.* that we are not easily if at all able to rule out some hypothesis that if the case threatens our knowledge.¹³ The AI form is the structure contemporary theorists utilize when considering EWS arguments. Accordingly it will be useful for us to have the argument from BIV in this form. Here is the BIV argument mapped onto the schematic structure of AI:

- P1. I don't know that I'm not a BIV.
- P2. If I do not know that I'm not a BIV, then I don't know that I have hands.
- C. I don't know that I have hands.

Responses to EWS

EWS has prompted a range of responses. These responses tend to seize on one element of the AI argument—either a premise or the conclusion: G.E. Moore, for instance, denies the conclusion outright, appealing to an intuitive sense of its absurdity—“here are my hands.” Putnam denied the first premise, his argument based on semantic externalism: on his view, the nature of language, reference, and meaning precludes a scenario such as BIV. Fred Dretske and Robert Nozick—both holders of a relevant alternatives view of knowledge—deny the second premise, the conditional that leads to the conclusion; they see this conditional as an instance of the “closure principle,” and in the face of an argument that otherwise seems compelling, feel obligated to deny this principle in order to deny the skeptic.¹⁴ All of these responses are characterized by an outright denial of one of EWS's premises or its very conclusion or some combination thereof. These responses are generally known as “Moorean” responses following Moore's example.

Contextualism, which we will explore in greater detail in the next section, takes the non-Moorean approach that each of the members of the argument *is* acceptable—at least intuitively in the sense we don't appear to have grounds to deny them—but that a consideration of the *contexts*

¹³ Keith DeRose, “Solving the Skeptical Problem,” *Philosophical Review* 104 (1995): 1. DeRose elsewhere refers to arguments like these as “arguments from skeptical hypothesis” Keith DeRose, “Introduction: Responding to Skepticism,” in *Skepticism: a contemporary reader*, ed. DeRose, Keith and Warfield, Ted A. (New York: Oxford, 1999).

¹⁴ For a discussion of both Dretske and Nozick, see DeRose's introduction to *Skepticism: a contemporary reader*: Keith DeRose, “Responding to Skepticism,” 18.

in which they are evaluated will clarify how arguments like BIV operate and thus allow us to disarm them. Once these types of arguments are properly understood, the contextualists hold that they no longer pose a threat to our ordinary knowledge claims.

Stewart Cohen's Contextualist Response to EWS

We sometimes have conflicting intuitions about what counts as knowledge: we find ourselves in situations in which two or more propositions, while independently plausible, appear to be contradictory when considered together. Situations of this type motivate an epistemological response that explains these contradictions while satisfying general intuitions on what counts as knowledge: epistemological contextualism is such a response. In this section, I explain the types of problems that recommend a contextualist approach by way of Fred Dretske's *painted mules* scenario, where the difficulty of knowing whether zebra are or are not painted mules brings conflicting intuitions to the fore.¹⁵ Taking Stewart Cohen's relevant alternatives derived contextualism as my contextualist representative, I then demonstrate the contextualist approach to solving such an epistemic quandary: roughly that S knows that p depending on the context in which knowledge is ascribed. Finally, the contextualist approach to the painted mules case is applied to the BIV scenario—giving us a contextualist response to EWS.

Painted Mules

Dretske's painted mules scenario illustrates the sort of conflicting intuitions about knowledge with which we're here concerned. If a zoogoer in front of a cage containing black and white striped, hoofed quadrupeds accompanied by a placard reading "zebra" claims "I know these are zebra" it seems *prima facie* a true statement: we can count it as knowledge. If another person standing nearby were to overhear the zoogoer, we would expect this person to ascribe the zoogoer's propositional knowledge claim with the truth value "true": it is the case that S (the zoogoer) knows that p (that these animals are zebra). And this seems unproblematic.

But things can get tricky: animals that look like zebra can possibly be painted mules. However unlikely, this alternative hypothesis to the animals actually being zebra is a distinct possibility until it is satisfactorily ruled out. The particularly tricky part, as Dretske observes, is that an alternative hypothesis like this characteristically plays on the *neutrality* of the available evidence: the evidence both supports the alternative hypothesis (these animals are painted mules) as well as the more obvious conclusion (the animals are zebra)—the evidence is *neutralized*.¹⁶ Cohen adds that the alternative hypothesis is *immune* to repudiation by appeal to this shared (thus neutralized) evidence for thinking the more pedestrian claim; the alternative hypothesis has *immunity* due to the evidence's *neutrality*.¹⁷

Imagine that the particular zoo our zoogoer is visiting is run by a *nefarious zookeeper*, one who only yesterday dispatched a *henchman* to replace the zebras with painted mules. The nefarious zookeeper, eavesdropping from his secret hideout as the zoogoer claims to know that the animals are zebra, exclaims: "That statement is false—those are not zebra." This is an interesting development: the zoogoer's statement, formerly true, is now clearly false—he does not know that the animals are zebra. But, unbeknownst to the nefarious zookeeper, his

¹⁵ Fred Dretske, "Epistemic Operators" in *Skepticism: a contemporary reader*, ed. Keith DeRose and Ted Warfield, 138 (New York: Oxford, 1999).

¹⁶ Dretske, 138.

¹⁷ Cohen, "Contextualist Solutions to Epistemological Prolems: Skepticism, Gettier, and the Lottery," 67. Elsewhere, DeRose, after Nozick, characterizes these hypothesizes as having the effect that you would still not believe them if they were true: one would still believe the zebra in the cage were zebra even if the alternative was true. This is termed the sensitivity of a belief: how amenable is the belief to contrary evidence establishes its sensitivity. This can be seen to play on the immunity/neutrality characteristic. This issue of the sensitivity of a belief is discussed in DeRose, "Responding to Skepticism," 18.

henchman couldn't bring himself to perpetrate such zoological fraud—the zebra *are* zebra after all. The henchman overhears the nefarious zookeeper and thinks: “Well, the nefarious zookeeper is incorrect, those *are* zebra, but it's not clear that the zoogoer is correct either—it may be the case that they are zebra, but they *were* almost painted mules, and it's odd to say someone knows something when it seems so by accident.”¹⁸

In this scenario, the same propositional knowledge claim (the zoogoer knows the animals to be zebra) has shifted through a range of truth values without any fact of the matter having changed. The zebra never stopped being zebra—rather we moved through various vantage points and saw how the truth-values of a knowledge claim about the zebra were affected by virtue of the situation.

One propositional claim being both true and false is problematic for an epistemic view—indeed a contradiction in any area of philosophy asks for some scrutiny. Recognizing that an epistemological view should not casually assent to inherent contradictions, we can see that dissolving the painted mules case would recommend an epistemological viewpoint that could do so. Cohen's relevant alternatives-based contextualism promises to be able to do just that.

The relevant alternatives (RA) elements of Cohen's view explain how knowledge is destroyed in a scenario such as painted mules. Given a zoogoer's usual situation—visiting a zoo *not* run by a nefarious zookeeper—a knowledge claim that S knows that animals are zebra *can* be taken to be true, *despite* the possibility that these animals could be painted mules. In general, we don't have to rule out such an outlandish claim—the alternative just simply is not *relevant*, and due to its irrelevance, there is no need to consider it in light of evaluating a knowledge claim.

For Cohen, an alternative is relevant if a subject S is in an epistemic position where an alternative to p would prevent S knowing p.¹⁹ On Cohen's view, relevance is determined by statistics: the probability of an event or situation occurring determines its relevance.²⁰ For example, if you come across a barn in the countryside, the possibility of that barn being a replica made of papier-mâché is generally absurd enough to be irrelevant. But if there was a big-budget blockbuster about barn-burning being filmed in the area, then the alternative takes on a significant relevance due to its statistically significant possibility. And for Cohen, that relevance can (at least theoretically) be determined computationally by factoring in the ratio of fake barns to real barns in the area and other related information.

Now, there is a flipside to the issue of statistically determined relevance: an alternative can be *made* relevant by drawing a subject's attention to a hitherto irrelevant alternative—one can make an irrelevant alternative hypothesis relevant by making the alternative *salient*.²¹ For example, given a fair lottery, to say “I know this is a losing lottery ticket” *before* there has been a drawing seems intuitively false: one doesn't seem to *know* a ticket is a losing ticket until the drawing is actually held—people who buy lottery tickets do so because they know there's *some probability* that a ticket is a winning ticket. At first this appears to be a problem for Cohen's theory of statistically-determined relevance—there is a very low probability that the ticket is a winning ticket (and note that we can make the probability as low as we like by manipulating the total number of tickets), but yet it *is* still a very relevant alternative (people buy tickets so that they might win)—but actually this seemingly paradoxical situation illustrates quite nicely

¹⁸ This could more clearly made to be an example of a Gettier case.

¹⁹ Black, Tim. “Contextualism in Epistemology,” *The Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/c/contextu.htm> (Accessed September 2004).

²⁰ Cohen, “How to be a Fallibilist,” 106

²¹ Cohen, “How to be a Fallibilist,” 106

Cohen's rule of salience: an otherwise irrelevant alternative can be made relevant by drawing a subject's attention to that alternative.²²

Statistically-determined relevance and the rule of salience serve to explain the conflicting intuitions about what is true in the painted mules scenario. For a zoogoer visiting a zoo *not* run by a nefarious zookeeper the probability of zebra being painted mules is extremely low—low enough not to endanger the knowledge claim “The zoogoer knows that these animals are zebra.” In a scenario where the nefarious zookeeper periodically replaces animals with painted stand-ins, however, this becomes a relevant alternative. First, notice that in a scenario where the nefarious zookeeper *has* replaced the zebra, we say that the probability is (clearly) high enough to make the alternative relevant and we clearly cannot count the zoogoer's claim as knowledge. But also notice that even when the henchman does *not* replace the zebra as ordered, it *still* seems correct to deny the zoogoer's knowledge claim: the zoogoer does not know the animals are zebra because he has not ruled out what is a very relevant alternative—not only *could* these animals be painted mules in some far-reaching sense, they *almost were*.

RA Contextualism and Preserving Closure

If relevant alternatives theory alone can explain the variety of scenarios encountered above, then the question arises: What recommends Cohen's theory over RA? Cohen's theory can be seen to be preferable over RA alone on account of two features: *context* and *closure*. For our purposes, context can be taken to indicate the sum of relevant elements (such as place, time, and other considerations) involved in fixing an environment of ascription. Closure is a bit more complicated a matter. To understand the issue, we will first put the painted mules scenario into its AI form.

- P1. S does not know that the zebra-like animals in the zebra cage are not cleverly painted mules.
- P2. If S does not know that the zebra-like animals in the zebra cage are not painted mules, then S does not know that they are zebra.
- C. S does not know the zebra-like animals in the zebra cage are zebra.

The *closure principle* is exhibited in P2: the principle states that to count as knowing p (the animals are zebra) when p entails that q (the animals are not painted mules), then you do not count as knowing p if you do not know q . Dretske's RA theory rejects “the principle that if you do not know that Q is true, when it is known that P entails Q , then you do not know that P is true.”²³ Cohen, however, believes that closure “expresses something deep about the nature of knowledge.”²⁴ His contextualism preserves closure while maintaining the RA approach.²⁵

Earlier, we took the slight variations of the painted mules scenario individually—in turn the zoogoer, the nefarious zookeeper, and the henchman; Cohen's contextualism accounts for these variations all at one instant by considering the contexts in which these knowledge claims are assessed, and in doing so Cohen's theory maintains the closure principle. The contextualist analysis says the zoogoer both counts as knowing *and* as not knowing the animals are zebra,

²² Cohen “How to be a Fallibilist,” 92.

²³ 139 skepticism

²⁴ Cohen, “Contextualism, Skepticism, and the Structure of Reasons,” 64. In addition, Cohen asks “If the closure principle is false, then what explains the appeal of skeptical arguments?”

²⁵ Cohen is following RA theorist Gail Stine in preserving closure against Dretske and Nozick, although he disagrees with her non-evidential-based approach (Cohen, “How to be a Fallibilist,” 98-101).

depending on the context in which knowledge is being ascribed. On a contextualist view, knowledge is an *indexical*—it takes its meaning by virtue of the context in which it is uttered (or assessed),²⁶ the context determining the standards for ascription—so the same propositional knowledge claim can have radically different truth values depending on shifts in the contexts of ascription: in the context in which his statement is assessed by a bystander, the zoogoer counts as knowing the animals are zebra—for two unsuspecting zoogoers, the claim that these animals are zebra is common sense.²⁷ In the context of the nefarious zookeeper, the zoogoer does not count as having knowledge, for there is now a highly relevant alternative that the zebra are painted mules, and as the zoogoer does not know these animals are not painted mules, he by virtue of closure, does not know they are zebra. And in the henchman’s context, the zoogoer does not count as knowing because, although the zoogoer does hold a true belief, the belief is not properly justified on account of the salient possibility that has not been properly ruled out—again by closure, he does not count as having the knowledge in question.

There is no over-arching fact of the matter whether or not the zoogoer has knowledge—it is entirely dependent on context. And this is how Cohen preserves closure: on account of shifts in context, there is only the *appearance* of closure failure—when the zoogoer’s claim is evaluated in light of different epistemic standards the zoogoer both knows that *p* (when the alternative is not relevant) and does not know that *p* (when the alternative is made salient)—and these shifts allow for closure to function dependant on different standards in different contexts. The closure principle instantiated in P2 is always the case, but only affects the zoogoer’s knowledge claim *qua* P2 when the alternative hypothesis that the zebra are painted mules is under consideration such as in P1. But in an ordinary, everyday context, standards for knowledge are relaxed, and closure isn’t in effect *qua* P2.

Painted Mules and BIV, Cohen’s Contextualism

Notice what we’ve been considering with the painted mules scenario is an instance of skeptical argument and what we’ve been discussing with Cohen’s contextualism is a response to a skeptical argument. It should be clear as well that the argument and the response generated are concerned specifically with arguments and responses of the AI form. Another observation to be made is that if one skeptical argument of the AI form (the painted mules argument) can be satisfactorily overcome through a certain approach (contextualism), then another skeptical argument of the same form should also be susceptible to such an approach. If contextualism works with the painted mules skeptical argument, then it should work with an EWS argument of the same AI form. In other words, contextualism should be able to provide an answer to the AI argument from the BIV scenario. This is of course not lost on the contextualists, and accordingly, Cohen has worked out just such an argument. Below I will clarify the structural similarity of the painted mules scenario and the BIV scenario, examine the differences of these two situations, and explain how, just as Cohen’s contextualism works to provide a solution for

²⁶ For support, Cohen points to the word ‘flat’ and similar context-dependant natural predicates used in language. The lawn is flat in regard to whether or not it’s a good space to serve as a soccer pitch. The lawn is not flat in regard to whether or not I can spin a coin on it. Cohen says the land around Princeton is not flat, but recognizes a giant would disagree. When we consider flatness ascriptions, we quickly see different contexts – and thus different standards – for flatness. Flatness simpliciter depends on the context of the attribution. According to the contextualist, ascriptions of ‘S knows that p’ likewise vary from context to context. For them, as there is no over-arching flatness simpliciter, there is likewise no over-arching knowledge simpliciter.

²⁷ Cohen concludes that a “social component of knowledge...[indicates] that attributions of knowledge are context sensitive.” Cohen, “Knowledge and Context,” 574.

the painted mules argument, so in the same way does it provide a solution to the BIV scenario, and thus EWS.

We can see that the two skeptical arguments from BIV and the painted mules scenarios are structurally similar in that they are both arguments of the AI form. We'll look at the structure of AI in two ways: one that focuses our attention on the closure principle as encountered above, and a second that focuses on the alternative hypotheses aspect of AI.

The AI-type argument is effective because it focuses our attention on the intuition about knowledge that if q follows from p and it's the case that S knows that p , then S counts as knowing that q . Say every time the power goes out in my house (p), then my bedside lamp is not on (q). If I know the power is out in my house (p), then I count as also knowing that the bedside lamp is not on (q). Arguments of the AI-type seize upon this aspect of knowledge. In the way of looking at the AI-type arguments that emphasizes the closure principle we see that these arguments place doubt on whether we know something that follows from knowing something else. In the way of looking at the AI-type arguments that emphasizes alternative hypotheses, we see that these arguments place doubt on whether we've ruled out some alternative to what we ordinarily take as knowledge. In both cases, p refers to the *ordinary* explanation for a propositional knowledge claim while q stands for the *negation* of an *extraordinary* explanation of the claim. If S know that p , then S knows that q : if S knows that the ordinary proposition is the case, then S knows that the extraordinary proposition is not the case (the negation of the extraordinary proposition). The argumentative structure of AI can be put like this:

- P1. S does not know q
- P2. if S does not know q , then S does not know p
- C. S does not know p

Reducing the argument further, removing the representation of knowledge claims and presenting the structure in its barest form, gives us:

- P1. not- q
- P2. if not- q , then not- p
- C. not- p

With "ordinary" for p and "not-extraordinary" for q we get

- P1. not-(not-extraordinary)
- P2. if not-(not-extraordinary), then not-(ordinary)
- C. not-(ordinary)

In the way of looking at AI that emphasizes closure, we can see P2 as the contraposition of "if p , then q " where if S does not know the consequent q , then S does not know the antecedent p (a case of modus tollens). In emphasizing this aspect of the argument, it is offered that a subject S does not know that q , and that since q follows from p , thus S does not know p . In the case of BIV, if S does not know she's not a BIV (q), then S does not know she has hands (p), for if S were to know she has hands, then S would know she is not a BIV (if p , then q), since a BIV of course has no hands. Likewise, in the case of the painted mules, if S does not know the animals are not painted mules (q), then S does not know the animals are zebra (p), for if S were to know

the animals were zebra, then S would know they were not painted mules (if p , then q), for painted mules are of course not zebra. As we've seen, this forces some relevant alternative theorists to deny closure.

In the second way of looking at the AI argument, we focus on the aspect that brings two alternative hypotheses, the ordinary and the extraordinary, to the fore and the fact that if you cannot rule out the extraordinary, *i.e.* if you don't know that it's not the case that the extraordinary hypothesis obtains, then you can't claim to know the ordinary proposition. In this way of looking at AI, the emphasis is on how an ordinary belief that we might take to be justified on a certain line of evidence is in conflict with an alternative, extraordinary explanation that also fits the evidence. In this sense, what's at issue in the AI-type argument is " $(p$ or not- q) & not- $(p$ & not- q)": it's exclusively either the case that p (the *ordinary* proposition, and therefore by modus ponens, q , the negation of the *extraordinary* proposition) or it's the case that not- q (*not* the negation of the *extraordinary* proposition²⁸ and therefore by modus tollens, not- p , the negation of the ordinary proposition). So it's either the ordinary proposition is the case or the extraordinary proposition is the case, and since we cannot rule out the extraordinary proposition due to its immunity, the argument delivers us with the conclusion that we do not have knowledge of the ordinary proposition. So in both of the skeptical scenarios we're examining, BIV and the painted mules, it plays out that one claim or the other is correct; you either know what you ordinarily take to be knowledge or you do not. The evidence for thinking I know I have hands is either from an external world that includes my hands or the result of an alternative hypothesis like the BIV scenario; the evidence for thinking I know that zebra-like animals in a zebra cage are zebra is either derived from the fact that the animals are zebra or some alternative hypothesis like the painted mules scenario. These arguments play on the fact that there possibly could be another explanation for our epistemic experience.

Although they are structurally isomorphic in argument form, Cohen significantly differentiates the painted mules and BIV scenarios: the painted mules scenario is an example of a *restricted* version of skepticism contrasted with the *global* BIV scenario.²⁹ As we'll see, these distinctions hinge on the extent of the neutrality of the evidence of each scenario and the corresponding immunity of the alternative hypotheses involved. As will also become clear, the variance of knowledge in a range including ordinary and extraordinary accompanies a varying degree of context, ordinary and skeptical: in an ordinary context, those things that pass for knowledge are quite ordinary, the things people commonly presume; in a skeptical context, however, when extraordinary hypotheses are being considered, those things that ordinarily pass for knowledge fall short of the mark.

Restricted skepticism is a form of skepticism that applies in a restricted manner, such as a line of argument limited to a particular topic (e.g. skepticism about the existence of God) or situation (e.g. S doesn't know the animals that look like zebra aren't actually painted mules). In these cases, the evidence cited for knowing about the particular topic, situation, or process is rendered neutral; the same evidence is simultaneously support for the alternative hypothesis.

Since any positive evidence for an ordinary claim is neutralized in this manner, Cohen tells us that we can appeal outside of the scope of the particular restricted skeptical issue, outside of the neutralized domain, and utilize *inductive evidence* in order to circumvent the skeptic. In regard to the restricted skeptical problem of the painted mules, we can simply appeal to things we know inductively about the world, which includes such things as being unfamiliar with any

²⁸ (In other words, not- q is, following the elimination of the double negation, the extraordinary proposition, itself).

²⁹ Cohen "Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery," 67.

widespread deception occurring in zoological exhibits, especially in such a crude fashion as painting mules. In the case of another example of restricted skepticism, the papier-mâché barns scenario encountered earlier, we can generally appeal to inductively acquired ordinary knowledge about the world: people don't regularly create fake barns. We can count as knowing the falsity of these hypotheses based on past experience of the behaviors of people, institutions, and so on: it's difficult to imagine the motivation for such schemes.

Where restricted skepticism is localized, global skepticism is all-encompassing, its reach is universal—global. In a global skeptical situation, sufficient evidence can never be acquired to establish the falsity of the skeptical argument—all evidence is neutralized and the skeptical thesis is completely immune to refutation on the grounds of this neutralized evidence. Witness the BIV argument, Cohen's example of a global skeptical thesis: all evidence remains neutral in an attempt to determine whether or not one is a BIV—any appeal to empirical evidence is restricted.

Restricted and global skepticism are further affected by consideration of context: in what sorts of context are these hypotheses entertained? A restricted skeptical argument—on the contextualist view—is inconsequential in an ordinary context: by and large a subject can appeal to inductive evidence to establish a justified true belief.³⁰ In a skeptical context, however, that same argument successfully denies knowledge to the subject—she cannot appeal to inductive evidence. With a global skeptical argument—such as an argument based on the BIV scenario—no amount of evidence will secure an ordinary claim.

Whereas with cases of restricted skepticism in ordinary contexts a subject can appeal to inductive evidence, Cohen initially insisted that with cases of global skepticism it is “non-evidentially rational” for a subject to deny a skeptical thesis: without any evidence, a subject S can deny a global skeptical hypothesis. “The skeptic would seem to be correct in claiming that we do lack evidence against radical skeptical hypotheses,” Cohen acknowledges. “However, it does not follow that it is not reasonable or rational to deny such hypotheses.”³¹ According to Cohen, it can be rational to hold a belief despite not having any evidence for such a belief.³²

Cohen noted, however, that it is *prima facie* odd for a subject to know particular details of a world without experience of that world: he recognized that non-evidential rationality leads to a-priori contingent knowledge. To correct this, Cohen found it necessary to conjoin his claim that it is rational to deny a global scenario *a la* non-evidential rationality with empirical evidence of that world: the *a priori* (non-evidential rationality) commingled with the contingent (the *empirical evidence*).³³

And so Cohen's RA contextualism falls into place: through a variety of considerations we can secure knowledge claims in an ordinary context while explaining how skeptical arguments destroy knowledge, albeit temporarily. Cohen's analysis involves examining the type of skeptical argument the argument is—restricted skepticism, global skepticism, or both—and showing how the skeptic makes salient an otherwise irrelevant alternative, lifting us out of an ordinary context and into a skeptical one. With closure, the normally irrelevant alternative

³⁰ “Thus in everyday contexts, I can know that I don't see a cleverly-disguised mule, on the basis of the inductive evidence I have against such a scenario.” in Cohen, “Contextualism, Skepticism, and the Structure of Reasons,” 66.

³¹ Cohen, “How to be a Fallibilist,” 112.

³² Stewart Cohen, “Contextualist Solutions to Epistemological Prolems: Skepticism, Gettier, and the Lottery,” *Australasian Journal of Philosophy* 76 (1998): 68. “Maybe we are unable to demonstrate to a skeptic that our beliefs are rational. But that does not mean that we cannot satisfy ourselves that they are. If it seems right to say that it is to some degree rational...then we can appeal to that fact in our attempt to resolve the paradox.” (Cohen, “Contextualist Solutionsto Epistemological Prolems: Skepticism, Gettier, and the Lottery,” 69).

³³ Cohen, “Contextualist Solutions to Epistemological Prolems: Skepticism, Gettier, and the Lottery,” 76.

destroys knowledge. When ordinary knowledge is threatened in a skeptical context, by virtue of the extraordinary possibility, the ordinary knowledge is lost; otherwise, what we take ourselves to know is safe from challenge, as, depending on the type of skeptical threat, we can appeal either to inductive evidence or non-inherent rationality in conjunction with empirical evidence. Understanding Cohen's contextualist solution to the skeptical problems of the painted mules and EWS provides a template to be deployed against AI-type skeptical arguments that may come our way.

The Problem of Induction

The problem of induction is a particular type of skeptical problem concerned with the justification of inductive inferences. Many philosophers have tried to solve the problem of induction, but no solution has gained wide acceptance. As we've seen, contextualists claim to answer another skeptical problem, the problem of the external world. This suggests that one can devise a contextualist solution to the problem of induction.

A person makes an inductive inference when she expects a future event to occur based on past events of the same sort having occurred under sufficiently similar conditions—such as presuming an A-type event will be paired with a B-type event after witnessing a history of such pairings. We make this type of inference all the time: thinking the sun will rise in the morning; relying on gravity to pull a ball tossed up back down again; expecting bread to provide nourishment.³⁴ However, David Hume raised the question of whether one can provide a rational basis for a belief in induction, despite the apparent widespread success of these inferences. Hume provides an argument against the possibility of justifying induction; the effects of this argument are enough to merit the “problem of induction” to become known by some as “Hume’s problem.”³⁵ In this section my main aim is to set out the problem of induction, whether we can be justified in extrapolating general inductive conclusions from corresponding inductive premises. To this end I will explain Hume’s argument against induction. In the following section I will consider the fallout if Hume’s argument is indeed correct.

Justification Requirement and Inductive Inference

The traditional three-part account of knowledge requires a knower to have a justified belief that is true. A subject S knows some proposition p just in case S believes that p , S is justified in believing that p , and p is actually true.³⁶ I will understand justification in regard to induction to require having satisfactory reasons to hold a belief acquired through inductive inference. As an instance of such inductive reasoning, we might believe that the sun will rise tomorrow because it has reliably done so in the past, and this belief very well might be true—let’s say for argument the sun *is* going to rise tomorrow. But this leaves the *justification* of this true belief in question: on what grounds are we justified in going beyond premises (past events) to some conclusion (a future event, something that is not yet an experience)? Laurence Bonjour recommends viewing the collected instances of observed inductive evidence as comprising a “standard inductive premise” and the conclusion we draw solely on the information in this premise as a “standard inductive conclusion.”³⁷ The *problem of induction* is the challenge to justify the *leap* from the standard inductive premise to the standard inductive conclusion: what is the rationale for concluding from all observations of the sun’s rising, the inductive premise, that it will do so yet again, the inductive conclusion? As we will see from Hume’s argument, there appears to be an additional, tacit assumption that begs the question in arguments to justify induction, a *tacit premise* in the space between the inductive premise and inductive conclusion that assumes the very point at issue.

³⁴ This last example is one of Hume’s examples. David Hume, *An Enquiry Concerning Human Understanding* (Oxford: Oxford University Press, 1999), 114.

³⁵ Howson, Colin. *Hume’s Problem: Induction and the Justification of Belief* (New York: Oxford University Press, 2003), 10.

³⁶ Again, I’m here setting aside any Gettier-type concerns.

³⁷ Laurence Bonjour, *Epistemology: classic problems and contemporary responses* (Lanham, Maryland: Rowman & Littlefield Publishers, Inc., 2002), 56.

Hume recognizes two types of knowledge: *a priori* and *a posteriori*.³⁸ Given this, Hume's argument against justifying induction takes the form of a dilemma: any proposed solution to the problem of induction must rest on *a priori* grounds or *a posteriori* grounds. As Hume intends to show induction to be unjustified, Hume *qua* skeptic must show that an argument justifying inductive inferences can not be justified by *a priori* or *a posteriori* means.

Before looking at Hume's argument we will look at what constitutes a case of inductive inference, how we presume the future will resemble the past, in more detail. Consistently paired types of events lead us to infer the continued pairing of these types of events: an A-type event followed by a B-type event leads us to expect the B-type event when we encounter the A-type event. As a result of encountering these event types in tandem, we develop the (generally tacit) belief that if an A-type event occurs then a B-type event will follow: events of this sort will unfold in the future in the same manner as events of this sort have unfolded in the past. When it has rained in the past, the pavement on my street has gotten wet. If looking out the window I see that it is beginning to rain, I inductively infer that the pavement is going to get wet in the same way it has each time before. Hume's argument challenges this common sense expectation. The tacit premise common sense relies upon to go from observed particular events to general inductive conclusions must be either justified on *a priori* or *a posteriori* grounds—if it can be justified at all. We'll look at the trouble Hume takes an *a priori* justification of induction to face.

Hume explains that “The mind can always *conceive* any effect from any cause, and... whatever we *conceive* is possible, at least in a metaphysical sense.”³⁹ For an argument to be justified on *a priori* grounds the conjunction of the premises and the negation of the conclusion must not be metaphysically possible. That is, on Hume's view, an argument is *a priori* justified where if the conclusion was denied it would contradict the premise(s).⁴⁰ In justifying induction, an *a priori* justification runs aground because it is possible that what inductive premises lead us to expect might not occur: that which has been regular quite simply might change and this will not contradict the prior events. A possible event, no matter how unlikely its occurrence might seem, can not be ruled out.

Most contemporary philosophers hold that Hume requires deductive validity for *a priori* justification.⁴¹ By this interpretation, in order to justify inductive inferences on *a priori* grounds one must demonstrate that the inference is logically entailed by the premises. It is important to stress here that an argument attempting to justify induction on *a priori* grounds would thus have to be demonstrably correct by virtue of its premises alone; it would have to be shown that a belief in induction is rational by an appeal to deductive proof alone.⁴² We can look at the rain scenario to see how this works.

³⁸ Although Hume does use the term *a priori* in the *Treatise* and the *Enquiry*, the designations *a priori* and ‘*a posteriori*’ come after Hume. Hume referred to “relations of ideas” and “matters of fact”: “All reasonings may be divided into two kinds.” (Hume, “Enquiry,” 115). In Hume's terminology, this is *demonstrative* (*a priori* deductive reasoning) or *probable* (*a posteriori*—reasoning by reference to experience). For clarity I use the current terminology, ‘*a priori*’ and ‘*a posteriori*.’

Charles Landesman presents Hume's division as such: on *a priori*, “Hume claims that such truths do not bring with them any ontological commitments; they imply nothing about any existing objects; the mind operating independently of experience, has no way of gaining access to any objects at all.” Charles Landesman, *Skepticism: the central issues* (Malden, MA: Blackwell Publishers, 2002), 143.

Landesman presents the *a posteriori* as facts that can be verified “by experience and that imply ontological commitments. These are neither intuitive nor demonstrable because their existence is contingent.” Landesman, 143.

³⁹ David Hume, *A Treatise of Human Nature* (Oxford: Oxford University Press, 2000), 410.

⁴⁰ Bonjour, 59.

⁴¹ “According to Hume, any true proposition about the relations among our ideas is provable by deduction, because its negation will imply a contradiction.” Landesman, 143.

It would take an extraordinary turn of events for it to rain in my neighborhood and the pavement not to get wet. But we can easily conceive of such a circumstance. For example, it's possible—if not probable—that rain could fall from the sky and slip into another dimension just as it was about to reach the street. This scenario shows that it's possible for it to rain in my neighborhood without the pavement on my street getting wet. Despite the seemingly good reasons we have for expecting a B-type event to follow an A-type, it's metaphysically possible that we could have the one, the A-type event, without the other, the B-type event, following. The pavement *could* remain dry on a rainy day.

So, as we can conceive that things can change—rain slipping into another dimension—intuitively we can recognize that an appeal to past events to justify belief in a future event can not provide *a priori* justification for induction. As Hume puts it, “that there are no demonstrative arguments in the case seems evident; since it implies no contradiction that the course of nature may change.”⁴³ Inductive inference is therefore not justified on *a priori* grounds and the tacit premise therefore doesn't consist of *a priori* justification. This leaves *a posteriori* justification to be considered. For if it is not through reason alone, Hume tells us, then rather it must be through experience that we can justify inductive inferences—that is, if they are justified at all.

In attempting to justify our belief in the correctness of inductive inferences, then, the other tact to take is to argue from experience, on *a posteriori* grounds. The common sense response to the problem of induction is to point to the innumerable instances where belief in induction seems to have paid off: the sun actually rising, objects tossed up falling back down again, the cases of bread providing nourishment, and so on. A more sophisticated version of this common sense approach is to argue that we can 1) review a sufficient number of past cases where a variety of conclusions based on corresponding inductive premises were found true; 2) justifiably infer from 1 that these many particular cases of induction can justify a belief in a general, law-like inductive thesis; and 3) use this general inductive thesis to justify future particular cases where we expect B-type cases to follow A-type cases.⁴⁴

This *a posteriori* approach doesn't get one as far as it might first seem to suggest. There are two problems with the approach: an attempt to verify step 1 will never succeed as one could never be in an epistemic position to observe all of the future, past, present, and possible elements in the various inductive conclusions. To take one example, one could never verify all (*future, past, present, and possible*) instances where a medium-sized body was affected by gravity.⁴⁵ The second issue facing an *a posteriori* justification for induction is that to justify inferring a general inductive thesis based on an appeal to experience (in steps 1 and 2) *is itself an example of inductive inference*: in order to justify making an inductive inference we are inferring inductively, begging the very question at issue. In short, the reason given to justify a particular inductive inference is itself inductive. The tacit premise we identified earlier is playing the question-begging role—we are using the tacit premise in the argument to justify the tacit premise. But this is strictly forbidden. In Hume's words, “our experience in the past can be a proof of nothing for the future but upon a supposition, that there is a resemblance betwixt them.”⁴⁶ But it is this supposition, the tacit premise, for which we require justification.

⁴² But, as Howson explains, “Hume commences by pointing out that no inference from the observed to the as-yet unobserved is deductive.” In Howson, 11.

⁴³ Hume, *Enquiry*, 115.

⁴⁴ Bonjour, 60.

⁴⁵ Hume does not mention this first problem: it is taken from Bonjour, 60.

⁴⁶ Hume, *Treatise*, 410.

Examining the rainy day case, we get a better understanding of this second difficulty faced by an *a posteriori* justification of induction. We come to expect the pavement to get wet in a rain shower because, based on experience, that's precisely what's happened when it's rained in the past. We can point to many cases for support. What Hume asks for here is to identify where exactly we're getting the justification to make such an inference; where in the range of cases of B-type events following A-type events do we find the means to project into the future that a yet-to-occur B-type event will follow this A-type event? In the case of the rainy day, for all observed instances when it's rained the pavement has gotten wet and from this it is inferred that it shall do so likewise in the future. In an attempt to justify this inference, we must reach farther than the particular instances where it has rained on my block. Eventually we are pressed to justify the general thesis. The move here is to invoke other instances of the sort that support a general belief in inductive inference, instances like bread nourishing or the sun rising. The thought is that one can justifiably rely on this particular conclusion because it fits within a larger schema of inductive inferences that seem to work quite well. The argument becomes: I can justifiably infer that the rain shower in my neighborhood will wet the pavement on my street because 1) it has done so every time I've observed it rain in my neighborhood and this is an instance where many particulars of a specific sort can justify a corresponding, inductive conclusion; 2) inductive inferences like this are justified in general because we can point to the numerous A-type followed by B-type cases that resemble the rainy day scenario (the sun rising, gravity's effect on medium-sized bodies, bread nourishing...) for support; and 3) one can justifiably infer from this general inductive thesis that a future event of a sort will resemble previous cases of that sort, in this case that the rain shower in my neighborhood will wet the pavement on my street because it has done so every time I've observed it rain in my neighborhood.

We can see that the *a posteriori* argument to justify an inductive inference presumes the very tacit premise it's trying to demonstrate. One way to frame this issue is in terms of evidence: the approach begs the question about inductive inference based on evidence by pointing to evidence in order to provide grounds for inductive inference. This approach—and Hume tells us any *a posteriori* approach—uses induction to justify induction, which is circular. Anyone who attempts to argue in this way, he says, “must be evidently going in a circle, and taking that for granted, which is the very point in question.”⁴⁷ For if faced only with limited knowledge, any conclusion about future events is one gained by the following presupposition: instances of past events grant the ability to anticipate the recurrence of such an event.

We can thus formulate in a neater form the two prongs of Hume's argument:

- 1) An *a priori* attempt to justify induction requires that a contradiction results from maintaining premises and negating the conclusion. But there is no contradiction in asserting that the future might change.
- 2) An *a posteriori* attempt to justify induction presupposes that the future will be like the past, making the very inference that it attempts to justify, thus begging the question.

On the first, it's not contradictory to suppose a belief in induction is false. On the second, making an appeal to any evidence of past inductive success presupposes that we can have prescience based on experience, the question at issue. Hume therefore seems well-equipped to deny attempts to justify induction through either *a priori* or *a posteriori* means.

⁴⁷ Hume, *Enquiry*, 115.

Consequences: Scientific and Personal

Hume wrote the *Treatise* with the expressed intent of becoming the Isaac Newton of the mind.⁴⁸ There's some irony here. For, in his project Hume appears to undermine the very basis of scientific success. As we've seen, Hume convincingly argues against any justification of induction. Thus, Hume's work on "inference," "probable arguments", and "reason from experience"⁴⁹ seems with its evident paradoxes to carry a serious task for science to reformulate its theoretical basis. Moreover, his argument reaches beyond the laboratory: it is fundamental to human life that conclusions about the world can be drawn and applied over a range of space and time; his argument threatens even our most basic beliefs: anytime we rely on a regular feature of the world we do so based on an inductive inference, whether the inference is articulated or not. We'll look first at the scientific consequences and then at the mundane, personal consequences.

If we were to take a complex system of physical events, such as the astronomical activity involved in the sun's apparent rising, and whittle it down to more basic elements, it would appear that we had found at root basic tendencies of the physical objects in the system. And (quite naturally it seems) we would have a most solid expectation of the continued existence of these causal laws—both particularly and universally. This presumed knowledge is implicit in our daily lives, where we work into the quotient of our physical movements the steadfast presence of gravity, for example. But it is an explicit component of science.

There are two basic processes in scientific information gathering: the summative and ampliative.⁵⁰ The former involves a claim about a group as a whole based on the inspection of its component members. A sample of emeralds might be collected and properties such as the color, hardness, density, and chemical composition measured with the greatest precision. Conclusions such as that emeralds are green with a density of such-and-such are made, this information is compiled, and a summary record is then available for other scientists to consult.

An instance of an ampliative process might consist of utilizing some of this stored information for its predictive or explanatory power. An ornithologist might evaluate the pigment found in the wings of ravens and, in addition to drawing conclusions as to the color properties of raven feathers, the scientist then might further amplify this basic information to include cases in the potential future or the most remote past. The qualities analyzed in the pigment might reveal to the scientist how ravens have fared in different climates in ages past or perhaps the composition of the pigment bears a certain reactive quality that can be realized for an industrial use: non-fading dyes, a catalyst in alumina reactions, or a general product that will find its way into everyone's home. We see in this way how basic information can be extracted and applied to unlimited cases future and past, well beyond the particular experience in which the information was gathered.

It is with this ampliative induction that science has known its greatest triumphs. Abstracting generalizations from particular instances and applying this knowledge to other, discrete, particular instances not only keeps humans upright and their thirsts quenched, but the careful measurements and thorough evaluation of data has also led to predictions of lunar eclipses and expeditions to the very surface of the moon. And with the advent of quantum mechanics, predictive power of the utmost precision has been gained.

⁴⁸ S.J Frederick Copleston, *A History of Philosophy Volume V* (New York: Doubleday, 1994), 261.

⁴⁹ Hume never used the word 'induction' as we're using it here (although Bacon and Locke did). Livingston, D.W.; J.T. King, *Hume : A Re-Evaluation* (New York: Fordham University Press, 1976), 69-70.

⁵⁰ Jonathan L. Cohen, *An Introduction to the Philosophy of Induction and Probability* (Oxford: Clarendon Press, 1989), 9-10.

But, as we saw above in Hume's analysis, what we have grown to expect may one day cease or change: despite what we might think likely, it is theoretically possible that a fair coin might consistently land heads.⁵¹ We have no rationale for ruling out the possibility of the otherwise, no matter how many events witnessed or heard about, no matter how many specimens collected and examined.

This strikes an awkward chord. This is *just* how we know, comes the response. We test and test, and then we conclude. This aspect of the inductive inference is known as *enumerative induction*: the tendency to extrapolate from many instances of particulars a general conclusion, with the greater amount supposedly supporting more strongly this inference, the larger number increasing the probability of whatever is in question. This has been the basis of empirical science. But, Hume appears to have revealed, that there is a lingering problem behind the seeming successes of science and indeed our daily habits. If Hume is correct, then there is no rational justification for what we ordinarily take to be given. As Bonjour starkly explains: "if inductive inference is unjustified, so all apparently are all such scientific beliefs [explanations for patterns or regularities]."⁵² By this light, a fundamental part of our lives is without any rational basis.

Science depends on the continued success of its research and analysis, explanations and predictions. Reliability and predictability are its hallmarks. Causal laws and past evidence are relied upon throughout technology, the implementation of scientific discovery. This of course impacts our personal lives through, among other things, manufactured equipment, the food we consume, and the energy we use. As we've seen, though, there are more personal, less refined ways that we use induction. Merely the participation in our daily lives without suddenly reeling into chaos depends on and instantiates our faith in these inferences. Our nourishment, our basic safety, our ability to walk across the room or use language: these all clearly depend upon inductive inferences. Any evidence about the past is rendered void and ineffective. Empirical evidence as a concept becomes meaningless: the evidence cannot be used as evidence, nothing can be taken as a reliable indicator of anything. As Bonjour asks, if Hume is correct and induction is unjustified,

how could I be justified in thinking that any of these sorts of alleged evidence are genuinely reliable indicators of the sorts of events they are alleged to be evidence for without relying in some way on inductively established generalizations pertaining to the relations between such events and the production of the corresponding evidence...?⁵³

Indeed on this account we don't seem to be justified in making knowledge claims about the external world. Many argue also that we cannot even assert with full confidence that an external world exists, whatever properties we may want to assign to it.⁵⁴ Without the knowledge that our inductive inferences are correct, we seem capsized, left with little but doubt, paranoia, and

⁵¹ This possible world can be made our world. According to Daniel Dennett, with 1,024 participants in a ten-round elimination coin-toss tournament, one person is bound to successively win ten successive coin tosses. Using this method, there can presumably be made-to-order a winner of any desired length of random coin-tosses Daniel C. Dennett, *Darwin's Dangerous Idea: evolution and the meaning of life* (New York: Touchstone, 1995), 54.

⁵² "Thus if inductive inference is unjustified, so all apparently are all such scientific beliefs [explanations for patterns or regularities]" Bonjour, 62.

⁵³ Bonjour, 62.

⁵⁴ Bonjour, 63.

solipsism. As Russell succinctly put it, “If this were not known to us, we could never extend our knowledge beyond the sphere of our private experience; and this sphere, as we have seen, is exceedingly limited.”⁵⁵

Faced with such a situation, it is clear that there is great motivation to assert the contrary, great intuitive compulsion to say we *can* make such inductive inferences. The philosophical skeptic, on the other hand, is appropriately skeptical of any success in this matter.

⁵⁵ Bertrand Russell, *The Problems of Philosophy* (New York: Oxford University Press), 60.

A Contextualist Solution to the Problem of Induction

The problem of induction is another argument from ignorance; as such, we have very good reason to expect that a contextualist solution to the problem of induction can be constructed along the same lines as our solutions to the AI arguments based on the BIV and painted mules scenarios. In this section, I develop a contextualist solution to the problem of induction and assess the success of that solution.

We've seen that the problem of induction is a skeptical problem, but moreover it is also an argument from ignorance.⁵⁶ Recall that the central feature of these arguments is that the skeptic presents us with an alternative hypothesis—some other explanation to the observed phenomena—that we are unable to rule out. In the case of BIV, we can't rule out the possibility that we are a BIV and, because of open possibilities for explanation, are therefore unable to grant knowledge of knowing we have hands.

In the problem of induction, we can't rule out there is some other explanation aside from induction—for example, all of the past correlated events could be the result of *random chance* (RC, where RC indicates random chance has until now provided for the appearance of a regular predictive feature(s) of the world). A consequence of Hume's argument is that if this alternative hypothesis RC cannot be ruled out—and as the course of nature could possibly change, it cannot be—then it follows that one cannot justifiably make an inductive inference.

Let's put Hume's problem in the AI schematic form we've already developed. We already have RC as an alternative to induction (H in the schematic form). We'll also need a restricted example of what we ordinarily take to be the case—O in the schematic form. An O is a particular aspect that follows from the generally accepted thesis (in the BIV scenario, the O is having hands, which we would know if we were to know that we were not a BIV). Tossing a tennis ball up and expecting it to fall down again (TB) is a specific, restricted instance of inductive inference and will serve our purposes.

P1. I do not know that not-RC.

P2. If I do not know that not-RC, then I do not know that TB.

C. I do not know that TB.

With the problem of induction in its AI schematic form, we can begin to see how a contextualist solution suggests itself. We can look back at how Cohen tackled EWS to see how a contextualist solution to the problem of induction should proceed.

Recall the template we developed in our analysis of Cohen's solution to EWS. A Cohen-based contextualist solution to the skeptical argument against induction proceeds in the same way: the solution evaluates the type of skeptical argument it is—restricted skepticism, global skepticism, or both—and establishes the (ordinary and skeptical) contexts in which the argument operates, identifying how the skeptic makes salient an otherwise irrelevant alternative and through the operation of the closure principle makes the conclusion seem irrefutable. Once this analysis gives us the component pieces, we can follow through with the contextualist explanation of the skeptical argument and determine how our ordinary knowledge is safe from the challenge.

⁵⁶Hume's problem is an ideal candidate, for it is, as Bonjour has commented, "an exceptionally clear example of the general form epistemological problems typically take." Bonjour, 56.

The AI argument in the problem of induction involves both a restricted form of skepticism—a challenge to the expectation that the tennis ball tossed up will fall down again—and a global form of skepticism—that an induction-undermining alternative is the case and thus inductive inferences are made without justification. This yields a universally threatening challenge to inductive-based knowledge, one that can extend to other restricted particulars: the example of the tennis ball can be replaced with another process or event where we regularly apply inductive inference. Note, also, that the presence of both global and restricted forms conveniently fits with Cohen’s analytical tools—we can deploy his analysis in full force, tackling both the restricted and global aspects of this problem.

In a restricted form of skepticism, Cohen tells us that when in an ordinary context we can appeal to inductive evidence gathered from past events to deny the particular, restricted skeptical threat: a subject can justifiably assert the zebra-like animals to be zebra based on experience of zoos, animals, and other related things. By the same token, in the case of the problem of induction, we can rely on experience playing the game of tennis to justifiably expect a tennis ball tossed up to come back down. If the ball has come back down every time I’ve tossed it for a serve this afternoon (and there is no skeptical argument present), then I can claim “when I toss this ball up it will come back down” with full justification.

For a restricted scenario in a *skeptical* context, however, that same justification is lacking. Earlier we saw that in the painted mules scenario, a subject in a skeptical context could *not* justifiably claim to know the zebra-like animals to be zebra. The skeptic has made salient the possibility of an otherwise improbable alternative, and through the closure principle this undermines our knowledge of the particular case in question.

Recall that on Cohen’s view it is inherently rational to deny the purchase of a global skeptical argument when in an ordinary context; this is particularly emboldened (and, as we saw, serves to allow Cohen to avoid *a priori* contingent knowledge) when we combine this with empirical evidence. We saw in the global form of EWS—the argument from BIV—that Cohen holds it to be inherently rational to deny BIV (combined with empirical evidence of hands to account for contingent knowledge). It is likewise in the case of the problem of induction. On Cohen’s line of argument, it is inherently rational in an ordinary context to insist that a subject is justified in presuming the world to consist of measurable, predictable processes and events (and, in light of Cohen’s concern, couple this with evidence of having seen the process in effect—TB, for example). It is irrational on Cohen’s line of thinking to disbelieve the validity of inductive inference, particularly when one can point to such instances as TB. In this way, Cohen preserves justification for relying globally on inductive inference in an ordinary context.

In a skeptical context, however, the force of the global skeptical argument is not tempered—knowledge is undermined: a subject cannot justifiably rely on inductive inferences. Just as BIV globally destroys knowledge about an external world, so does an argument based on RC deny a subject’s ability to knowledgeably predict an occurrence of an event based upon experience.

According to our analysis, then, the skeptic makes the normally irrelevant alternative RC relevant by making it salient, and thus, through the closure principle, knowledge is destroyed, as evidenced by our loss of TB. As Cohen would have it, though, this does not destroy our knowledge of induction as a reliable process simpliciter. It is only in the skeptical context that we are not justified in relying on inductive inference; in a more ordinary context, we can justifiably rely upon these processes we take for granted everyday.

Our contextualist solution to the problem of induction does just what we expected a solution to do: it preserves the knowledge claims under consideration while valuing our intuitions about knowledge—it seems correct to both deny and assert justified, true belief depending on the context. The apparent threat is defused, our knowledge is secure, and the closure principle is upheld. A successful contextualist solution to Hume’s problem evaluates its efficacy in terms of contexts, establishes the plausibility of the propositions and the shifting contexts involved in the arguments success, and this seems precisely what our contextualist solution to the problem of induction does.⁵⁷ But there is a problem—two problems: namely, the horns of Hume’s problem.

The Two Horns

Recall that the dilemma of Hume’s problem is that you must solve the problem of induction by either *a priori* or *a posteriori* means, but that neither are viable routes to justifying induction. Notice that this means that you cannot rely on *a priori* rationality to justify induction (for as we’ve seen, there is no contradiction involved with a change in the course of nature) and, on the other hand, one of course cannot appeal to empirical evidence either (as this relies upon inductive inference). Cohen’s contextualism, when deployed to solve the problem of induction, clearly violates these restrictions: it explicitly appeals to induction in the restricted scenario and explicitly appeals to rationality in the global. The component of Cohen’s contextualism that establishes knowledge in the face of a global scenario—that it is inherently rational to deny the skeptical hypothesis—is an *a priori* argument. The component of Cohen’s argument that establishes knowledge in the face of a restricted skeptical scenario—that we can appeal to inductive evidence—is an *a posteriori* argument.

These are precisely the two approaches Hume has shown to be deficient in attempting to refute his problem. Additionally, recall that the first component bothered Cohen with its result of *a priori* contingent knowledge. To remedy this, he asserted that one must combine inherent rationality with empirical evidence. As we’ve established in our treatment of induction, however, empirical evidence entails inferring inductively. With this addition, the global scenario alone—in the case of the problem of induction, a hypothesis that undermines justification for inductive inference—now violates both horns of Hume’s argument. So, even when Cohen bolsters his non-evidential rationality with inductive evidence, he is still begging the question.

These background assumptions might be acceptable with EWS, but in the case of inductive skepticism, they undermine the conclusion. Cohen’s contextualism fails as a solution to the problem of induction: when his contextualism is put against the problem, you get Hume’s problem all over again.

One who sympathizes with contextualism might argue that the contextualist analysis is designed only for EWS and examples like the painted mules and papier-mâché barns. But this doesn’t fly, for as I have shown, the problem of induction is of the same AI type as the other arguments. Even if we were to grant this *ad hoc* exemption, however, there is still a lingering problem: through the problem of induction alone, we can generate EWS. This suggests that, at minimum, these two issues are conceptually entangled, which would require from a contextualist an explanation for why we should consider them to be separate, and, at maximum, it suggests

⁵⁷ The first premise creates a skeptical context. The second premise can occur in a skeptical or ordinary context, since for the contextualist this instance of the closure principle states a logical truth that follows independent of context. The conclusion is valid only in a skeptical context. When the skeptical argument is not present, the skeptical context is not maintained, and the conclusion has no purchase.

that, at least for a relevant alternatives contextualism like Cohen's, none of the contextualist solutions for AI-type arguments is successful without a full and proper treatment of the problem of induction.

Conclusion

In this paper I have argued that a contextualist solution to the problem of induction built on the work of Stewart Cohen is untenable. The “External World Skepticism” section explained the skeptical problem of the external world, showing that the skeptic denies we have the justification we normally presume to have for knowledge claims about an external world. The following section explained Stewart Cohen’s contextualist argument against EWS, the upshot of which is that depending upon the context, the standards for justification for knowledge claims differ and therefore, although knowledge of a proposition *is* lost in a skeptical context, knowledge of the same proposition is secure in an ordinary context. The following section outlined the problem of induction and then in “A Contextualist Solution to the Problem of Induction” we saw how utilizing our earlier analysis of Cohen’s contextualist solution to EWS we were able to construct a similar solution to the problem of induction. In the end, though, we saw that there are particular problems with Cohen’s contextualism that invalidate an attempted solution to the problem of induction. If my arguments are correct, this further suggests that Cohen’s contextualism—and any theory similar to Cohen’s—must account for the problem of induction in a contextualist solution to EWS. This perhaps too suggests that contextualism in general may contain some significant deficiencies—even beyond its explicit question-begging.

Contextualism has many champions and many detractors. Some dismiss its concerns as epistemologically irrelevant. Others focus on its practice of openly begging the question, questioning its philosophical relevance. If Cohen’s contextualism could have been successfully deployed against the problem of induction, then this would have bolstered the contextualist thesis in general. However, as we’ve seen the presumptions in Cohen’s theory prevent it from successfully addressing the problem of induction, we should weigh whether or not this adds to the detractors column in the debate on contextualism’s relevance. Cohen’s failure here should not be taken to damn contextualism outright, but it certainly recommends examining other forms of contextualism in light of Cohen’s and strongly speaks against an RA-based form of contextualism.

I’ve suggested ways in which the contextualist might respond to my criticism, but I’ll also note here that one can see a built-in safety feature in contextualism: it seems a contextualist will attempt to avail herself of a contextualist defense from critics of contextualism. In the context of skepticism about contextualism, a contextualist might nod assent—the criticism is accurate in that context. But allow the lax standards to resume and the contextualist can say by these everyday standards contextualism just seems right. In the specific case of justification for ordinary knowledge claims, the contextualist will try to say that only by investigating these claims in light of epistemology (the practice of which David Lewis suggests undermines knowledge) does this requirement for justification pose a problem. In an everyday context, where we aren’t skeptical about everyday knowledge claims, and truth is gauged by the way we use words, the contextualist will say that our claims are safe. The contextualist seems to think herself in quite an untouchable position (despite our skepticism about such a position). But maybe we should simply ask: Is a philosophical position that is immune to critical debate a satisfactory position?

References

- Black, Tim. "Contextualism in Epistemology," *The Internet Encyclopedia of Philosophy*. <http://www.iep.utm.edu/c/contextu.htm> (Accessed September 2004).
- Bonjour, Laurence. *Epistemology: classic problems and contemporary responses* (Lanham, Maryland: Rowman & Littlefield Publishers, Inc., 2002).
- Cohen, L. Jonathan. *An Introduction to the Philosophy of Induction and Probability* (Clarendon Press, Oxford: 1989).
- Cohen, Stewart. "Knowledge and Context," *Journal of Philosophy* 83 (1986): 582.
- Cohen, Stewart. "How to be a Fallibilist," *Philosophical Perspectives* 2 (1988): 113.
- Cohen, Stewart. "Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery," *Australasian Journal of Philosophy* 76 (1998):
- Cohen, Stewart. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13 (1999).
- Copleston, S.J Frederick. *A History of Philosophy Volume V* (New York: Doubleday, 1994).
- Dennett, Daniel C. *Darwin's Dangerous Idea: evolution and the meaning of life* (New York: Touchstone, 1995).
- Descartes, Rene. *Selected Philosophical Writings* (New York: Cambridge University Press, 1998).
- Dretske, Fred. "Epistemic Operators" in *Skepticism: a contemporary reader*, ed. DeRose, Keith and Warfield, Ted A. (New York: Oxford, 1999).
- Howson, Colin. *Hume's Problem: Induction and the Justification of Belief* (New York: Oxford University Press, 2003).
- Hume, David. *A Treatise of Human Nature* (Oxford: Oxford University Press, 2000).
- Hume, David. *An Enquiry Concerning Human Understanding* (Oxford: Oxford University Press, 1999).
- Kornblith, Hilary. "The Contextualist Evasion of Epistemology," *Philosophical Issues (Nous)*, 10 (2000).
- Landesman, Charles. *Skepticism: the central issues* (Malden, MA: Blackwell Publishers, 2002).
- Putnam, Hilary. *Reason, Truth, and History* (Cambridge: Cambridge University Press, 1981).

Livingston, D.W.; King, J.T. *Hume : A Re-Evaluation* (New York: Fordham University Press, 1976).

Pryor, James. "Highlights of Recent Epistemology," *British Journal for the Philosophy of Science* 52 (2001).

Russell, Bertrand. *The Problems of Philosophy* (New York: Oxford University Press, 1965).

Schiffer, Stephen. "Contextualist Solutions to Skepticism," *Proceedings of the Aristotelian Society* 96 (1996): 317-333.

Sosa, Ernest. "Knowledge in Context, Skepticism in Doubt," *Philosophical Perspectives* 2 (1988): 139-155.

Stanley, Jason. "Context and Logical Form," *Linguistics and Philosophy* 23 (2000): 391-434.

Vita

William Galloway Osborne III is an artist, writer, musician, teacher, and thinker who recognizes the pretentiousness inherent in these types of designations. He has a lovely, smart wife named Annie and a handsome grey cat named Isaac (after the Woody Allen character in *Manhattan*). He has earned a terminal degree in painting; taught at the primary, secondary, and university level in a variety of disciplines; and cooks in the Thai, Indian, Latin, and Continental traditions. He aspires to live in a functional community with an agreeable climate along the lines of San Francisco, New York, Boston, Paris, or Geneva. Meanwhile he attempts to temper ambition with gratitude and to practice equanimity in the face of ephemera.