

# Live Skeptical Hypotheses

Revised version forthcoming in John Greco, ed., *Oxford Handbook of Skepticism*

Bryan Frances

Fordham University

## 1. Skeptical Threats Real and Merely Academic

Those of us who take skepticism seriously typically have two relevant beliefs: (a) it's plausible (even if false) that in order to know that I have hands I have to be able to epistemically neutralize, to some significant degree, some skeptical hypotheses, such as the brain-in-a-vat (BIV) one; and (b) it's also plausible (even if false) that I can't so neutralize those hypotheses. There is no reason for us to also think (c) that the BIV hypothesis, for instance, is *plausible or probably true*. In order to take skepticism seriously it's sufficient to hold (a) and (b); one need not hold (c). Indeed, philosophers who accept (a) and (b) *never* endorse (c). Show me a philosopher who suspects that he is a brain in a vat and I'll show you someone who is deranged!

That's one thing that bothers undergraduates in philosophy. They object: why on earth do some philosophers take the BIV hypothesis to pose *any* threat at all to our beliefs given that those very same philosophers think that there's *no* real chance that the BIV hypothesis is true? Sure, the BIV hypothesis is *formally inconsistent* with my belief that I have hands, so if the former is true then my belief is false. But so what? Why should that bare inconsistency matter so much? Is this strange attitude amongst philosophers the result of some logic fetish infecting the philosophical community? It is sometimes said that the skeptical hypotheses are not only inconsistent with our beliefs but are *explanatory* of our experiences, which is supposed to make them more of a threat. But students aren't fooled: although the skeptical hypotheses may attempt to explain why our experience is as it is, it's the kind of attempt appropriate for science fiction movies that are all special effects and virtually no plot. No one with any sense of reality will take the evil demon hypothesis to be even tenuously explanatory.

The students would understand the fuss over the BIV hypothesis if there were some decent reason to think that the BIV hypothesis was really true. If you believe P, a contrary hypothesis Q has some reasonably good backing—perhaps endorsement by legitimate experts in the relevant field—and you are quite familiar with Q's good status as well as the conflict between P and Q, then the Q possibility *does* seem to mount a threat to one's belief in P, a threat that if left unneutralized does ruin one's chance at knowledge of P's truth. If the BIV or evil demon hypotheses were like Q, then we would have a real threat to our belief that we have hands.

Here is an example, one that I'll appeal to later in the essay as well. When Jo was a teenager she learned the theory that a huge meteorite wiped out the dinosaurs. She learned this theory in the usual way, hearing it from her parents and teachers, and reading about it in 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; lucky Jo! 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 studied the dinosaurs until she graduated with a Bachelor's degree in paleontology she was perfectly aware that there were two "live" contrary hypotheses, ones 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 these alternative hypotheses as well as their liveness and how they conflict with the meteor hypothesis, then as long as the competing hypotheses weren't ruled out 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. Perhaps she doesn't have to completely *demolish* the supervolcano and solar flare hypotheses in order to know that the meteor story is true, but she certainly has to knock them down a few epistemic notches.

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. There aren't many ideas in philosophy that can't be coherently and cleverly challenged, but it is a brave philosopher who thinks that the above story is misguided. Surely, one would think, in *those* circumstances Jo's true and partially justified belief in the meteor hypothesis didn't amount to knowledge.

Let's return to our students' skeptical attitude towards skepticism. If we could find some skeptical hypotheses that were false but *live*, in something like the way the supervolcano and solar flare hypotheses were, then we would think that they pose a significant threat to our beliefs that are obviously inconsistent with them: if we couldn't epistemically neutralize those skeptical hypotheses (whatever neutralization amounts to), then our beliefs obviously inconsistent with those false but live hypotheses would fall short of knowledge (even if they were true beliefs). At least, this lack of knowledge would follow if the people in question were well *aware* of the live status of the skeptical hypotheses as well as the inconsistency between their beliefs and the live skeptical hypotheses. For in that case we would have the threat to our beliefs staring us right in the face as it were. The threat would be out in the open. Our students would understand our skeptical worries better.

Well, *are* there any such skeptical hypotheses—live hypotheses inconsistent with huge portions of our commonsensical beliefs? It's easy to adjust the BIV story so that it becomes live in another possible world. BIVs become technologically and fiscally possible; a team of mad scientists ends up ruling the world; certain top secret documents detailing plans for massive kidnapping and BIVing are discovered and leaked to the public; impenetrable 100,000 square mile complexes of laboratories are set up in the wilds of Canada, Brazil, and elsewhere; thousands of people start mysteriously disappearing in the middle of the night to become BIVed.... Well, you can see how at least one BIV hypothesis—'I have become a BIV as the result of these crazy mad scientists'—could become something close to a real, live possibility for most any person (not the mad scientists themselves, provided we tell the story right). Even if you're never BIVed in that world, does your true belief that you have a bicycle in your garage amount to knowledge?

The students won't be satisfied with this skepticism, not fully anyway. For we have yet to find any real skeptical threat. No *actual* version of skepticism is on the horizon. None of our students' actual beliefs are threatened by any skeptical hypothesis. Skepticism remains merely academic and highly dubious. They're still skeptical about skepticism.

However, there are skeptical hypotheses that are actually live; others are at least very nearly actually live. In this essay I will list some of these hypotheses and show why they do generate a "real"

skeptical threat, one quite different from the usual, merely academic, threats. The skeptical theory that results I call *Live Skepticism*.<sup>1</sup>

## 2. The Live Skeptic's Argument

The Live Skeptic uses an argument just like the dinosaur one: 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 H.

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). (Thus, this condition (7) entails (1)-(5).) In fact, you are as aware as just about anyone of the fact that H is live.

---

<sup>1</sup> I have introduced and evaluated Live Skepticism in more detail elsewhere (2005a, 2005b).

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 X’ in (7) means that X is true and S has all the usual good warrant for believing X, warrant that usually suffices for knowledge. I can’t say that it entails *knowledge* of X, 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 X, she also has some “negative” warrant that offsets or vetoes that positive warrant enough so that her belief in X no longer is, on balance, warranted or sufficient for knowledge.

The conjunction of (6) and (7) (which as noted in (7) includes by entailment (1)-(5)) says, very roughly, ‘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 roughly, ‘You do not have what it takes to defuse the threat.’

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 implausible. 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 virtually 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 well-informed 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 Sceptic now makes three key claims.

*The Modesty Principle:* if S is A mere mortal with respect to live hypothesis H (so conditions (1)-(9) hold), 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 an extremely realistic, very close possible world many members of our actual intellectual community of contemporary analytic philosophers, cognitive scientists, and their advanced students 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, e.g., Patricia Churchland (1986), Paul Churchland (1989), Stephen Stich (1983), Daniel Dennett (1978), Paul Feyerabend (1963), Richard Rorty (1970), and Quine (1960, 1985)).<sup>2</sup>

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).

---

<sup>2</sup> 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.

- Color Error Theory: no ordinary physical objects are colored (endorsed by, e.g., 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 from these three claims 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 believe that they are (utterly, completely) *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.

### 3. The Liveness-Mortality Premise

Consider belief error theory, or eliminativism as it's commonly called. Eliminativism is currently in our intellectual community (of contemporary analytic philosophers, cognitive scientists, and their advanced students) a live possibility. The hypothesis is a going-concern among the experts. That is just a brute fact about the community of philosophers of mind and cognitive scientists. 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 and other areas of philosophy 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 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 most of her students and colleagues 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.

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 can marshal for casting doubt on 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 and arguing against 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—(9) in particular—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, briefly 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. 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.

The Live Skeptic needs a weak 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 nearly universal common sense could easily permeate the philosophy of mind while eliminativism had more expert supporters and fewer detractors.<sup>3</sup> 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. 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 some objections to the Live Skeptic's argument. I can't consider more than a small fraction of all the reasonable objections (see my 2005a and 2005b for more thorough treatment)! So, your favorite might not get investigated here. And with regard to the few I do consider I don't treat them as thoroughly as I should. But I hope to provide some initial reason for doubting that they succeed.

The objections will fall into two classes. First, there are reasons for thinking that we don't have to be able 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 several ways, that to a certain extent I lose the *motivation* to find fault with the Live Skeptic's argument. I will list three key marks of Live Skepticism.

---

<sup>3</sup> 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 even come close to meeting 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, none of which I can go into here. Suffice it to say that being a live hypothesis involves meeting a very demanding set of conditions. In part, I stick with hypotheses that have scientific backing because we defer to scientists more than we defer to philosophers, and this difference seems wise.

First, the Live Skeptic is claiming that we lose our knowledge by means of a perfectly familiar mechanism.<sup>4</sup> 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', they will claim that what you say is fine in some respects but 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-is-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 assuming that 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.

---

<sup>4</sup> 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.

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 interesting feature of Live Skepticism. 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 in science or philosophy. 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 and procedures 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, 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.

The Live Skeptic's crucial claim is this: *the overall warrant possessed by the beliefs targeted by the live skeptical hypotheses is anemic compared to that had by our other beliefs that amount to knowledge and aren't targeted by any live skeptical hypotheses.*

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. 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).

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.

Now for the third interesting feature of Live Skepticism. 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?

## 5. 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.<sup>5</sup>

*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.

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."<sup>6</sup> 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 knew full well of the inconsistency of H and P. If what I've argued above is wrong, then some

---

<sup>5</sup> Of course, this principle is usually stated 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.

<sup>6</sup> I have in mind contextualist and relevant alternative theories, although what appears below addresses contextualism primarily. These theories are discussed in essays 2 and 3 of part 3 of this volume.

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 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 some extent. 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. 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 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'll include just one argument regarding this kind of contextualism.

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.

## **6. Is it Not Actually Easy to Rule Them Out?**

So perhaps we do need to be able 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.

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 most interesting 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 community.<sup>7</sup>

More to the point, the notion of belief is just more theoretically loaded than that of 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

---

<sup>7</sup> One can of course marshal good reasons for thinking that a certain kind of "phenomenal pain" doesn't exist, where a theoretically-loaded reference-fixing definition of 'phenomenal pain' is in play, but that is another matter entirely.

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 is sufficient for knowledge of facts about beliefs, but the liveness and mere mortality conditions generate enough negative warrant to sabotage that knowledge.

Maybe reliabilism will suffice to doom Live Skepticism. Suppose that reliability facts suffice to place a significant amount of positive epistemic warrant on our second-order, first person beliefs. Perhaps they supply enough warrant that in ordinary circumstances, when we aren't in a society fallen to serious consideration of eliminativism, true second-order, first-person beliefs are often good enough for knowledge. The Live Skeptic can agree with all of that. Perhaps they supply so much warrant that even in an eliminativist-live society people can have second-order knowledge if they are quite divorced from and ignorant of both philosophy and cognitive science. Even so, eliminativism is actually live, and I'm aware of and it and its status. Suppose I happen to reject the eliminativist hypothesis, and so believe that I have beliefs. Still, I'm no immortal regarding eliminativism. It seems that there is now a significant amount of *negative* epistemic warrant placed on my second-order, first person beliefs (this is the negative warranted revealed by the eliminativist skeptical argument). And there is enough of it to render the positive reliability-based warrant insufficient to render my second-order beliefs warranted enough for knowledge. Only an exceedingly strong form of reliabilism can back up a neutralization of the eliminativist hypothesis. 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.

A relevant point here is that if the degree of reliability of the belief *producing* process has significant epistemic weight (as the reliabilist reasonably holds), then the degree of reliability of the belief *sustaining* process also has significant epistemic weight. This is no mystery. 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 the recognizably live hypotheses 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 a mere mortal with respect to Q. 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. 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. Reliability of belief formation gives me 1218 positive warrant units; standing fast with the belief that P in face of appreciated, well-respected, expert evidence produces just 200 negative warrant units; I needed just 1000 units to have my true belief turned into knowledge. Alternatively: belief-forming reliability gives me positive warrant of sufficient quality for knowledge and belief-sustaining unreliability amounts to no veto. Should we swallow this extremism? I certainly don't feel any temptation to when it comes to philosophically uninteresting examples, so I don't see why we should when it comes to philosophically interesting examples.

## 7. 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. I will make two comments in response.

I have given some pretty good arguments that in theoretical contexts—ones analytic philosophers and cognitive scientists often find themselves in—the new skeptical arguments go through. I find that skeptical result 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.

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 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 making one crucial claim: *the overall warrant possessed by the beliefs targeted by the live skeptical hypotheses is anemic compared to that had by our other beliefs that amount to knowledge and aren't targeted by any live skeptical hypotheses*. 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.

#### References:

- 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.
- Dennett, Daniel. (1978) *Brainstorms*, Montgomery, VT: MIT Press.
- Doris, John. (1998) "Persons, Situations and Virtue Ethics," *Noûs*, 32: 504-30.
- Doris, John. (2002) *Lack of Character*, Cambridge: Cambridge University Press.
- Feyerabend, Paul. (1963) "Materialism and the Mind-Body Problem," *Review of Metaphysics*, 17: 49-66.
- Frances, Bryan. (2005a) *Scepticism Comes Alive*, Oxford: Oxford University Press.
- Frances, Bryan (2005b) "When a Skeptical Hypothesis is Live," *Noûs* 39 (2005), 559-95.
- 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.

Hyman, John. (2003) "Pains and Places," *Philosophy*, 78, 5-24.

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.