

When a Skeptical Hypothesis is Live

Bryan Frances
Fordham University

Revised version forthcoming in *Noûs*

I'm going to argue for a set of restricted skeptical results: roughly put, we don't know that fire engines are red, we don't know that we sometimes have pains in our lower backs, we don't know that John Rawls was kind, and we don't even know that we *believe* any of those truths. However, people unfamiliar with philosophy and cognitive science do know all those things.

The skeptical argument is traditional in form: here's a skeptical hypothesis; you can't epistemically neutralize it, you have to be able to neutralize it to know P; so you don't know P. But the skeptical hypotheses I plug into it are "real, live" scientific-philosophical hypotheses often thought to be actually true, unlike any of the outrageous traditional skeptical hypotheses (e.g., 'You're a brain in a vat'). So I call the resulting skepticism *Live Skepticism*. Notably, the Live Skeptic's argument goes through even if we adopt the clever anti-skeptical fixes thought up in recent years such as reliabilism, relevant alternatives theory, contextualism, and the rejection of epistemic closure. Furthermore, the scope of Live Skepticism is bizarre: although we don't know the simple facts noted above, many of us do know that there are black holes and other amazing facts.

1. The Dinosaur Argument

When Jo was a teenager she learned that a huge meteorite wiped out the dinosaurs. She learned this theory in the usual way, hearing it from her parents, teachers, and books. Now pretend that at the time she was told the meteor story as a child, say at the age of eight, the scientific community was sharply divided on the issue of what caused the demise of the dinosaurs. Although most scientists accepted the meteor hypothesis, many others subscribed to the idea that their death was caused by some enormous solar flare. A significant number of other scientists thought that it wasn't a solar flare or a meteor but a particularly nasty series of supervolcanos. These latter two classes of dissenters had decent evidence: evidence concerning the sun and supervolcanos that the meteor advocates took seriously. Both the solar flare theorists and the supervolcano theorists were highly respected professors, highly respected by the meteor theorists and at the top of their profession. Whole book series, conferences, and PhD dissertations were devoted to these competing hypotheses. Suppose

further that upon going to her university Jo found out about the rival and highly respected hypotheses. She didn't understand all the reasons why they were so well respected and endorsed, but she was well aware that they were well respected and frequently endorsed by the experts, even the best among them. Even so, she kept her meteor belief.

As it turned out, the meteor hypothesis was the right one. But although some experts may have known that fact, surely Jo was in no position to know it after she found out about the eminent status of the rival hypotheses. She could do little or nothing to defeat the rival hypotheses. Even after studying the issues as an undergraduate she couldn't *know* that a meteor wiped out the dinosaurs, for from the moment she first learned of the dinosaurs until she graduated with a Bachelor's degree in paleontology she was perfectly aware that there were two "live" contrary hypotheses that she was in no position to rule out—and they had not been ruled out for her, either, for instance by her teachers. And if she was aware of all this and the competing hypotheses weren't ruled out, then she couldn't know that the meteor hypothesis was correct. Perhaps in order to know the meteor truth she wouldn't have needed to rule out the possibility that she is a brain in a vat. And perhaps someone mostly outside the scientific community and thus not exposed to the *ultimately* misleading supervolcano and solar flare evidence could know the meteor hypothesis is true. Still, in order for Jo to know the meteor fact she *does* have to be able to rule out, to some degree anyway, the solar flare and supervolcano possibilities, for, unlike the crazy philosophical hypotheses, these are "real, live possibilities" *and* she is perfectly aware of their existence, live status, and inconsistency with her belief.

2. The Live Skeptic's Argument

The heart of the dinosaur argument is simple: because the supervolcano and solar flare hypotheses are real, live possibilities inconsistent with the meteor hypothesis, Jo is *aware* of all that, she is nothing even approaching an expert or genius on these matters, and those hypotheses are not ruled out, she doesn't know that the meteor story is true. The Live Skeptic uses the same argument form: if you believe P, hypothesis H inconsistent with P is a live hypothesis, and you are like Jo in being a well-informed "mere mortal" with regard to H, then you don't know P. Let's make the argument moderately precise.

We start out with a sufficient condition for a hypothesis H to be a "real, live" socio-epistemic possibility in our intellectual community.

1. In our intellectual community hypothesis H has been through a significant (not to say exhaustive) evaluation by experts over many years.
2. It is judged actually true or about as likely as any relevant possibility by a significant number of well informed, well respected, and highly intelligent experts in the field(s)

H belongs to. Also assume that there are lots of these experts and they are not crackpots.

3. Those experts reached that favorable opinion based on H's merits in a familiar, epistemically responsible way (that is, the way they reached their opinion seems as epistemically responsible as any).
4. Those experts consider there to be several decent and independent sources of evidence for H. So it is not the case that the only reason people pay serious attention to H is the presence of one weird experiment or line of reasoning.
5. Many of those experts consider H to be a "real, live possibility" (i.e., this is what they would say if you asked them).

The second task is to articulate a set of sufficient conditions for being a *well-informed mere mortal* with respect to a hypothesis.

6. You know that the hypothesis H is inconsistent with P (where you believe P).
7. You are at least somewhat familiar with H and the issues surrounding H, including the fact that H is live in the sense of (1)-(5). In fact, you are as aware as just about anyone of the fact that H is live.
8. Your intelligence, understanding, and knowledge are not extraordinary for people in your intellectual community with regard to H or the issues surrounding H. (So although you *may* be an expert, you are certainly no better than the other experts—in particular, the ones who insist that H is or could very well be true.)
9. If you have any reasons or evidence you can marshal for casting doubt on H, and if they were carefully considered by the members of that community—in particular, by the well informed, well respected, and highly intelligent experts who were thoroughly familiar with the hypothesis—they would be nearly universally and confidently rejected as clearly insufficient to rule out the hypothesis (although they may have other merits).

Roughly put, 'S is as aware as just about anyone that P' in (7) means that P is true and S has all the usual good warrant for believing P, warrant that usually suffices for knowledge. I can't say that it entails *knowledge* of P, however, for reasons that will become clear below. Briefly, the reason is that even though S has the usual "positive" warrant sufficient for knowledge of P, she also has some "negative" warrant that offsets or vetoes that positive warrant enough so that her belief in P no longer is, on balance, warranted or sufficient for knowledge.

The conjunction of (6) and (7) says, crudely put, 'You're fully aware of what is going on with the live hypothesis, so it is a real threat to your knowledge, a threat that must be defused.' The

conjunction of (8) and (9) says, again crudely, ‘You do not have what it takes to publicly defuse the threat.’ The latter conjunction doesn’t literally say that you haven’t ruled out H, but it does suggest it. The Live Skeptic’s argument develops that suggestion.

Please note the strength of (9): it’s not merely saying that you can’t convince experts that H is false. Condition (9) doesn’t just mean that you lack reasons that would be accepted as establishing beyond serious dispute that H is false. No, it’s saying that even the experts who *reject* H would say that your reasons for rejecting H are clearly inadequate by just about anyone’s lights and not just by the lights of those who accept H. For instance, you might be a graduate student in paleontology who is aware of the rival hypotheses about the demise of the dinosaurs and who happens to believe the true meteor hypothesis. You go to see your PhD supervisor and she asks you what you plan to say about the supervolcano hypothesis in your dissertation. You say that that theory isn’t very plausible but you’re happy to throw in a brief section showing why it’s wrong. She *agrees with you* that the meteor hypothesis is correct but she asks you what you plan to say against the supervolcano hypothesis. You give your spiel and she tells you flat out that what you’ve said is clearly inadequate and you should either do much better with a critical section or drop it entirely and say in a footnote that you’ll be merely assuming the falsehood of the supervolcano theory. After all, professor so and so right down the hall is an advocate of the supervolcano theory, he’s certainly no dope, he isn’t alone in his expert opinion, and you’ve said nothing to put any pressure on his view.

The graduate student is a well-informed mere mortal, ‘mere mortal’ for short. A child is not, as she fails to satisfy either (6) or (7). Another kind of mere mortal is an expert in the field but whose specialization lies elsewhere. Professor Smith teaches various science classes. She is perfectly aware of the supervolcano and solar flare hypotheses but wouldn’t be able to say anything interesting against them. She has the true meteor belief but like the graduate student her belief is too lucky to amount to knowledge. *That’s* the type of person I have in mind as a mere mortal; my hope is that conditions (6)-(9) capture the important aspects of their epistemic position. Needless to say, there could be many such individuals.

The Live Skeptic now makes three key claims.

The Modesty Principle: if S is mere mortal with respect to live hypothesis H, then H isn’t ruled out with respect to S.

The Live Hypothesis Principle: if S knows that P entails $\neg H$ and is as aware as just about anyone that H is live (so H really is live), then if H isn’t ruled out with respect to S, S doesn’t know P.

The Liveness-Mortality Premise: in the actual world or a very close possible world many members of our actual intellectual community of contemporary analytic philosophers and

cognitive scientists are mere mortals with respect to each of the following live hypotheses (so conditions (1)-(9) all hold in that world).

Belief Error Theory: no one believes anything (endorsed to some significant degree at various times by Patricia Churchland (1986), Paul Churchland (1989), Stephen Stich (1983), Daniel Dennett (1978), Paul Feyerabend (1963), Richard Rorty (1970), and Quine (1960, 1985)).¹

Pain Error Theory: pain is only in the brain (endorsed by Russell, Broad, Ryle, David Lewis, and David Armstrong; see John Hyman 2003 for references).

Color Error Theory: no ordinary physical objects are colored (endorsed by Galileo, Larry Hardin, Paul Boghossian, David Velleman, Emmett Holmon, and Frank Jackson as well as by scientists Semir Zeki, Stephen Palmer, Werner Backhaus, and Randolph Menzel; see Alex Byrne and David Hilbert 2003 for references).

Trait Error Theory: no one has any character traits (endorsed by Gilbert Harman (1999, 2000) and John Doris (1998, 2002)).

The Live Skeptic's validly inferred conclusion is that in either this world or some very close possible worlds many of us don't know that fire engines are red, we don't know that we sometimes have pains in our lower backs, we don't know that John Rawls was kind, and we don't even know that we *believe* any of those truths. This skeptic does not say that any of those four error theories are true; she can safely assume they are (utterly, completely, comically) *false*. The *error theorist* says that we don't know that fire engines are red because they aren't red; the *Live Skeptic* says that we don't know that fire engines are red because even though our belief may well be true, it's not sufficiently warranted. Obviously, if the error theories are true then the Live Skeptic's conclusion holds, assuming knowledge requires truth. And for the very same reason if the error theories are truth-valueless, the Live Skeptic still wins. So no matter what the status of the live error theories (true, false, without truth-value), we mere mortals about those error theories don't know that anyone has any beliefs, anything is colored, etc.

My purpose in this essay is not to show that Live Skepticism is true. Although I write with the voice of someone who endorses Live Skepticism, I do so merely to present the positive case for it. I think the balance of considerations favors the Live Skeptic, but at this early stage of evaluation only

¹ For my purposes 'eliminativism' indicates just the view that no one believes anything (although it is convenient to conjoin it with the thesis that knowledge requires belief; see below). It is silent—neither pro nor con—on the status of other folk psychological concepts. Stich offers a very helpful summary of eliminativist arguments (eleven by my count!) in his 1996, 16-29. Since the details of these positions and arguments are both distracting and irrelevant to my purposes, I will not examine them here.

the foolish take sides. Instead, I want to show that the Live Skeptic's argument presents us with a new paradox of a traditional form: a group of claims that are mutually inconsistent but individually highly plausible. The interest in the paradox is not just in its skeptical solution but the consequences of giving up the Live Skeptic's premises.² As far as I can see, any realist response to the argument (one that denies the skeptical conclusion) has far-reaching consequences for the nature of knowledge and warrant. I will be indicating some of these consequences as the argument proceeds.

3. The Liveness-Mortality Premise

Consider belief error theory, or eliminativism as it's commonly called. Eliminativism is currently in our intellectual community a live possibility. The hypothesis is a going-concern among the experts. That is just a brute fact about our profession. You might not like it, but that is just the way it is. The Churchlands, Feyerabend, Stich, Rorty, Quine, Dennett, and others have at various points in their exceedingly distinguished careers concluded upon careful reflection that it is very probably true or as serious a contender as any other hypothesis about folk psychology. They arrived at these views in the familiar and epistemically responsible way of looking hard at data and lines of reasoning, and then thinking about them in highly intelligent, relatively unbiased and open-minded ways. These people aren't dopes. And it goes without saying that no magic bullet has been found for it and people still take it virtually as seriously as ever. What this means is that the live status of a hypothesis is a *sociological* fact, not a matter for philosophical argument.

Okay, I lied; eliminativism isn't live today or even in the 1980s or 1990s. Too few supporters. But if Fodor, Dretske, Burge, and a few other realists had died in 1974, the reaction to the eliminativist-leaning writings of Stich, Dennett, the Churchlands, and others had been more positive in the 1980s and 1990s, the 'Who cares what common sense says?' attitude in contemporary metaphysics was alive in the philosophy of mind ... You get the idea. Surely in some possible worlds practically next door to ours eliminativism is a live hypothesis in both the cognitive science and philosophy of mind communities. In case you weren't aware of it, color error theory is practically the dominant view among color scientists today; the same holds for error theory for pain location. And character trait error theory, although not actually live, could easily become live in the same way eliminativism could. My interest lies in the epistemic status of mere mortals in those actual or almost actual communities: what do *they* know about beliefs, colors, pain locations, or character traits? The live skeptical argument concludes that a society virtually just like ours but with sufficient philosophical confusion regarding eliminativism, for instance, could have a significant period of time (decades, even centuries) in which most philosophers and cognitive scientists knew nothing of the

² With an appropriate change to the Live Hypothesis principle, the Live Skeptic's argument can be turned into an argument for truly universal skepticism, saying that mere mortals with respect to live eliminativism know nothing at all. I explore this argument in an unpublished manuscript.

form ‘S believes P.’ That is shocking enough. And yet, for all we know, *our* society will soon enter into philosophical depravity and wholeheartedly vote eliminativism as the theory of cognition most likely to be true.

The Live Skeptic need not think that eliminativism (or the other error theories) is remotely plausible. Even Fodor could be a Live Skeptic! In fact, the Live Skeptic herself could not only *know* that eliminativism is false but *know* that she has refuted it and know that she knows it’s false—even though she claims that in her own philosophical community eliminativism is live. She could hold herself to not be a mere mortal; and she could be right about that because she really does have a refutation of eliminativism. Exactly the same holds in the more familiar dinosaur case: a paleontologist genius could *know* full well that the supervolcano hypothesis is mistaken and the meteor theory true while holding that her students don’t know either of those facts due to their mere mortality with respect to the live supervolcano hypothesis. Throughout this essay I assume with the Live Skeptic that the error theories are false and nowhere in the vicinity of the truth. All the Live Skeptic is claiming, in endorsing the Liveness-Mortality premise, is that there are easily possible philosophical communities in which the error theory in question is live in the sense of (1)-(5) and there are many ordinary individuals in those communities who satisfy (6)-(9)—individuals just like us philosophers in the actual world. There is no tension at all in endorsing that claim while holding that eliminativism is utterly implausible and has even been ruled out by oneself. Perhaps Kripke, Burge, and a few others have unpublished but rock-solid proofs that eliminativism is false; even the eliminativists would admit defeat if they just had a chance to see them. It doesn’t matter in the least because those proofs are irrelevant to the epistemic status of mere mortals in communities in which those proofs don’t exist or are known of by just a few philosophical hermits.

Am I claiming that if some group of nuts takes over philosophy or cognitive science, just due to the fact that they are clever enough to publish often but not wise enough to have any sense regarding anything, then whatever crackpot idea they latch on to automatically becomes live—something the rest of us level-headed folk are now epistemically responsible to even though we can see how silly it is? No. Let it be stipulated that the experts in question, the ones taking eliminativism, for instance, so seriously must have distinguished themselves in other areas of philosophy or cognitive science. And remember that the hypothesis has to have survived a thorough evaluation by such experts. Further, the experts have to actually think the hypothesis is true, not just interesting. (Contemporary epistemologists take the brain-in-a-vat hypotheses to be interesting but not because they think they have any chance at truth.) Throw in other conditions on liveness and being a mere mortal as you feel appropriate, adding to the nine I gave in the previous section. I trust that when you’re finished, our society still comes pretty close to satisfying the conditions and in very nearby possible worlds—or the actual future—easily satisfies them.

I assume that the only reason to balk at the Liveness-Mortality premise lies in condition (9). One might think that even if eliminativism is live in the sense of (1)-(5), one could still easily rule it

out by merely reflecting on the fact that one has beliefs. But (9) is just saying that virtually all experts would *judge* the evidence you have against H to be clearly insufficient to rule it out. It doesn't say that the experts are *right* about that. Perhaps every one of us can rule out eliminativism, for instance, merely by saying to ourselves 'It sure seems to me that I believe that $2 + 2 = 4$ '. I'll be considering such ideas later in the essay. But in the worlds in which eliminativism is live philosophers sufficiently skeptical about the truth of common sense will not consider such a line of reasoning as sufficient to rule out eliminativism (so there aren't too many John Searles there!).

One might suspect that in the actual world a large majority of philosophers of mind really do think that one can refute eliminativism merely by breezily reflecting on one's own beliefs. (Recall that we are assuming that eliminativism is false.) If so, then perhaps worlds that make (1)-(9) true are quite distant from actuality, contrary to the Liveness-Mortality premise.

I think the objector's claim about the actual world is mistaken, but it will help the Live Skeptic's cause if we get clearer on just how a world could make (1)-(9) true and yet remain very similar to actuality.

A quick answer is that just as color error theory is *actually live today*, as anyone familiar with the philosophy and science of color can attest, belief error theory could have the same respected status. Both error theories are radically and comparably opposed to common sense; so that hurdle can be met. Perhaps more convincingly, one additional feature, mentioned above, easily makes (1)-(9) true at a world: more people being skeptical about the strength of the connection between common sense and truth. As it stands today, in the philosophy of mind and epistemology communities there is a strong confidence that statements that are conversationally appropriate in ordinary contexts are true. Epistemologists and philosophers of mind are typically loath to claim that large parts of common sense are radically mistaken. In metaphysics and the philosophy of language and logic, however, one does not find this faith in the inference from common sense to truth. (Think of theories of vagueness, truth (the semantic paradoxes), material constitution, and identity through time—not to mention the philosophy of physics!) Indeed, it is hard to find any contemporary metaphysician who *doesn't* hold claims that are about as outrageous, from the perspective of common sense, as eliminativism. An acquaintance of mine who is sympathetic to dialethism (and with a very impressive CV) thought Live Skepticism was pretty obviously true. In my experience, those who work only in epistemology or the philosophy of mind often find this attitude highly dubious or even hard to fathom, but there are good reasons why it is prominent in other areas of philosophy. Perhaps the paradoxes that engage the metaphysicians and philosophers of language, logic, and physics are much harder than those in epistemology or the philosophy of mind, which would causally account for the difference in attitude, but given that metaphysics, the philosophy of language, and the philosophy of logic have shown us that even some of the most entrenched parts of common sense simply have to be rejected, one wonders whether fidelity to common sense in epistemology and the philosophy of mind is a good constraint on theorizing.

I don't want to pass any judgment on the faith in the truth of common sense. I personally suspect it really is *mere* faith. But the Live Skeptic needs a much weaker claim, one operative in her rebuttal to the objection that worlds in which (1)-(9) are true are remote from actuality: the skeptical attitude we find in contemporary metaphysics, philosophy of language, and philosophy of logic regarding even the most well-entrenched and universal common sense could easily permeate the philosophy of mind while eliminativism had many more expert supporters and fewer detractors.³ I'm not claiming that such an attitude movement would be correct or fully justified; I'm just saying that it could easily happen. In those worlds, which I claim are highly realistic in being very similar to the current time of the actual world, eliminativism is live and many of us are mere mortals regarding it. And since *we* inhabit such possible worlds, the use of the second-person in motivating Live Skepticism is appropriate. So we have justified the Live Skeptic's claim that worlds satisfying (1)-(9) are next door to actuality.

4. The Nature of Live Skepticism

In later sections I will address the strongest objections to the Live Skeptic's argument. I can't consider all the reasonable objections, and with regard to the ones I do consider I don't always treat them as thoroughly as I should, but I hope to provide my basic reasons for thinking that none of them succeed.⁴

They will fall into two classes. First, there are reasons for thinking that we don't have to rule out, in any significant sense, the error theories in order to know truths obviously inconsistent with them. For instance, relevant alternative and contextualist theories are important here. These objections fault the Live Hypothesis principle, which says we do need to rule out the error theories. Second, one can admit the ruling out requirement but argue that it's actually pretty easy to do the ruling out. Reliabilism and introspective accounts figure here, among others. These objections target the Modesty principle, which says that mere mortals can't manage the ruling out.

But before we put our critical hats on we must understand the nature of Live Skepticism, for it is quite different from any traditional skeptical theory. In fact, it's so strikingly different, in many ways, that to a certain extent I lose the *motivation* to find fault with the Live Skeptic's argument. That is, upon realizing what Live Skepticism amounts to I have a hard time justifying my gut reaction that

³ How large does the group have to be? Surely precision is out of place here. There is a group of well-informed, et cetera philosophers who take seriously the possibility that there are no chairs. But there aren't enough of them to make it live, at least in my judgment. More to the point, that hypothesis does not meet all five of the liveness conditions. There are deep issues involved in determining what kinds of hypotheses generate decent Live Skepticisms when plugged into the argument template of the previous section.

⁴ In my 2005 book I address more objections.

Live Skepticism is counterintuitive—as odd as that sounds. What this means is the subject of this section.

It is a brute historical fact that we are confronted with three seemingly outrageous propositions when considering traditional skeptical arguments: (a) the skeptical hypothesis itself (e.g., the BIV hypothesis), (b) the conclusion that we know next to nothing (or nothing of a certain form, e.g., external world skepticism), and (c) the claim that some outrageous skeptical hypothesis dashes our hopes for ordinary knowledge. In this section I examine these three alleged “outrages”, (a)-(c), applied to the live skeptical arguments. I will argue that the live skeptical arguments significantly diminish all three. In particular, (b) and (c) are wiped out almost completely. So the nearly instinctive reaction ‘This is an argument for a kind of skepticism; so there’s a mistake in it’ is definitely not intellectually responsible (if it ever was). This is not to deny that the skeptical conclusions and arguments generated from eliminativism and color error theory (the two skeptical hypotheses I focus on most in the remainder of the essay) are less than wholly innocent; they have their own counterintuitive elements. Even so, the considerations to follow make me think that the option of simply accepting the soundness of the live skeptical arguments is at least as good as any of the anti-skeptical options.

A. Outrage: Skeptical Hypotheses

Traditional skeptical hypotheses are completely outrageous. The hypothesis that you’re a BIV is just about the most outrageous thing any philosopher has come up with. But skeptical hypotheses need not be nearly as outrageous as they have traditionally been. Suppose for a moment that the eliminativist and color error theories are actually true.

- Is there a physical world? Yes.
- Is there a planet Earth? Yes.
- Do you have two hands? Yes.
- Were the New York Yankees the best team in baseball in the 1950s? Yes.
- Are you a cognitive being functioning perfectly well in navigating the world? Yes.
- Are there any devious mad scientists or evil demons manipulating your mind? No.
- Are there any brains in vats enjoying sensory lives akin to yours? No.

Practically everything is the same in the true error theory world—a world in which all the live skeptical hypotheses are true—as it is in our world with the error theories all false. Ask a group of non-philosophers to make utterly random lists of claims about the world they think are definitely, expertly, and widely acknowledged to be true. Even if the live skeptical hypotheses are true, virtually everything on the lists can be perfectly correct—that is, true even though we have no beliefs, tomatoes

aren't colored, etc. These features of the live skeptical hypotheses demonstrate how utterly different they are from classic skeptical hypotheses.

Only someone literally insane would think that any of the traditional skeptical hypotheses is true; but clearly sane, intelligent, sober, and informed individuals support the live skeptical hypotheses.

A: You have no beliefs. Perhaps better yet: no tomatoes are red.

B: What? That's impossible!

A: For all the philosophers and cognitive scientists have been able to determine, you have a whole slew of mental states that pick out things in the world and successfully guide your actions, but they aren't technically beliefs. Not just any cognitive state has what it takes to be a belief, and it turns out that although you have a very rich mental life, just as you have always thought and experienced, the category of belief just doesn't apply. Instead, an entirely different set of categories applies to your mental states. The category of belief is a primitive yet persistent and oddly useful illusion. The real cognitive categories will come from science. Further, although you perceive things as being colored, they aren't really colored; the color is just in your perception for such-and-such reasons.

B: Oh. That explains it (a bit anyway).

B. Outrage: Skeptical Conclusions

According to the Live Skeptic most philosophers and cognitive scientists (the obvious candidates for mere mortals) know nothing about our beliefs, pain locations, character traits, or colors. Just why is it that skeptical conclusions, restricted or not, are judged to be so bad, so utterly shocking, even horrifying? The new skeptic's conclusions are not outrageous in my view, although they are philosophically significant. Live Skepticism is not only easier to swallow than the traditional forms of skepticism but isn't that hard to swallow, period. I will present four arguments for that claim.

First, the Live Skeptic is claiming that we lose our knowledge by means of a perfectly familiar mechanism.⁵ Pretend that it had been common sense for hundreds of years that whales are fish; pretend it's true too. Then scientists came along to challenge that bit of common sense. They admit that there are many fish and many whales, but they have some impressive arguments for the shocking claim that whales aren't fish—they are mammals instead. If you take a quiz listing kinds of fish and you write 'whale', scientists will mark your answer as incorrect. If you say to your child, 'Look at the whales Julia; they're the biggest fish there are', what you say is fine in some respects but

⁵ Here I'm assuming that before we become mere mortals, say as children, the liveness of the skeptical hypotheses did not sabotage our knowledge. More on this point later in this section.

really false. For the purposes of this essay pretend that the scientists are wrong: whales are fish. So one familiar and even highly reliable method for finding out that something is a fish—find out it’s a whale—would clearly be question-begging against the group of scientists. In such a scenario, if you were a mere mortal regarding the whale-fish controversy, then you would be familiar enough regarding the relevant issues that in order to know that Keiko (the whale) is a fish you would need to have some way of neutralizing the live whales-aren’t-fish hypothesis.

Consider the advocate of the pain-only-in-the-brain hypothesis and compare what she says with what the whales-aren’t-fish advocate says (keep in mind that we’re pretending that whales are fish and toes sometimes throb in pain).

- I realize that it’s common sense that whales are fish.
 - I grant that there are loads of fish and loads of whales and to a certain extent there’s nothing wrong with classifying whales as fish.
 - But they aren’t fish, not really.
 - Whales are actually mammals, not fish, for various technical reasons.
 - There are zillions of fish to be found, but they aren’t to be found among the whales.
-
- I realize that it’s common sense that toes sometimes throb with pain.
 - I grant that there are loads of throbbing feelings and loads of toes and to a certain extent there’s nothing wrong with classifying some toes as throbbing.
 - But they aren’t throbbing, not really.
 - The throbbing feeling is actually in your brain, not your toe, for various technical reasons.
 - There are zillions of throbbing feelings to be found, but they aren’t to be found among the toes.

I take it as intuitive that in the envisioned whale-fish scenario mere mortals wouldn’t know that Keiko the whale is a fish—even though they used the familiar, commonsensical, and highly reliable ‘If you find a whale, then you’ve found a fish’ method to form the true belief that Keiko is a fish. By analogy, just because you (a mere mortal) have a true belief that Mary’s toe is throbbing and you formed it in a common way—see her stub her toe, curse, and tell you it’s throbbing painfully—this method, as familiar, commonsensical, and reliable as it indeed is (in that world and our world), is not sufficient to make you know that Mary’s toe is throbbing. Not when you’re a mere mortal with respect to the live pain-only-in-the-brain hypothesis. Again: the *reason* for the lack of knowledge is exactly the same as in the whales-aren’t-fish case. The Live Skeptic is saying that our knowledge is sabotaged due to the operation of a perfectly familiar epistemic mechanism, the one governing the whales-fish or dinosaur cases.

Now for the second and most interesting reason the new skeptical conclusions are not so bad. Suppose once again that Live Skepticism is true. Suppose that tomorrow the definitive refutations of color error theory and eliminativism are published and widely digested. It's reasonable to think that we *then* know that fire engines are red and that Moore believed skepticism false, for at that point the error theories have been ruled out on our behalf. In addition, it is plausible to hold that since my father knows nothing of eliminativism, color error theory, professional philosophy, or cognitive science, *he currently* knows that fire engines are red and that my mother believes that motorcycles are a menace. It's also reasonable and consistent with the new skeptic's position to think that my father knew those facts years ago before eliminativism (and, let us pretend, color error theory) was even on the radar screen. Finally, we should hold that before we were mere mortals, say when we were children, we had loads of knowledge inconsistent with the skeptical hypotheses even if the latter were already live.

This is just to say that the reach of the epistemic threat posed by the live skeptical hypotheses isn't sufficient to affect those people for which it isn't live. In a way, this makes the new skepticism modest: by saying it rules out knowledge *only* for mere mortals, we are restricting its scope severely. In other ways, the new skepticism is now unlike anything we have ever encountered.

Here's why. Skepticism has always been considered a permanent, blanket, and highly negative condition. If you're in it, then everyone else is in it too (that's 'blanket'), forever ('permanent'); and this marks a significant epistemic deficiency on our part ('highly negative'). If skepticism is right, then it must reflect an everlasting, perhaps absolutely necessary, failure of our cognitive systems to achieve a certain result. The idea that some kind of radical even if restricted skepticism could be the result of a purely accidental, contingent, and temporary confluence of chance social events affecting a fully rational, cognitively well-off individual seems absurd. And it is absurd, at least for classic versions of skepticism. But what we have seen is that skepticism can be thus accidentally generated. The skepticisms attributed in this essay are relatively fleeting and are no indictment of our cognitive systems or evolutionary progress. Most surprisingly, the new skepticisms are the odd result of cognitive systems working *well*, not poorly: it is through the epistemically beneficial practices embodied in the premises of the new skeptical argument template that we have temporarily fallen into our regions of restricted skepticism. By being part of a community that pays due heed to expertly-produced contrary evidence we have temporarily robbed ourselves of large portions of knowledge, but we have done so as the result of following epistemic practices that almost always actually *buy* us knowledge. Once we understand how we got into the mess of not having knowledge of color, belief, or pain location, for instance, we see that it's not a bad thing. In fact, *we are epistemically better off than we were before!* Yesterday you knew; today you don't; but today you're smarter than you were yesterday regarding color, belief, and pain. People unacquainted with philosophy or color science or cognitive science may know more than we do, but this just shows that we need a new and improved measure of epistemic standing. Suppose I think that on balance color

error theory is very likely mistaken and so even in the full awareness of the live status of that theory I continue to hold ordinary color beliefs. My mother's belief that fire engines are red is warranted enough for knowledge; mine is not, as I remain a mere mortal; but in some sense I'm in the superior epistemic situation with regard to the color of fire engines despite my lack of knowledge and lack of immortality vis-à-vis color error theory.

That may sound contradictory: if token belief A (my mother's) is better warranted than token belief B (mine), and A and B have the same content (e.g., that fire engines are red), then of course A is epistemically better than B. In the previous paragraph I had the Live Skeptic deny this conditional, but she need not. Perhaps the *collection* of my attitudes regarding color is epistemically superior to hers, even though when it comes to the *particular* fire engine belief her position is superior to mine. The details on how best to describe the differences in the two believers are interesting, but in any case the end result is that a typical reaction to skepticism—the kind of epistemic failure attributed to us by the skeptic just couldn't be built into our cognitive systems—does not touch the live skeptical theses. Similarly for the objection that runs: according to the skeptic we are cut off from reality, unable to know it, and that's just implausible. On the contrary, by becoming a mere mortal and thereby falling into the live skeptical trap I am in a better epistemic position than I was before; skepticism is an *improvement*, something to be bragged about, not ashamed of. Falling “victim” to the live skeptical snare does not cut me off to the reality of color, belief, pain, and knowledge compared to those non-philosophical folk who in a real sense know more than I do. Instead, I'm the one better in tune with the facts; my opinions dig deeper into the nature of reality than theirs do. This is not too surprising: by knowing more about the possibilities regarding color and belief, I have a better appreciation of the “whole color thing”, or the “whole story regarding cognition.” The ordinary person may know that fire engines are red, while I do not, but her belief has a measure of accidentalness that mine will never have. If she had just been privy to the intelligent and sophisticated ruminations of some philosophers and scientists, then she would no longer know. It's just an accident that she does know. The nature of the accidentalness is not so great as to rob her of knowledge, or so the anti-skeptic says and the Live Skeptic may admit (more on this point in a few pages), but coupled with my expanded knowledge of possibilities and evidence for and against those possibilities, it is enough to render her overall epistemic situation less secure than mine.

Furthermore, *I've got all the positive warrant that she does* when it comes to our beliefs that fire engines are red. This point is so important I'll put it in a box, as if it were in a science textbook.

The Live Skeptic is not denying us any
of the warrant we think we have;
neither is she denying its quality.

Engage in the useful fiction that warrant comes in units (we can make the same point without the fiction). She has 1200 positive warrant units for her fire engine belief; I have just 800; one needs 1000 for the true belief to be knowledge. I have had all the same fire engine and color experiences that she has had; I have each of her 1200 units. My only problem is that I have 400 *negative* warrant units, coming from my foray into color science and philosophy. Just because she knows and I don't does not mean that I lack some positive warrant that she has. Perhaps we need to judge the epistemic standing of beliefs with additional measures, such as with positive warrants (I have all the good reasons, reliability facts, etc. that she has) and absolute value warrants (since the absolute value of my warrants is much greater than hers, I'm much more familiar with the relevant issues regarding the content of the belief).

We get the same result without thinking of competing warrants in quantitative terms. It seems plain that although as a mere mortal all my previous positive warrant for believing that fire engines are red is still in place and is still sufficient, *if left alone*, for knowledge (that was the 1200 "units"), there is negative warrant sufficient to offset or cancel or "veto" the positive warrant. It seems to me that some sources of negative warrant produce a veto of most any positive warrant for the belief in question. If H is live, S is a mere mortal with respect to H, and H obviously entails not-P, then H vetoes the warrant S has for her belief in P, where 'vetoes' means 'makes insufficient for knowledge.'

Thus: if you still believe that fire engines are red and you worry, at least a bit, about the possibility that you have fallen into the skeptical trap, be assured that this trap is quite unlike the traditional ones. You can hold your head up high and boastfully declare your arrival in the pit of skepticism. If you can then actually crawl out of it under your own power, by ruling out the skeptical hypotheses in question, then you should get a distinguished chair in the philosophy department of your choice.

These points on epistemic standing inside the skeptical snare are bolstered by the third reason that Live Skepticism is palatable. Suppose Live Skepticism is true, so you don't know that the table is brown; neither do you know that you believe that. Still, no one said that you didn't *know* perfectly well that the table looks brown and is for all practical purposes brown. Further, you still know that red is darker than yellow. In addition, I still know that I act like I believe P, that it seems to me that I believe P, that I feel like I believe P, that it seems as though Mary believes Q, etc. I also know that for all practical purposes Moore believed skepticism false. And I know that belief and truth are required for knowledge. If that's the case, then what have we lost, really, in losing our knowledge that fire engines are red? He knows that fire engines are red; we don't; but we're familiar with much more color theory than he does and we still know that fire engines look red and are red for all practical purposes. Now who is in the better epistemic position vis-à-vis the redness of the fire engine?

Still, one might find it odd that the uninformed know while we well informed mere mortals fail to know. Gilbert Harman (1980) and others have investigated so-called “social” Gettier cases that are somewhat similar to the cases I’m focusing on and that may seem to put pressure on my suggestion that uninformed non-mere mortals know while mere mortals don’t know. Jane and Jon read in the newspaper that Dr. Kirby has been assassinated. The story is a true account by a fully competent eyewitness. On the basis of that accurate report both Jon and Jane come to believe that Dr. Kirby was assassinated. However, starting the next day there is widespread publicity of credible (but ultimately misleading) denials of the assassination by well-placed public officials (who are lying). Jane is unaware of those denials as she didn’t read the newspapers that day; Jon did read them. Does Jane still know the assassination occurred? Jon is a mere mortal with respect to the no-assassination idea; Jane is the non-mere mortal as she is totally uninformed of that hypothesis (and we can pretend that she has no special knowledge that could show the denials to be mistaken). If we tell the story right, Jon doesn’t know the assassination took place, even if he retains his true belief, because he has read all about the live no-assassination hypothesis. Presumably, Jane doesn’t know either (although many people’s intuitions are not so strong on this matter). If so, then the live hypothesis rules out knowledge for the non-mere mortal as well as the mere mortal, contrary to my suggestion.

However, the Live Skeptic’s case is different in at least one crucial respect, one that Harman himself notes. Jane is a member of the epistemic community in which the denials are public. She read about the assassination in the *New York Times* on April 5th; the denials were published in the same paper on April 6th. She is epistemically responsible to the April 6th information in the sense that it affects her belief’s epistemic status. However, she doesn’t read any philosophy or science. She knows absolutely nothing of the philosophy of color or color science. She isn’t part of those epistemic communities. So there is less reason to think the color error hypothesis poses a threat to her belief that her socks are blue as there is reason for thinking that the no-assassination hypothesis poses a threat to her belief that Dr. Kirby was assassinated.⁶ Furthermore, the method she used to obtain her belief in the assassination is the very same method for obtaining the defeater for that belief: relying on the newspaper’s reports. This isn’t true in the color case: she relies on a relatively primitive perceptual experience to obtain her socks-are-blue belief and would have to endorse a completely different method for obtaining evidence against that belief (i.e., relying on the testimony of expert color scientists or philosophers).

In any case, I’m not claiming that Live Skepticism does not apply to people who know nothing of the error theories. I’m just allowing for that relatively likely possibility. If it does apply to them, then Live Skepticism is even stronger, covering a much wider range of beliefs, although it still doesn’t affect the millions of beliefs not targeted by live contrary hypotheses.⁷

⁶ I say a bit more about cases like this one later in this section.

⁷ Thanks to the referee’s comments that led to the previous three paragraphs.

The fourth and last reason for not getting too upset over the skeptical conclusion lies in the Live Skeptic's response to an intriguing anti-skeptical argument. Even a skeptic has to admit that we divide up knowledge (and justification) claims into two huge piles 'Appropriate' and 'Inappropriate' according to relatively stable, cross-cultural, intersubjective criteria. Further, we agree on what goes in the piles, at least for an important class of "commonsensical" claims; call those claims 'core.' Common sense says that for those core claims—and there are lots of these claims—the ones that show up in the pile marked 'Appropriate' are virtually all true and the ones that show up in the other pile are virtually all false. Now revisionist theory T comes along and it disagrees wildly with common sense, at least in a sense. It agrees with common sense that we sort the knowledge claims into the two enormous piles, 'Appropriate' and 'Inappropriate.' It agrees that we do this according to some stable criteria. It agrees with common sense regarding what core claims are in fact sorted into which piles. That is, when common sense says that core claim C gets sorted, as a matter of brute fact, into 'Appropriate' or that core claim D gets sorted into 'Inappropriate', theory T concurs.

What T vigorously disputes is common sense's claim that virtually all the core 'Appropriate' knowledge claims are true and virtually all the core 'Inappropriate' knowledge claims are false. For a significant and unified portion of the core claims common sense puts in the 'Appropriate' pile, the skeptical revisionist says that they are false.

To take a concrete example that targets the Live Skeptic, it's just plain weird that Sam the mere mortal regarding color is just as competent regarding color as Samantha the non-mere-mortal and yet Samantha but not Sam knows that some fish are purple—as they personally encountered the fish while scuba diving together and both believe it in the usual way based on those veridical perceptions. The two attributions 'Sam knows that some fish are purple' and 'Samantha knows that some fish are purple' seem equally good in some important and epistemically central sense. They are "warrantably assertable" to the same extent. That Sam knows that some fish are purple is the judgment of common sense and common language. These judgments are not to be dismissed frivolously. We end up with several questions the Live Skeptic needs to address:

- Why is it that 'Sam knows that some fish are purple' is false, as the Live Skeptic claims?
- In what sense are 'Sam knows that some fish are purple' and 'Samantha knows that some fish are purple' equally good, even though according to the Live Skeptic only the latter is true?
- Why is it that 'Sam doesn't know that some fish are purple' seems inferior to 'Sam knows that some fish are purple' even though according to the Live Skeptic the former is true and the latter is false?

These are good challenges. The traditional skeptic has problems answering the analogous questions for, e.g., ‘Sam knows he has hands’. But I think the Live Skeptic can answer them.

Regarding the first question, the Live Skeptic’s skeptical argument provides the explanation, as I was trying to indicate in the first of these four reasons why Live Skepticism is palatable. That argument says that Sam doesn’t have knowledge because of a wholly familiar mechanism: losing knowledge in the face of *recognized* expert contrary evidence and opinion. The phenomenon is perfectly ordinary.

The other two questions require answers that relate truth and the conversational appropriateness of knowledge attributions. We begin to answer them by noting that the phrase ‘knowledge attribution’ is ambiguous in the following relevant way. It may stand for the truth or falsehood attributed, the *proposition* that S knows P, or it can stand for the action of attributing knowledge: the *attributing* of a piece of knowledge. Whether an attribution is conversationally appropriate is obviously dependent on whether it serves certain purposes, especially the ones peculiar to the context of conversation. That’s pretty much what ‘conversationally appropriate’ amounts to. Now one purpose is to have the proposition expressed by the use of ‘S knows P’ true. But there are other purposes and often those are more important. (Think of the idiosyncratic ways you converse with someone you’ve known for many years.) So a use of ‘S knows P’ is conversationally appropriate if and only if it serves the purposes relevant to that use. We cannot complain that the right-hand side of the biconditional in the previous sentence is vague: it’s that vague because the left-hand side is that vague.

The Live Skeptic can employ the action-proposition distinction to help answer the two questions raised above (where we assume that the scuba diving subject Sam correctly believes, in the usual experiential way, that some fish are purple). For convenience, let ‘practicality’ be short for ‘conversational appropriateness.’

1. Any use U of ‘Sam knows that some fish are purple’ expresses a false proposition when uttered by another mere mortal. This is a sufficient condition for live skepticism for color (given the quality of Sam’s epistemic situation relative to that of other people). This says nothing regarding the practicality of the action U.
2. A use U of ‘Sam knows that some fish are purple’ is an action that is practical if and only if U serves the purposes relevant in U’s context.
3. As a matter of brute fact, for zillions of potential uses U of the sentence ‘Sam knows that some fish are purple,’ U is practical. That is, *those knowledge claims succeed at being practical*. These practicality facts are the truths that common sense is tracking. The skeptic can say that a use of ‘Sam knows that some fish are purple’ is

“warrantedly assertable” if and only if practical if and only if it’s conversationally appropriate if and only if it serves the purposes relevant in U’s context.

4. Why is (3) true? Why do all those false knowledge claims succeed at being practical? The answer has three parts.

- (a) Sam has *all* the usual good reasons for believing it that people who know it have. More generally, Sam has all the positive warrant Samantha has.
- (b) Sam has *all* the usual abilities you would expect in someone who knows it. For instance, he can point out the purple fish, distinguish them from red ones, tell you that they look a different color when exposed to sufficiently different light, etc.
- (c) The attributor of the knowledge claim ‘Sam knows that some fish are purple’ most likely isn’t currently thinking about color error theory and so does not currently and consciously appreciate how it threatens Sam’s belief. The same very likely holds for her audience. Since that threat is the *only* significant threat to Sam’s belief, it follows that the only significant threat to Sam’s belief is very likely currently and consciously out of everyone’s mind—thus contributing to the appropriateness of the knowledge attributor’s action.

5. A use U of ‘Sam knows that some fish are purple’ seems to express a true proposition (even though it doesn’t) because:

- (a) U expresses an action that is practical (as explained in (4)),
- (b) we don’t usually distinguish the practicality of the action from the truth of the proposition, *especially* when (4c) is true, which it is in virtually *all* ordinary contexts,
- (c) the practicality of the action is often dependent, but only in part, on the truth of the proposition, and
- (d) a use of ‘Sam *doesn’t* know that some fish are purple’ seems to express a false proposition.

6. Why does a use U* of ‘Sam *doesn’t* know that some fish are purple’ seem to express a false proposition, as noted in (5d) (even though it expresses a true proposition)?

Because:

- (a) U* expresses an action that is impractical,

- (b) we don't usually distinguish the impracticality of the action from the falsehood of the proposition, *especially* when (4c) is true, which it is in virtually *all* ordinary contexts, and
 - (c) the impracticality of the action is often dependent, but only in part, on the falsity of the proposition.
7. Why is a use U* of 'Sam *doesn't* know that some fish are purple' impractical, as noted in (6a)? The answer has three parts.
- (a) It strongly suggests to us that Sam doesn't have all the usual good reasons (more generally, warrant) for believing it that average people who know it perfectly well have. But that claim is false; Sam has the usual good reasons. Thus, since U* strongly suggests something that is false, and recognizably so, we think U* is false; and that makes U* impractical.
 - (b) It strongly suggests to us that Sam doesn't have all the usual abilities you would expect in the average person who knows perfectly well that some fish are purple. But that claim is false; Sam has all those abilities. Thus, since U* strongly suggests something that is false, and recognizably so, we think U* is false; and that makes U* impractical.
 - (c) If someone did employ a use U* of 'Sam *doesn't* know that some fish are purple,' the audience would most likely not be currently thinking about color error theory and so would not consciously appreciate how it threatens Sam's belief. Since that threat is the *only* significant threat to Sam's belief, it follows that the only significant threat to Sam's belief is very likely out of everyone's mind—thus contributing to the inappropriateness of the use of U*.
8. A use U of 'Sam knows that some fish are purple' seems to express a false proposition—but only to those of us who have thought hard about the new skepticism, and only in some of our moods—because, well, U *does* express a false proposition and we can easily come to suspect that fact by thinking hard about the new skeptical arguments.

What we have here in (1)-(8) is a good part of an explanation of how the skeptic can stick to her skepticism while admitting the undeniable practical irrelevance of her result as well as the undeniable utility of the false knowledge attributions.

It's important for us to be clear that the explanation does *not* rely on some controversial distinction between semantic content and pragmatic content—in fact, it relies on no such distinction,

controversial or not. Neither does it rely on any controversial claims about implicature. It does employ implicature-like claims as described in (7a) and (7b), but these are entirely commonplace as well as highly intuitive. Obviously, ‘S doesn’t know P even though P is true and S believes it’ often strongly suggests that S doesn’t have the warrant ordinarily sufficient for knowing P; and just as clearly it suggests that S doesn’t have the abilities you’d expect someone to have who does know that P. Nothing mysterious in those suggestions. Further, it is clear that the connections in (7a) and (7b) are defeasible, in that it is easy to imagine circumstances in which their suggestions prove false (i.e., ‘S doesn’t know P’ is true even though S has all the usual warrant and abilities of someone who does know P). Many “warranted assertability maneuvers” are pretty suspect because they rely on controversial semantic-pragmatic content distinctions and controversial implicatures. Fortunately, my explanation avoids those problems.

Here is an uncontroversial and analogous use of (7), loosely based on examples present throughout the literature. Suppose senator Smith is in Bigtown, Smith’s aides have told the media that Smith is in Bigtown (the media aren’t in Bigtown so they don’t have direct evidence), and the media have dutifully told the masses that Smith is in Bigtown. George has seen the newscasts and read the relevant newspaper articles that say Smith is in Bigtown. But while he was seeing those newscasts and reading those newspapers, his sister, who is one of senator Smith’s aides, told George that although Smith’s people were going to tell everyone Smith would be in Bigtown that weekend, in reality Smith was going to be in Bribetown on secret business. George’s sister has let him in on secrets like this before and she has always been proven right later, when the media found out about the deception. This time, however, George doesn’t believe his sister. He believes that Smith is in Bigtown, based on what the media have reported. He disregards his sister’s assertions because, perhaps, he unwisely thinks that Smith wouldn’t take bribes. He knows that most senators take bribes all the time, but for some bad reason he suspects Smith wouldn’t do it (which is false). It turns out that Smith is in Bigtown because the secret business in Bribetown fell through at the last minute so Smith didn’t go there at all (George’s sister didn’t know of these last minute developments).

It’s pretty clear that George doesn’t know that Smith is in Bigtown *even though it’s true, he believes it, and he has all the positive warrant for his belief that the average person does*—the average person who *knows* that Smith is in Bigtown based on that warrant that George has. He doesn’t know because he has the expert testimony of his sister, her perfect track record on this sort of thing, his knowledge that most senators take bribes, etc. And this isn’t a case of someone *first* knowing P and *then* losing knowledge because of the emergence of good but ultimately misleading contrary evidence that some people may have but *the subject knows absolutely nothing about*. That kind of case, an instance of which is Harman’s assassination story given earlier, is controversial, as that subject is unlike George in that she knows nothing about the contrary evidence.⁸ Since George

⁸ See Lycan 1977, Harman 1980 (especially note 7), Ginet 1980, and Conee 2001.

got his sister's advice at the same time he was watching television and looking at the newspapers (his sister was there in the room with him), he was never in any position to come to know that Smith is in Bigtown.

The assertion 'George doesn't know that Smith is in Bigtown' is true but a bit misleading. And it's misleading because it suggests that he doesn't have the warrant people typically have for knowing that Smith is in Bigtown—seeing the television and newspaper reports, for instance. That is, that sentence suggests that he didn't see those news reports, or if he did see them he didn't believe them. But he did see them, and he understood them perfectly well and even came to believe them. So 'George doesn't know that Smith is in Bigtown' is quite misleading in many ordinary contexts even though it is uncontroversially true. That's (7a). (The other claims in (1)-(8) apply to George as well.) The same holds for 'Sam doesn't know that those fish were purple': he did see the fish, he understood what was going on when he saw them, and he came to believe that the fish were purple. So 'Sam doesn't know that the fish were purple' is quite misleading but true anyway—just like in the case of George. We could tell an analogous story for (7b).

C. Outrage: Skeptical Hypotheses Preventing Knowledge

What seems outrageous, and probably *correctly* so, is the claim that the bare existence of some crazy hypothesis ruins your warrant for your beliefs. How could such a nutty, irresolvable, and scientifically baseless hypothesis such as the BIV hypothesis take away your knowledge? Indeed, that would be shocking.⁹ But none of this holds for the live skeptical hypotheses. Those hypotheses are live, scientific hypotheses. So it is entirely reasonable that their status as live and obviously inconsistent with P means that you need to rule them out in order to know P.

Now the new skeptical results look odd but not debilitating; strange but not indicative of any significant epistemic failure. Does that make the results philosophically uninteresting? Yeah, right: you don't know that you believe that Moore thought skepticism false; neither do you know that fire engines are red or that pains occur in chests. Despite all that, you are in a better epistemic position than those who aren't in the skeptical traps (other than the immortals of course). You never would have accepted any of those results without fighting tooth and nail.

5. Precedents for Live Skepticism

In this section I would like to show that skeptical arguments akin to the new ones can overthrow knowledge of color and pain location. This is important, I think, because it is reasonable to doubt that

⁹ In fact, it's a reason for thinking contextualism—or any other theory—is mistaken in giving in to the classic skeptical arguments in *any* actual context, even ones having to do with skeptical epistemological reasoning.

any set of circumstances could overthrow such fundamental and effortless knowledge. Overthrowing George's knowledge of Senator Smith's whereabouts or the student's knowledge of the cause of the demise of the dinosaurs is one thing; overthrowing knowledge of the color of your socks is something else entirely, as the latter seems *much* more secure and immediate given that we're assuming that the socks are colored. If I come to a true belief regarding simple color, belief, and pain location facts and do so for the "right" reasons or experiences—the reasons or experiences that in ordinary circumstances are abundantly sufficient for knowledge—then how could I not have knowledge? It's obvious that this can happen in other cases (e.g., the cases of George and the student paleontologist), but can it happen in the case of one's own beliefs or sock colors? Are there any *clear* and realistic examples in which a live hypothesis overthrows one's knowledge of one's own beliefs or sock colors? It would be nice to see some examples before we subscribe to the Live Skeptic's idea that there are worlds in which *vast amounts* of such knowledge are overthrown due to live skeptical hypotheses. Before you convince me that there are dozens of people over ten feet tall, I'd like to see at least one.

You see a sock in the usual excellent viewing conditions. It looks, and is, blue. But it's your colleague's sock, and his wife is a color scientist and he insists that he is wearing some of her "trick" socks she uses in her experiments, in that although they look blue and normal, they're actually very weird and really green. We can suppose that he's made an innocent mistake in that the socks he is wearing are entirely normal and blue. You mistakenly think he's trying to fool you even though he's actually a pillar of honesty, so you stubbornly persist in your belief that the socks are blue. Suppose his wife comes in and says 'Well *there* are those trick socks! We were looking for them all morning in the lab! What are you doing with them on?'. Other people concur with her (her lab assistants and children say). She and other color theorists have created various other strange objects, strange in ways having to do with their color appearances. You are somewhat aware of these objects, involving rotating disks, unusual materials, and the like. So you know of the existence of such objects. Your blue socks belief is true and reliably produced in the entirely ordinary way, but does this belief constitute knowledge once you've encountered the weird-green-socks story, especially given that you've heard and understood loads of intelligent, sincere, and honest experts saying that the socks are really green—not just his wife, but her assistants, other professors, etc.? Don't you have to rule out the weird-green-socks hypothesis to retain your knowledge? I think you would be in violation of some significant epistemic duties if you retained your belief. Do you keep your knowledge in these circumstances, in spite of your epistemic violations? It seems not. You believe the truth, and you do so for the right reasons, but in circumstances such as these that isn't enough.

We can tell analogous stories for some of the other skeptical hypotheses. During surgery you have a particular pain that seems to be, and really is, in your spine. The neurologist correctly asserts that this illusion happens all the time with this kind of operation but goes on to mistakenly say that your pain is really in your brain as your spine is perfectly sound and not even being operated on. For each of these "one-off" cases we get the result that persisting in your belief is in serious violation of

some epistemic duties. Further, it seems as though you no longer *know* that the socks are blue or that the throbbing pain is in your spine. If this is right, then we have one-off skeptical results, but skeptical results all the same, each based on *highly realistic* skeptical hypotheses and each denying instances of the knowledge denied by the Live Skeptic. We're all aware of the clever devices made by color scientists, so the trick socks are easily understandable; we're aware that the spine and brain are intimately connected and that people feel pains "in" phantom limbs. There's nothing bizarre about any of these one-off skeptical hypotheses—limited to just one pain or pair of socks. So it seems that we should accept these one-off skeptical arguments. And if we're going to accept them, thereby accepting that one could have come to have the belief that P in the perfectly ordinary, upstanding manner and yet later not know that P—where these Ps are exactly the targets of the live skeptical arguments—then why should we balk at the conclusions of the more encompassing live skeptical arguments? For we now have *precedents* for the live skeptical results: relatively uncontroversial cases in which you once knew, in the ordinary upstanding way, that P, and then due to expert contrary evidence you no longer know that P—where the P in question is targeted by the live skeptical arguments.¹⁰

6. Do We Really Need to Rule Out the Error Theories?

Now we move on to consider objections to the Live Skeptic's two principles. If we answer 'no' to this section's title question, then we are rejecting the Live Hypothesis principle, which says that we mere mortals do need to rule out the error theories in order to know things obviously inconsistent with them.

The Live Hypothesis principle is akin to but much more plausible than any of the standard epistemic closure principles. Compare what might be called Simple Closure with the Live Hypothesis Principle.

Simple Closure: If S knows that P entails $\neg H$, then if H isn't ruled out with respect to S, S doesn't know P.¹¹

¹⁰ The eliminativist case is harder. That is, it is difficult to find a one-off case in which you start out knowing that you believe P but then you *clearly* no longer know that you believe P because of a live, *familiar, realistic* hypothesis to the effect that you don't have that belief. The problem is in finding a case that satisfies the italics. I have an argument that there are such cases, but it's too long to include here (it appears in my 2005).

¹¹ Of course, this principle is usually stated a bit differently, but those differences won't matter. Minor objections to Simple Closure, e.g., the subject has to "think through" the connection between the entailment and her belief in P, can be avoided by building into the antecedent additional conditions. Of course, these can be applied to the Live Hypothesis principle as well.

The Live Hypothesis Principle: if S knows that P entails $\neg H$ and is as aware as just about anyone that H is live (so H really is live), then if H isn't ruled out with respect to S, S doesn't know P.

The two principles differ in two crucial ways: only the Live Hypothesis principle demands that H be live *and* that the subject be fully aware of H's live status. As we are about to see, by incorporating these conditions, the Live Hypothesis principle is immune to the plausible objections raised against Simple Closure.

A. Relevant Alternatives and Contextualism

It is sometimes said that one need not, in order to know P, rule out (or know the falsehood of) counterpossibility H provided H is appropriately "irrelevant."¹² Maybe so, but it is hard to see how this would apply if H were live in the manner described above, you were fully aware of the liveness, and you were knew full well of the inconsistency of H and P. If what I've argued above is wrong, then some hypotheses do not need to be ruled out for you in order for you to know P even though all the following conditions obtain:

- (i) Most people including yourself believe P;
- (ii) Everyone believes (correctly) the obvious fact that H is inconsistent with P;
- (iii) You've actually put together P and H and know as well as anyone that H is inconsistent with P;
- (iv) H is a real, live contender in your intellectual community;
- (v) You're aware that H is a real, live possibility actually endorsed by plenty of top experts;
- (vi) Lots of people worry about their attitude towards P in light of what they think about H; and
- (vii) Even the experts who think H is rubbish would pretty much laugh at any reasons you could give against H.

If just conditions (i)-(iii) hold, then perhaps I don't have to rule out H in order to know P.

Such a situation may obtain when P is 'I have hands' and H is some traditional skeptical hypothesis.

We *might* (I don't know) get the same answer if just the first four conditions hold. Perhaps one could

¹² I have in mind contextualist and relevant alternative theories, although what appears below addresses contextualism primarily. These theories are offered in several excellent works, e.g., Lewis (1996), Cohen (1988, 1999), DeRose (1995, 1999), Dretske (1970, 1981), Stine (1976), Goldman (1975), Heller (1999), and others. For reasons of space I don't address other, similar, theories such as various kinds of invariantism (e.g., Hawthorne 2004).

know P without ruling out H provided one believed P for the “right” reasons and was unaware that H was a live possibility endorsed by experts. But when all seven conditions hold, one can’t know P without being able to rule out H. If all seven conditions hold, then H needs to be ruled out as it is now a “relevant alternative” to what I believe; I am “epistemically responsible” to H. It seems to me that that is a reasonable constraint on relevant alternative theories.

Contextualists hold that different uses of ‘S knows P’ can have different truth-values even when concerning the same person, truth, and time. The idea is that depending on the context of utterance, different levels of epistemic standing for S’s belief in P are necessary for the truth of a use of ‘S knows P’ in that context. In some contexts, for instance when a lot is riding on whether P is true or S can be relied on, then S’s warrant has to meet a high standard to make the knowledge attribution true. In other circumstances, when for instance little rides on the truth of P, S’s warrant need meet only a low standard.

This theory-schema can be filled out in many ways depending on how one fills out ‘standards,’ what one says about which contextual factors can raise or lower standards, and what one says about how much those factors can raise or lower standards. There are many examples that support contextualism to a significant extent (e.g. DeRose 2005). Independently of that matter, contextualism is also thought to offer a plausible response to skepticism by making room for the claim that in discussions of skepticism the standards can, if the conversation proceeds just right, get raised absurdly high, so when the skeptic concludes with ‘So, we don’t know anything,’ she speaks the truth—even though in other contexts of utterance, governed by low standards, we speak the truth when we say ‘I know she hates my motorcycle.’ But can the theory provide a decent response to Live Skepticism?

Take as our example eliminativism. It’s natural to say that in theoretical contexts (e.g., discussions of connectionism or Fodor’s theory of content), eliminativism definitely poses a threat that can be neutralized only with some strong epistemic factors. The Live Skeptic seems right about knowledge attributions in those contexts. But what about the completely ordinary conversational contexts in which mere mortals discuss their alleged knowledge of a friend’s beliefs? Or the color of their socks? Or whether the doctor is aware of the new pain in their grandfather’s hip? Here the error theories are far from anyone’s mind. Sure, the liveness factors still exist, but they have been disarmed, so to speak, by contextual factors involving practicality concerns and the like. So the live eliminativist hypothesis doesn’t have to be ruled out in order for the mere mortal’s ‘I know she thinks motorcycles are dangerous’ to be true. Or so the objection to Live Skepticism goes.

For the sake of argument in the remainder of this subsection §6A I will assume that contextualism of some sort is true. The issue here is the truth-value of a *specific* kind of contextualism, one that implies that, for instance, ‘I know you think motorcycles are dangerous’ is true when the mere mortal conversational participants are not thinking or talking about live eliminativism (that’s the objection to Live Skepticism) even though when they discuss eliminativism

or other relevant theoretical topics appropriately then the very same knowledge sentence is false (that's the concession to Live Skepticism). I have three things to say about this kind of contextualism. Two appear here; the third is the subject of §9.

First point. It is plausible to think that the BIV hypothesis needs to be ruled out in order to save the truth of ordinary knowledge attributions only if it is being *discussed or thought about* appropriately. It is hard to see what else could make the BIV hypothesis a threat to one's belief that one has hands. However, that doesn't mean that occurrent discussion and thought are the *only* ways to make a hypothesis threatening. On the face of it, liveness conditions make a hypothesis just as or even more of a threat than occurrent attention. We already saw this point when briefly discussing relevant alternative theories. So we can't just say 'Well, we've already seen that skeptical hypotheses can be truth conditionally irrelevant in some contexts; so that's what must be happening with the error theories too.' Neither can we get away with 'Well, we know that if one is just minimally rational then one has neutralized the BIV hypothesis; so the same must be true for the color error theory hypothesis.' The liveness conditions (1)-(5) and mere mortality conditions (6) and (7) make the error theories truth-conditionally relevant as they are "highlighted" or contextually salient.

I don't deny that one can *formulate* a contextualism according to which if we're not talking or thinking about the error theory then it isn't highlighted and so presents no threat even if it is live and we are mere mortals. But this strikes me as an inferior kind of contextualism. Consider how such a theory would work with a different and less philosophically charged example. Suppose that every baseball book says, correctly, that Hank Aaron hit 755 home runs in his career. Sarah knows that that's what everyone thinks; and she agrees with them because she has read those books. Then Smith and Jones get archival evidence that suggests that in a 1958 game in which Aaron hit two home runs in the first four innings, the game was halted after four innings on account of rain. And the game was never continued. And the newspapers made a mistake the next day and reported the game as rained out after six innings. And a game is official if and only if it goes at least five innings. And because the newspapers said the game went six innings, the statisticians treated the game as official. And unless a game is official, what happens in it doesn't count towards one's statistics. (Most of this is almost true.) So if Smith and Jones are right, Aaron hit only 753 home runs. Suppose further that many people are very impressed by the evidence uncovered by Smith and Jones; the Smith-Jones hypothesis is live (but still underdog). Sarah hears about the new hypothesis and becomes aware of its respected standing. Still, she thinks it's probably false, although she is a mere mortal regarding it. So she still believes Aaron hit 755 home runs. Since Smith and Jones are wrong, her belief is true. Can you think of a context in this story in which 'Sarah knows that Aaron hit 755 home runs' is true even after Sarah becomes a mere mortal with respect to the live 753-homer hypothesis?

I can; but only if I give the sentence some non-literal interpretation. If it meant something better expressed by 'Setting aside the whole Smith-Jones thing, Sarah knows that Aaron hit 755 home runs,' then maybe it would express a truth. The 'Setting aside' sentence may be true, but I don't think

it conflicts with (the literal reading of) ‘Sarah does *not* know that Aaron hit 755 home runs,’ as it doesn’t even entail that Aaron hit 755 home runs. That is, ‘Setting aside the whole Smith-Jones thing, Sarah knows that Aaron hit 755 home runs’ doesn’t entail the truth of ‘Aaron hit 755 home runs.’ The ‘Setting aside’ sentence is more like ‘Sara knows that: assuming Smith and Jones are wrong, Aaron hit 755 homers,’ which obviously doesn’t entail that Aaron hit 755 homers. Even Smith or Jones could comfortably say ‘Setting aside our hypothesis, Sarah knows that Aaron hit 755 home runs.’ They would say it, believe it, and in saying it they would not be contradicting their belief that Aaron didn’t hit 755 home runs.

But other than taking ‘Sarah knows that Aaron hit 755 home runs’ to have some non-standard interpretation, I don’t see how it could be true unless she can rule out the 753-homer hypothesis, and the reason for this is that the latter hypothesis is live, Sarah is aware of that fact, and she is aware that it conflicts with her belief. If neither Fred nor his audience, both mere mortals, are thinking of the 753-homers hypothesis, he says ‘Sarah knows that Aaron hit 755 home runs,’ and he meant it literally, then what he says is false even if it’s conversationally appropriate. It simply doesn’t matter that the 753-homer hypothesis isn’t flashing on a screen over Fred’s head; whether Fred is thinking of or has recently mentioned the 753-homer hypothesis just doesn’t seem relevant.

I’m not beating up on contextualists here; this isn’t to say that contextualism is false. The Live Skeptic is free to go along with the contextualists when they claim that in some contexts ‘Sarah knows that Aaron hit 755 homers, assuming Smith and Jones are wrong’ and ‘Sarah knows that Aaron hit over 700 homers’ will be true while in others they are false due to various skeptical hypotheses becoming highlighted for standard contextualist reasons. I’m addressing only what I consider to be a relatively extreme kind of contextualism.

I’ve overstated the case slightly. I can see the attraction of allowing that a god or suitably wise 25th-century person would speak truly using ‘In the 21st century Sarah, like most people, knew that Aaron hit 755 homers.’ Perhaps they would speak truly because they know that the 753-homer idea was a false and short-lived aberration and from *their* god’s-eye perspective can be ignored. I doubt it, but in any case for anyone in Sarah’s community—especially the mere mortals—a use of ‘Sarah knows that Aaron hit 755 homers’ is false in every way.

Second point against this kind of contextualism. If the objection assumes that the *practical irrelevance* of the error theories is so extreme that they cannot really pose an epistemic threat large enough to significantly challenge one’s everyday, non-theoretical knowledge attributions, then the objection is dubious. We must remember the distinction between the *practicality* of the attribution-action (the action of saying ‘S knows P’) and the *truth* of the attribution-proposition (the proposition expressed by one’s utterance of ‘S knows P’). The objection can confuse the two. I fail to see why we should think that practicality concerns introduced by context affect truth conditions *to such an extreme extent* when it’s obvious that they affect in the first instance the conversational appropriateness of actions. In the previous section I admitted and even accounted for the

conversational and functional worth of attributions such as ‘I know that the table is brown and the cause of the pain in Jan’s knee after she ran into it; and I know that Jan believes all this too.’¹³

B. The Double Standard Objection

When I warmed us up with the dinosaur story I used scientific hypotheses for illustration: the supervolcano, solar flare, and meteor hypotheses. That was adequate to motivate the operative principles behind the Live Skeptic’s argument. But when I turned to squeezing some philosophical juice out of those principles, I plugged in four philosophical hypotheses, not scientific ones. And that’s the rub. We defer to *scientific* experts and liveness. But we don’t to philosophical ones. We should, the objector says, hold a *double standard*: the skeptical argument schema is correct for scientifically live hypotheses, but not for merely philosophically live ones. Philosophical speculation just doesn’t have the evidential weight of scientific opinion; that is the basis for the double standard. There may be no sharp distinction between philosophy and science, but there is a distinction nonetheless. Since the error theories aren’t scientifically live, we don’t need to rule them out.

However, cognitive scientists (in some possible worlds) consider eliminativism a live hypothesis. That hypothesis is live in the *scientific* community. The same holds for color error theory; in fact, it is clearly one of the dominant scientific views today in the color science community. The same holds for the pain location error theory.

So perhaps the Double Standard objection should conclude like this: the argument template is correct for hypotheses that are live due to scientific reasons, not those whose liveness is traced to philosophical reasons. But that won’t do either, because there are non-philosophical reasons for all four error theories.

Finally, is it really appropriate to have a double standard? If the Double Standard objection is sound, then philosophical reasoning can pose no threat to your beliefs. That is, even though you’re perfectly aware that many of your philosophical *superiors* claim that not-P, they have done so after years of reflection, and you couldn’t punch a significant hole in their considerations if your life depended on it, even if that’s all true, still their hypothesis that not-P poses no threat to your belief that P. Surely this says something interesting and highly negative about philosophy. If you advocate it, then why are you a philosopher anyway?

C. The Robustly Live Objection

¹³ I don’t mean to say that contextualists make this mistake; remember that in this subsection I’m assuming contextualism is true. Much more could be said about practicality and truth conditions and much that is worthwhile has been said by DeRose 1999, 2002, and 2005. Jessica Brown 2005 is a good rebuttal to DeRose 2002.

The kind of liveness that works in the dinosaur argument is *truth endorsement* on behalf of experts: the solar flare and supervolcano hypotheses are significant threats to my meteor belief because experts think that those alternative hypotheses are *really true*. But one might think that all we have in the case of our four error theories is some importantly weaker kind of endorsement that runs ‘Well, given the state of play in this area, this error theory is definitely one that we need to examine in order to further the field’. Let’s call a hypothesis *robustly live* if plenty of experts rate its probable truth more than 50% say. Then revise the Live Hypothesis principle so that it turns on robust liveness only. Now we see where the Live Skeptic has gone wrong: the four error theories are not robustly live. Call this the *Robustly Live* objection. The solar flare hypothesis is robustly live; eliminativism is not.

First problem: the objection has its facts wrong about color error theory. Actual experts really do think it’s true. Same for pain location error theory.

Second problem: some possible worlds in which the error theories are judged to have a 50% chance of truth by a large and diverse group of experts (based on several strands of evidence, et cetera; see the liveness conditions (1)-(5)) are very realistic. We saw this earlier when we discussed how the philosophy of mind could be permeated (if you like, infected) with the skeptical attitude towards common sense that one finds in many other central areas of philosophy.

7. Is it Not Actually Easy to Rule Them Out?

So perhaps we do need to rule out the live error theories. Can we mere mortals do it? If we can, then the Modesty principle is false, as it says we can’t.

A. The Effortless Ruling Out Objection

Maybe we can neutralize eliminativism without really doing anything. To see how, suppose philosophers and cognitive scientists proclaimed that there are no feelings at all. These eliminativists about feelings argue just as the eliminativists about belief argue:

No one has any beliefs/feelings. I realize that there are all sorts of cognitive/experiential processes or states in our heads and bodies, but none of them have what it takes to be a belief/feeling. The notion of a belief/feeling is a muddled folk notion that has been constructed in such a way that anything that is a belief/feeling must satisfy certain conditions. But in all probability nothing comes close to satisfying those conditions, which is why there are no beliefs/feelings.

I assume that eliminativism about feelings really is crazy. Even if all the philosophers of mind and cognitive scientists were screaming at my door ‘No one has ever had a painful feeling! There are no

feelings! No one has ever been in pain!' I would still *know* that I have had loads of painful feelings. My "access" to my painful feeling is so reliable or direct that I can epistemically neutralize, without even thinking about it, any hypothesis that says I have no feelings. I might be wrong about the location of the feeling, but there's no way in hell I'm wrong that I have a painful feeling.

Now it might be thought that we have given the anti-skeptic all she needs. Since I can know without any effort that the eliminativist about feeling is wrong, since I automatically have evidence (very broadly construed) that rules it out, similarly I can know without effort that the eliminativist about belief is wrong, since I automatically have evidence that rules it out. Perhaps this maneuver won't work for the other error theories, but it works for the key one, eliminativism.

I don't think so. For one thing, I simply can't imagine any remotely plausible reasons for thinking that no one is ever in pain. The no-pain hypothesis just could not be live in a rational philosophical community. One can of course marshal good reasons for thinking that a certain kind of "phenomenal pain" doesn't exist, but that is another matter entirely.

More to the point, belief is just more theoretically loaded than feeling. For belief isn't some subjective state like feeling. Perhaps a painful feeling just is something that seems painful; but the analogous claim for belief doesn't work. Sure, there is a feeling of subjective certainty towards a claim; to that extent belief can involve something akin to sensation. I can know, no matter what the experts say, that I have a feeling of certainty that is directed towards a claim. But there is more to belief than that. Beliefs aren't subjective feelings, of certainty or anything else. They might include such feelings, but they are more than that. At the least, it is pretty unreasonable to reject this view about belief going beyond feeling. In fact, I don't anyone who rejects that view. So even though I do have "subjective evidence" that neutralizes the hypothesis that I don't have any feelings of certainty or at least approval of a claim, I don't have subjective evidence that rules out eliminativism about belief. Similar points hold for pain locations, character traits, and colors. For those reasons I don't think the anti-skeptic can reasonably hold that we all automatically, or upon easy reflection, possess warrant that is, on balance, sufficient to neutralize the error theories. We possess positive warrant that when left alone does the trick, but the liveness and mere mortality conditions generate enough negative warrant to sabotage our knowledge.¹⁴

B. The Externalism Objection

¹⁴ A pseudo-Moorean might want to argue thus: I believe that $2 + 2 = 4$; thus, eliminativism is false. Since I know the premise and the argument's validity, I have a "proof" that eliminativism is false but no such proof that I believe that $2 + 2 = 4$. But of course the Live Skeptic denies that we know the first premise to be true. The objector (who is a mere mortal in an eliminativism-live world) could then try this argument: I know that believe that $2 + 2 = 4$; thus, Live Skepticism for belief is false. However, there is good evidence that Moore would not argue this way; and even if he did, his argument would fail (see my 2005).

Maybe the processes by which our second-order, first person beliefs (your beliefs about your own beliefs) are formed are incredibly reliable. The processes by which we judge colors are pretty reliable too; the same for pain locations. One might think that these reliability facts provide enough warrant for our true beliefs about colors, pain locations, and beliefs to amount to knowledge.

I'll grant that these reliability facts suffice to place a significant amount of positive epistemic warrant on our second-order, first-person beliefs. Even so, only an exceedingly strong form of reliabilism can back up a neutralization of the eliminativist hypothesis. No reliabilist I know of has endorsed such a version of reliabilism. Everyone has admitted that factors like the recognition of expert counter opinion, liveness, independent contrary evidence, and whatnot can defeat a reliably formed belief—even highly reliable ones. We saw some instances of this in previous sections (George's belief about the senator and, especially, your beliefs about the color of your trick socks and the pain in your spine).

Another problem with the reliability objection is that if the degree of reliability of the belief *producing* process has any epistemic weight (as the reliabilist reasonably holds), then the degree of reliability of the belief *sustaining* process also has that weight. The belief sustaining process has to be reliable to *retain* knowledge, and in the cases described in this essay—liveness, mere mortality—to continue to have beliefs that are known to be in contradiction with hypotheses one appreciates as live is clearly *highly* unreliable. That is, the following belief-sustaining procedure—which applies to the beliefs in question—is unreliable: continue to believe that P even though a contrary hypothesis Q has become live and one is well aware of the inconsistency and liveness. The reliabilist might be right that the reliability of the belief forming process is often enough to make a true belief knowledge; so as children those beliefs amount to knowledge primarily in virtue of reliability facts. Our question has to do with the *retention* of that knowledge upon coming aware of contrary expert evidence. The reliabilist who wants to use reliability factors to defeat the live skeptical arguments has to claim not only that the belief-producing reliability confers a significant amount of warrant, but also, mysteriously, belief-sustaining reliability is irrelevant. I don't see any good reason to swallow this extremism.

The generality problem for reliabilism is relevant here, but I won't take the space to address it properly. I will make just one comment. One cannot object to my argument of the previous paragraph by claiming that the mere mortal isn't aware that she is doing something that falls under the belief-sustaining process description 'continue to believe that P even though a contrary hypothesis Q has become live and one is well aware of the inconsistency and liveness'. Since she is a mere mortal she is aware of the live error theory, its liveness, and its inconsistency with her threatened belief. To that extent, the belief-sustaining method is particularly relevant to the epistemic status of the belief.

8. The Internalism Objection

One might think that the Live Skeptic is implicitly assuming externalism. After all, consider a world just like the eliminativism-live world except that eliminativism isn't live. You exist in both worlds and possess the very same evidence for your belief that Moore believed that skepticism is false: in both worlds you have read some of what he wrote about skepticism, you have read what Barry Stroud wrote about Moore in his 1984 book, etc. You are intrinsically identical in the two worlds, the only difference being in how favorably philosophers of mind and cognitive scientists view eliminativism. Clearly, in the world in which eliminativism is not even remotely live—no Stich, no Churchlands, no Rorty, no Feyerabend, et cetera—you are free to know that Moore believed skepticism to be false. The Live Skeptic has allowed that knowledge, since there is no live hypothesis ruining your belief. Well, if we're internalists, then since you are intrinsically physically identical across worlds, your true belief is warranted enough for knowledge in both worlds. Thus, even in the eliminativism-live world you know that Moore believed skepticism to be false; so Live Skepticism is false.

We haven't put our finger on just *where* the Live Skeptic went wrong, so we can't classify the internalist response as targeting the Live Hypothesis principle or rather the Modesty principle, but we know there has to be a mistake somewhere.

This is confused, and matters won't improve if we fiddle with definitions of internalism. Take another look at the conditions on being a mere mortal, (6)-(9) given in §2. If eliminativism hasn't even been thought of in your possible world, then it's hard to see what those conditions amount to. For one thing, (8) says that you are *quite aware* that eliminativism is live and you are familiar with H and the issues surrounding H, including some of the evidence for H. But how could that happen to you in the world in which eliminativism isn't even on the radar screen in philosophy or cognitive science? If eliminativism has never even crossed anyone's mind, then (8) and (9) are odd to say the least.

Let's back up. The internalist wants an objection that runs like this.

- A. In the eliminativism-dead world (the world in which eliminativism hasn't crossed virtually anyone's mind) the subject has internal access to all the same considerations (as to what Moore believed) as the mere mortal in the eliminativism-live world. She has read the same literature about Moore, etc.
- B. In the eliminativism-dead world she knows Moore believed that skepticism is false.
- C. Internalism is true. And it entails that if the warrant she has internal access to in the first world is sufficient for knowledge, then since she has access to all and only that warrant in the eliminativism-live world, she has the same knowledge in that other world.
- D. So, she knows she's a believer in the eliminativism-live world.

In order to make (A) true the subject needs to be just like us, roughly put, vis-à-vis eliminativism. Pluck a mere mortal from an eliminativism-live world and set her down in the eliminativism-dead world with exactly the same attitudes and philosophical abilities; that's what it takes to make (A) true. Thus, the truth of (B) depends on the answer to the following question.

Suppose you *sincerely think* that (1)-(5) are all true of your community of philosophers of mind and cognitive scientists; you know that your belief P (e.g., that Moore believed skepticism to be false) is inconsistent with eliminativism; your intelligence, understanding, and knowledge are not extraordinary for people in *our* eliminativism-live intellectual community with regard to eliminativism or the issues surrounding eliminativism; and if you have any reasons or evidence you can marshal for casting doubt on eliminativism, and if they were carefully considered by the members of *our* eliminativism-live community—in particular, by the well-informed, well-respected, non-crackpot, and highly intelligent experts who are thoroughly familiar with eliminativism—they would be nearly universally and confidently rejected as clearly insufficient to rule out eliminativism. (That's what it takes to make (A) true.) If that's all true of you, then do you *know* P?

I don't see any reason to answer affirmatively, even if internalism is true. I admit that someone utterly unacquainted with eliminativism could know many truths inconsistent with eliminativism; that issue was explored in detail above. But (B) in the internalist's argument is just plain dubious. So the internalist objection fails. And frankly, even if it succeeded we would still have the interesting conclusion that on externalist theories Live Skepticism is true.

9. Watering Down the Conclusion

Suppose that despite everything I just argued, some anti-skeptical solution for ordinary, everyday contexts is right even though *perhaps* in theoretical contexts the Live Skeptic wins. So in ordinary-life contexts 'I know my socks are blue' is true. It doesn't matter why the Live Skeptic is wrong; let's just assume she is wrong. I will make three comments in response.

I have given some pretty good arguments that in theoretical contexts—ones analytic philosophers and cognitive scientists find themselves in virtually every day—the new skeptical arguments go through. I find that skeptical result still pretty amazing. If it's right, then you don't know much of anything *right now* about the color of objects around you, anyone's beliefs, anyone character traits, or the pains in your knees. And you don't know any of that in many, many philosophical and even scientific contexts. In spite of the recent endorsement of skeptic-friendly versions of contextualism, according to which we know very little if we're discussing a classic skeptical hypothesis and argument in the right way (thereby making those hypotheses mount real

epistemic threats, despite being *nowhere near* live), I still find it almost unbelievable that we don't know, *right now*, simple facts about the colors of our socks, our aches and pains, or what we believe.

For the sake of argument pretend that you're the mere mortal. You continue to believe P even in the recognition that it conflicts with the live hypothesis H. You're aware that if P is true, then H is false; you've actually thought this through. And you have reflected on the fact that you believe P. So you believe not-H, as you can see as well as anyone that P's truth straightforwardly rules out H—where H is the real, live hypothesis endorsed by experts, etc. If what I've argued above is wrong, so you *know* that P, you *know* that if P, then not-H, and you've thought this all through properly, then it seems that you know not-H, especially after you go through this reasoning. *So you know that the live hypothesis is false—even though there are loads of experts better positioned than you who don't, you're no genius regarding P or H, you are as aware as anyone that these experts disagree with you and are real experts, and no expert would take seriously your reasons against H.* This seems faintly ridiculous.

Last point. Suppose once again that some anti-skeptical solution is right; so Live Skepticism is defeated. Then perhaps the anti-skeptic may have won a battle but lost the war. Knowledge attributions are often true (that is the won battle), but the truth conditions of those attributions are so meager that we should feel a little ashamed (that is the lost war). The *persistent* Live Skeptic will say that even if the relevant ordinary knowledge attributions are true, skepticism wins because the truth conditions for those knowledge claims are impoverished. It should be clear that the persistent skeptic's position is not this: although 'S knows P' is true, S's knowledge is not "high standards" knowledge. *This* skeptic isn't complaining that our knowledge doesn't fulfill some super-duper high-octane condition that only a philosopher could love. The persistent Live Skeptic who objects that even if the relevant knowledge attributions are often true the knowledge states would exist but be pathetic isn't whining that knowledge isn't what she always fantasized it should be. She has admitted that we can have all the positive warrant we thought we had; she can even admit that we are often have super-duper high-octane knowledge of contingent matters of fact. She is targeting a special set of beliefs, ones targeted by the live error theories, and claiming that the warrant possessed by these beliefs is anemic *compared* to that had by our other beliefs that amount to knowledge. *There are seven reasons—(1) through (7)—for thinking that I need some powerful epistemic factors to defuse the expertly endorsed, highly respected contrary scientific error theories plus two reasons—(8) and (9)—to think I have no such factors; and yet you are telling me that I can know lots of things inconsistent with those error theories?* This cannot be a victory the anti-skeptic should celebrate. If this realism about knowledge is accurate, then it's a gaunt realism and the Live Skeptic has been closer to the truth than the realist.¹⁵

¹⁵ Thanks are due to the University of Leeds for granting me a semester of research leave to work on a boring project that I abandoned in favor of doing the research that led to this essay. Thanks also to a *Noûs* referee, Paul Bloomfield, David Chalmers, Richard Fumerton, Sanford Goldberg, John Greco,

References:

- Brown, Jessica. (2005) "Adapt or Die: the Death of Invariantism?" *Philosophical Quarterly*.
- Byrne, Alex and David Hilbert. (2003) "Color Realism and Color Science," *Behavioral and Brain Sciences*, 26: 3-21.
- Churchland, Paul M. (1989) *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*, Cambridge, MA: MIT Press.
- Churchland, Patricia S. (1986) *Neurophilosophy: Toward a Unified Science of the Mind-Brain*, Cambridge, MA: MIT Press.
- Cohen, Stewart. (1988) "How to Be a Fallibilist," in *Philosophical Perspectives*, 2, Tomberlin, ed., 91-123. Atascadero, CA: Ridgeview Publishing Company.
- Cohen, Stewart. (1999) "Contextualism, Skepticism, and the Structure of Reasons," in *Philosophical Perspectives*, 13, Tomberlin, ed., 57-89. Oxford: Blackwell.
- Conee, Earl. (2001) "Heeding Misleading Evidence," *Philosophical Studies*, 103: 99-120.
- Dennett, Daniel. (1978) *Brainstorms*, Montgomery, VT: MIT Press.
- DeRose, Keith. (1992) "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, 52: 913-929.
- DeRose, Keith. (1995) "Solving the Skeptical Problem," *Philosophical Review*, 104: 1-52.
- DeRose, Keith. (1999) "Contextualism: An Explanation and Defence," in *The Blackwell Guide to Epistemology*, Greco and Sosa, eds. Oxford: Blackwell.
- DeRose, Keith. (2002) "Assertion, Knowledge, and Context," *Philosophical Review*, 111: 167-203.
- DeRose, Keith. (2005) "The Ordinary Language Basis for Contextualism and the New Invariantism," *Philosophical Quarterly*.
- Doris, John. (1998) "Persons, Situations and Virtue Ethics," *Noûs*, 32: 504-30.
- Doris, John. (2002) *Lack of Character*, Cambridge: Cambridge University Press.
- Dretske, Fred. (1970) "Epistemic Operators," *Journal of Philosophy*, 69: 1007-23.
- Dretske, Fred. (1981) "The Pragmatic Dimension of Knowledge," *Philosophical Studies*, 40: 363-378.
- Feyerabend, Paul. (1963) "Materialism and the Mind-Body Problem," *Review of Metaphysics*, 17: 49-66.
- Frances, Bryan. (2005) *Skepticism Made Live*, Oxford: Oxford University Press.
- Frances, Bryan. (ms) "Universal Skepticism, Live Hypotheses, and Ruling Out Principles."

Ken Himma, Sarah McGrath, Joseph Melia, Peter Millican, Mark Nelson, Duncan Pritchard, and Tom Stoneham for written or oral comments, and the faculties at the University of Liverpool and University of Connecticut for helpful discussion. Special thanks are due to my generous friend Andrew McGonigal, with whom I discussed the issues raised in this essay on many fruitful occasions, and my wife Margaret Frances for offering substantive comments on the entire essay.

- Ginet, Carl. (1980) "Knowing Less by Knowing More," in *Midwest Studies in Philosophy Studies in Epistemology*, V, French, Uehling Jr., and Wettstein, eds. Minneapolis: University of Minnesota Press.
- Goldman, Alvin. (1975) "Discrimination and Perceptual Knowledge," *Journal of Philosophy*, 73: 771-791.
- Harman, Gilbert. (1980) "Reasoning and Evidence One Does Not Possess," in *Midwest Studies in Studies in Epistemology*, V, French, Uehling Jr., and Wettstein, eds. Minneapolis: University of Minnesota Press.
- Harman, Gilbert. (1999) "Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error," *Proceedings of the Aristotelian Society*, 99: 315-331.
- Harman, Gilbert. (2000) "The Nonexistence of Character Traits," *Proceedings of the Aristotelian Society*, 100: 223-6.
- Hawthorne, John. (2004) *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Heller, Mark. (1999) "The Proper Role for Contextualism in an Anti-Luck Epistemology," in *Philosophical Perspectives*, 13, Tomberlin, ed., 115-129. Oxford: Blackwell.
- Hyman, John. (2003) "Pains and Places," *Philosophy*, 78, 5-24.
- Lewis, David. (1996) "Elusive Knowledge," *Australasian Journal of Philosophy*, 74: 549-67.
- Lycan, William. (1977) "Evidence One Does Not Possess," *Australasian Journal of Philosophy*, 55: 114-126.
- Moore, George Edward. (1962) *Philosophical Papers*, New York: Collier Books.
- Quine, Willard.V.O. (1960) *Word and Object*, Cambridge, MA: MIT Press.
- Quine, Willard.V.O. (1985) "States of Mind," *Journal of Philosophy*, 82: 5-8.
- Rorty, Richard. (1970) "In Defence of Eliminative Materialism," *Review of Metaphysics*, 24: 112-121.
- Stich, Stephen. (1983) *From Folk Psychology to Cognitive Science: The Case Against Belief*, Cambridge, MA: MIT Press.
- Stich, Stephen. (1996) *Deconstructing the Mind*, New York, NY: Oxford University Press.
- Stine, Gail. (1976) "Skepticism, Relevant Alternatives, and Deductive Closure," *Philosophical Studies*, 29: 249-261.