

The Impossibility of Skepticism

Forthcoming in *The Philosophical Review*

Not final draft: please cite published version

Epistemologists and philosophers of mind both ask questions about belief. Epistemologists ask normative questions about belief—which beliefs ought we have? Philosophers of mind ask metaphysical questions about belief—what are beliefs, and what does it take to have them? While these issues might seem independent of one another, there is potential for an interesting sort of conflict—the epistemologist might think we ought to have beliefs that, according to the philosopher of mind, it is impossible to have.

In §1, I argue that this conflict does arise, and that it creates problems for traditional skeptical views in epistemology. In particular, I will argue that on certain popular views about the nature of belief, it is impossible to adopt the near-global agnosticism recommended by the skeptical epistemologist.¹ On other plausible views, it is only possible in special circumstances, and this limitation undermines skeptical epistemological claims. The only views about the nature of belief on which there are no metaphysical hurdles to adopting the agnosticism

¹I refer to *near* global skepticism because the sort of traditional skeptical epistemologist I have in mind does not aim to undermine introspective or *a priori* knowledge. This restriction only makes things harder for me; as should become clear, if my arguments entail that *near* global skepticism is impossible, they also entail that *global* skepticism is impossible.

recommended by the skeptic are views that face powerful objections—objections that are completely independent of anti-skeptical epistemological considerations.

This might not seem to be any conflict at all—perhaps the beliefs that we ought to have are ones that, unfortunately, it is not possible to have. However, in §2 I defend an epistemic “ought implies can” principle, according to which the metaphysical impossibility of being a skeptic implies that it is not the case that one ought to be a skeptic. With this principle in place, my arguments from §1 constitute a novel objection to traditional skeptical epistemological views—novel in part because it does not depend on any controversial positive epistemological claims about which beliefs we ought to have.

1 Skepticism

A certain sort of skeptical argument proceeds roughly as follows. It starts with the observation that various bizarre possibilities—possibilities in which the earth was created five minutes ago, or in which the laws of nature will radically change tomorrow, or in which we are handless brains in vats—are in some sense similar to the possibilities we take to be actual. The skeptical epistemologist might claim that these bizarre possibilities are ones in which everything appears the same to us as it actually does, or that these possibilities are phenomenally indistinguishable from the actual one. She goes on to claim that when various possibilities are alike in this sense, you ought to be totally agnostic about which one is actual—for none of them should you be more confident that it is actual than any other.²

It is not completely obvious how we should understand the state the skeptical

²Vogel (2004) discusses some arguments along the above lines.

epistemologist thinks we ought to be in.³ If we work in a qualitative framework and talk in terms of belief, disbelief, and suspension of judgment, we will say that she is recommending near-global suspension of judgment. According to the skeptical epistemologist, the only matters about which we are licensed in having beliefs are the contents of our own minds, perhaps along with things like logic and mathematics—on everything else we must suspend judgment. Things are trickier if we try to model the recommended state quantitatively—the question of how to best represent agnosticism in a quantitative framework is a controversial one and I will take it up only briefly in this paper.⁴ I will try to remain as neutral as possible between different ways of representing neutral doxastic states; my arguments will, I hope, be translatable into any attractive framework for thinking about agnosticism.

In the first part of this paper, I will ask whether it is possible to be a skeptic, and if so what it would take, on three different accounts of the nature of doxastic attitudes. On the first account, which I will call *the direct pragmatic picture* following Robert Stalnaker (1984), I will argue that it is impossible to be a skeptic. On the second account, *the indirect pragmatic picture*, being a skeptic will turn out to be possible, but only in special empirical circumstances. This leads to the result that one cannot consistently accept the indirect pragmatic picture, be a skeptic, and believe that one is a skeptic—if one believes that one is a skeptic

³Terminological note: I will call the person who endorses an argument like the one sketched in the above paragraph the *skeptical epistemologist*. I will call the person who is in the doxastic state recommended by the skeptical epistemologist the *skeptic*.

⁴Joyce (2005) argues for representing doxastic neutrality using sets of probability functions, a strategy that White (2010) criticizes. Norton (2008) argues against standard Bayesian methods for representing ignorance, as well as methods using sets of probability functions, and defends an alternative approach to representing ignorance. I will briefly discuss this issue in §1.1.1.

and believes the indirect pragmatic picture, one must also believe some specific empirical hypotheses, belief in which is inconsistent with being a skeptic. On the final account I will consider, the *phenomenological picture*, being a skeptic is easy. However, there are strong objections to this picture that have nothing to do with its allowing the possibility of skepticism.

First, however, I would like to deflect a potential criticism. One might worry that I've set up the skeptical position in a non-standard way; standard skeptical arguments purport to establish conclusions about what we *know*,⁵ while the skeptical epistemologist I discuss purports to establish conclusions about what we ought to believe. I doubt that this distinction will diminish the interest of my discussion, however. Plausibly, there are close connections between what we know, and what we ought to believe; in particular, becoming convinced that we lack knowledge of a given domain should move us to agnosticism about matters in that domain.⁶ Moreover, philosophers who have argued for skeptical theses concerning knowledge have typically tried to show that their case extends to justified belief.⁷

I don't claim that there is *no* interesting way of raising skeptical challenges that exclusively concern knowledge, and which do not have implications for what

⁵See, e.g., DeRose (1995).

⁶Some writers have defended views according to which the concept of a knower is closely connected with the concept of an informant on whom we ought to rely in forming our beliefs. See, e.g., Craig (1990) and Gibbard (2003, chap. 11). I suspect such views will have an easy time explaining connections between what we know and what we ought to believe. Other ways of securing the connection might involve appealing to the Williamsonian thesis that one ought to believe only what one knows (Williamson, 2000).

⁷Unger (1975) argues first for skepticism about knowledge, but then argues that his case extends to skepticism about reasonability and justification as well.

we ought to believe.⁸ I do claim that the skeptical position I am considering is not of merely parochial interest due to its focusing on claims about what we ought to believe, nor should the discussion of such a position be thought of as irrelevant to the central questions about skepticism.

1.1 The Direct Pragmatic Picture

In chapter 1 of *Inquiry*, Robert Stalnaker sketches and defends a way of thinking about a certain class of mental states:

Here is one impressionistic picture of human activities which involve mental representation—call it the pragmatic picture. Rational creatures are essentially agents. Representational mental states should be understood primarily in terms of the role that they play in the characterization and explanation of action... One explains why an agent tends to act in the way he does in terms of beliefs and [pro and con] attitudes. And, according to this picture, our conception of belief and of attitudes pro and con are conceptions of states which explain why a rational agent does what he does. Some representational mental states—for example, idle wishes, passive hopes, and theoretical beliefs—may be connected only very indirectly with action, but all must be explained, according to the pragmatic picture, in terms of their connections with the explanation of rational action. (Stalnaker, 1984, p.4)

⁸Though see Burnyeat (1983, pp. 118-19) for the view that contemporary discussions of skepticism tend to focus too much on knowledge, and not enough on belief.

Other philosophers who have endorsed similar positions about the nature of representational mental states include Daniel Dennett (1987) and David Lewis (1974). I will postpone discussion of why one might be attracted to this picture until §1.3 when I discuss the phenomenological picture—one of the main motivations for accepting either the direct or the indirect pragmatic picture is that they seem to be capable of solving problems that, on the phenomenological picture, look impossible to solve.

My strategy in this section will be to identify two principles that I take to be motivated by the idea that there's a close constitutive link between belief and action, and which will define what I'll call the "direct" pragmatic picture, and to argue that these principles together rule out the possibility of being a skeptic. A forewarning—insofar as these principles seem to the reader to go beyond the picture sketched in the quote above, I suspect it is because the reader associates the quote with the indirect pragmatic picture, rather than the direct pragmatic picture. As the quote says, the picture is impressionistic, and depending on how it's fleshed out, the connection between belief and action may be more or less tight. That is fine—the reader should take the following two principles to be definitive of the direct pragmatic picture, and should wait for the next section for discussion of the indirect pragmatic picture.

The two principles I will use to argue against the possibility of skepticism are the *weak rationalization principle* (WRP), and the *strong rationalization principle* (SRP).

If S has a set of doxastic attitudes B , then

WRP: B rationalizes some of S 's behavioral dispositions together with

S 's desires.⁹

SRP: There is no proper subset $B' \subset B$ such that B' rationalizes S 's behavioral dispositions just as well as B , together with S 's desires.

First a terminological note: I take an action ϕ to be rationalized by a set of beliefs and desires just in case the set of beliefs and desires rules out some alternative action ψ available to the agent as irrational, while not ruling out ϕ as irrational. So if I have normal beliefs and desires, and I am offered \$100 from box A , or \$100 from box B , or \$0 from box C , taking the contents of box A would be rationalized by my beliefs and desires, as would taking the contents of box B . Taking the contents of box C would not be rationalized. However, if I believed that every option open to me in some situation had the same payoff—e.g., if I would get \$100 no matter what I did—then my beliefs and desires would not rationalize any action in that situation. I take it that the motivation for WRP is straightforward,

⁹According to Strawson (1994), there could be creatures with beliefs and desires, but without dispositions to engage in any outward, observable behavior. He imagines a race of weather watchers with beliefs and desires about the weather, but no dispositions to act in ways that would satisfy these desires. While (at least as I read him) he grants that a set of beliefs and desires must rationalize an agent's behavior in order for the set to truly be the agent's beliefs and desires, he holds that some of this behavior may be irreducibly mental—the weather watchers' beliefs and desires may explain when and why they're happy or sad, afraid or relaxed, without explaining when and why they perform bodily actions (since they never act, and nor are they disposed to). While this is a significant departure from some versions of the direct pragmatic picture, I think it is still enough for my argument to get off the ground. Roughly, the same considerations that (as I argue later in this section) show that no goal-directed *actions* could be rationalized by a skeptic's beliefs and desires, also show that no mental episodes of happiness, sadness, fear or relaxation could be rationalized by a skeptic's beliefs and desires.

if we are inclined to accept anything in the neighborhood of the direct pragmatic picture. But why accept SRP? An advocate of the direct pragmatic picture might offer the following analogy to motivate the acceptance of SRP.

In chemistry textbooks, atoms are described as having valences, or valency numbers, which measure the number of bonds they can form. For example, an atom with a valence of 1 is missing an electron, and can form a bond with another atom with a valence of -1 , since the latter would have an electron to spare. The entire theoretical role of valency numbers is to explain bonding behavior along the above lines. However, we might imagine a chemist who proposed a revisionary theory of valence. According to this chemist, valency numbers do not only have real, integer-valued parts. They also have complex parts. He might hold, e.g., that the valency number of hydrogen is not 1, as in the standard theory, but $1 + 3i$. However, on his theory, the complex part of an atom's valence would play no role in explaining its bonding behavior—according to his theory, an atom with valence $1 + 3i$ would interact with other atoms just as an atom with valence 1 would according to the standard theory.

I take it as obvious that the chemist imagined above would simply be confused about what work the theory of valences is supposed to do—there is no possible scenario in which atoms' bonding behavior is just like it is, but their valences have complex-valued parts in addition to real, integer-valued parts. The complex parts would be an explanatory fifth wheel, and since valences *just are* whatever plays a certain explanatory role, a true theory of valences could not include explanatorily irrelevant parts along the lines of the one proposed by our hypothetical chemist. The defender of the direct pragmatic picture sees the relationship between doxastic attitudes and intentional action along the same lines as the relationship between

valences and bonding behavior. If there are two hypotheses about which doxastic attitudes an agent has, and they both rationalize the agent's behavioral dispositions equally well, and they differ only in that one ascribes the agent extra doxastic attitudes that the other does not, then these extra attitudes are explanatorily irrelevant to the agent's behavior. If the direct pragmatic picture is correct, then an agent could not actually have beliefs that play no role in explaining his behavior—the more complicated of the two theories could not be correct. But this is just SRP.¹⁰

Now that we have explained why someone sympathetic to the quote at the beginning of this section might accept WRP and SRP, we can ask how these principles bear on the possibility of skepticism. In broad outline, my argumentative strategy in this section runs as follows:

Premise 1: If the direct pragmatic picture is correct, then if a subject has a set of doxastic attitudes B , those attitudes rationalize the subject's behavior.

Premise 2: No body of behavior could be rationalized by the near-global agnosticism recommended by the skeptical epistemologist.

¹⁰I take it that the above considerations do not mean that a defender of the direct pragmatic picture has to think that beliefs *supervene* on behavior. Even if we accept the direct pragmatic picture and the above argument for SRP, unless we're comfortable with *extreme* indeterminacy in which beliefs an agent has, we will need to posit extra, tiebreaker principles governing belief and desire. For familiar reasons, there will always be multiple, distinct sets of beliefs and desires that do equally well at explaining the agent's behavior, each of which includes only attitudes that do play a necessary role in explaining the agent's behavior. Stalnaker (1984) appeals to causal considerations to resolve indeterminacies like these, and Lewis (1983b) appeals to considerations of naturalness.

Therefore:

Conclusion: If the direct pragmatic picture is correct, no subject could have the attitudes recommended by the skeptical epistemologist.

The notion of rationalization used in the argument is rationalization in accordance with WRP and SRP. Premise 1 is true by the definition of the direct pragmatic picture, so the challenge is in establishing premise 2. While I will offer a detailed defense of premise 2 below, the basic motivation behind it is easy to state. The following is plausible as a first pass characterization of what's required for rational action: in order for an agent's attitudes to rationalize an action, she must both have some desires, and believe that in acting in a particular way she will be likely to satisfy those desires.¹¹ But if she believes that if she acts a certain way then she will satisfy her desires (or at least will be more likely to do so than if she acts in other ways), she is not agnostic about everything the skeptical epistemologist says she must be. She has some views about what outcomes her actions are likely to lead to, so she is already too opinionated to count as a skeptic.

While the above considerations give the general flavor of the reasons to think that the direct pragmatic picture rules out the possibility of skepticism, in the rest of this section, I will refine the argument for their incompatibility. Readers who are willing to take my word that the direct pragmatic picture rules out the possibility of skepticism may skip to §1.1.3.

We do not quite yet have the materials necessary to mount a careful argument against the possibility of skepticism on the direct pragmatic picture—first, we

¹¹I call this a “first pass” characterization in order to make room for the idea that ultimately, it is credences and utilities, rather than beliefs and desires, that rationalize action.

will need to draw a distinction between two types of desire, and their roles in rationalizing actions. Call an agent's desire *normal* if its satisfaction conditions do not depend on how things seem to the agent at the moment he or she acts on it. Suppose I desire not to be rained on in the afternoon, and I act on this desire by picking up an umbrella as I leave my home in the morning. My desire not to be rained on in the afternoon in this case is a normal desire—whether it is satisfied depends not on how things seem to me at the moment I pick up my umbrella, but instead on whether I get rained on in the afternoon.

Distinguish normal desires from *capricious* desires. An agent's desire is capricious if its satisfaction conditions depend only on how things seem to him or her at the moment he or she acts on it. Perhaps the desire to picture something red *right now* is a capricious desire—I can satisfy it by constructing a mental image of something red at the very moment that I have the desire, no matter what things are like with respect to matters that outstrip my present appearances.¹² There are some reasons to be dubious of capricious desires—perhaps action and thought must take time, and so it is somehow illegitimate to speak of the moment at which I act on a desire—but admitting the possibility of such desires only makes things more difficult for me, so I will set aside worries like this.

¹²This example is actually controversial because of issues involving semantic externalism—perhaps my desire could not be satisfied if I had never had any (even indirect) causal interaction with the color red, and so was incapable of thinking about it (Putnam, 1981). We still might say that, in an attenuated sense, my desire counts as a capricious desire. Once its content is fixed—whether it's about red, schmed, or whatever—whether it is satisfied will depend only on my thoughts, and not on any further external world conditions being satisfied.

An astute reader will notice that the normal/capricious distinction is not exhaustive. For instance, one might have desires that depend on coordination between present appearances and facts that outstrip these appearances. For instance, I might desire to have the exact same sensation now that I had when I first tasted chocolate, or I might have the conditional desire to feel miserable now if and only if I committed a sin in the past. Call these *hybrid* desires.

In the remainder of this section I will try to establish the following claim: there is no body of behavior such that a) it might be the behavior of some agent S , b) S is a skeptic, and c) WRP and SRP are true.

Suppose for reductio that S is a skeptic. By WRP, some of S 's actions must be rationalized by her beliefs, together with her desires. Could normal desires help rationalize any of her actions? I will argue that they could not. Take my normal desire not to get rained on in the afternoon. What would my beliefs have to be like in order for this desire to rationalize the action of taking my umbrella in the morning? Clearly, I would have to believe something like the following: I am less likely to get rained on in the afternoon if I take my umbrella than if I do not. However, holding such a belief is incompatible with being a skeptic—whether taking my umbrella will make me less likely to get rained on in the afternoon is exactly the sort of matter on which skeptics are agnostic.¹³

¹³Clearly a skeptic can't have outright beliefs about whether she will be rained on if she takes her umbrella, but can she believe that it is *less likely* that she will get rained on if she takes her umbrella? I don't see any principled way of using the sort of skeptical arguments I mentioned at the beginning of this paper that leaves room for this position. Such arguments imply that *any* kind of doxastic discrimination between competing hypotheses that are compatible with one's experiences, including taking some such hypotheses to be more likely than others, is unjustified.

This is not a special feature of the example. In general, for a desire ϕ to help rationalize an action ψ , it must do so in the presence of something like a belief that ψ is more likely than alternative actions to lead to the satisfaction of ϕ .¹⁴ But when ϕ is a normal desire, claims about which actions will lead to the satisfaction of ϕ are claims that do not depend on how things currently appear, nor are they plausibly *a priori*. So, according to the skeptic, we must be neutral among all such claims. But if we are neutral among these claims, then our normal desire cannot rationalize any action—it cannot combine with our doxastic attitudes to make any action look more attractive than any other.¹⁵ Similar considerations show that hybrid desires could not rationalize any actions for a skeptic. Take the example of the desire to experience the exact same sensation *now* that I experienced when I first tasted chocolate. Even if I could will myself to feel any experience I liked, without some non-skepticism about matters that outstrip present appearances (e.g., whether my apparent memory is reliable, what I experienced in the past, etc.), this desire will not rationalize any action.

¹⁴What might the analogue of this principle look like in a quantitative framework? I suggest that instead of thinking of beliefs and desires together rationalizing actions, we would talk about features of an agent's probability function rationalizing actions together with features of the agent's utility function. The principle in the text would turn into the principle that assigning some proposition O a high utility could only rationalize an action ψ over some alternative action γ in the presence of a probability function P with the feature $P(O \mid \psi) > P(O \mid \gamma)$.

¹⁵In the context of offering a different sort of argument against skepticism in his dissertation, Jacob Ross (2006) comes to a similar conclusion. Very roughly, he argues that whatever the evidence that might seem to mandate taking agnostic *doxastic* attitudes of the sort the skeptic recommends, we can rationally *act* as if we believe in common sense, since the practical impotence of the skeptical position means that we can safely ignore it in deciding what to do.

We supposed for reductio that S was a skeptic. By WRP we concluded that his doxastic attitudes, together with his desires, must rationalize some of his actions. We then concluded that neither normal or hybrid desires could rationalize any of his actions together with his doxastic attitudes. So the only remaining possibility is that some of his actions are rationalized, but only by capricious desires. Call an agent who only ever acts on capricious desires a *capricious agent*. I will now appeal to SRP to rule out the possibility that a capricious agent might be a skeptic.

If we assume that actions, when rationalized, are rationalized always by a combination of a desire-like attitude and a belief-like attitude, then there must be some doxastic attitudes that come into the picture when capricious desires rationalize actions. But what might such doxastic attitudes be? Perhaps if a capricious desire to see something red now rationalizes my action of picturing something red now, then it must do so in conjunction with the belief that picturing something red now is a way of seeing something red now. What seems clear, however, is that doxastic attitudes to propositions whose truth is contingent on matters that outstrip present appearances will *not* help rationalize our skeptic's actions. While our skeptic may be completely agnostic about empirical questions such as whether there have ever been any black dogs, her agnostic attitudes will play no role in rationalizing any of her actions. But SRP rules this out—the actions of an agent who only ever acts on capricious desires could be just as well explained by attributing to her no doxastic states at all (rather than neutral ones) concerning matters that outstrip appearances—neutral ones do not help explain any of her behavior. A few examples will make this clear.

Imagine that Sue bets on various political elections. However, there is one election in which not only does Sue not bet, but when given the offer of specifying

a bet which she would take, she's unable to answer. That is, for each candidate in the race she is asked what odds it would take for her to be willing to bet on that candidate, and after much time spent pacing back and forth, she explains that she cannot come up with a response that she's willing to stick to. In this case, interpreting Sue as agnostic about which candidate will win the election seems to best explain her behavior. She clearly understands elections, and betting, and bookies—interpreting her as having a typical set of concepts and beliefs rationalizes most of her actions. Furthermore, she has thought about which candidate will win this particular race—it is not as if she has no doxastic attitudes at all towards this subject matter. Moreover, there is no non-neutral doxastic attitude towards the question of who will win that explains her actions; only an agnostic one will do.

Contrast Sue's relation to the proposition that some particular candidate will win with that of a typical dog, Fido. Fido digs for bones, chases cats, and barks at strangers. If we accept the direct pragmatic picture, we should allow that Fido has beliefs about matters such as where his bone is, and whether his master is a stranger. But we should not allow that Fido has beliefs about politics. It is not that he is agnostic about which candidate will win—he does not have some neutral doxastic attitude towards this subject matter, the way that Sue does. Rather, he has no intentional relations at all to any propositions about which candidate will win. By appealing to SRP, we can explain this—a body of doxastic attitudes which included agnosticism towards political matters would explain Fido's behavior no better than a restricted set of attitudes that included no attitudes at all towards propositions about politics.

The relationship between Fido and subject matters like politics is like the relationship between a capricious agent and all subject matters that outstrip its

present appearances. If we accept the spirit of the direct pragmatic picture, we will not think that Fido has any representational mental states—concepts, beliefs, whatever—that play no role in explaining his behavior. This will rule out representational mental states that concern matters like politics. By the same line of thinking, a capricious agent will lack representational mental states that concern matters other than its present appearances. If the direct pragmatic picture is true, a capricious agent is not *agnostic* about the external world, or the past, or the future—rather, it is a conceptually impoverished creature that has no attitudes towards these subject matters at all.¹⁶

Our reductio against the possibility of skepticism, assuming the direct pragmatic picture, is complete. We assumed *S* was a skeptic. We used SRP to infer that some of *S*'s actions must be rationalized by her beliefs and desires. We concluded that no normal or hybrid desires could rationalize any of her actions, so she must be a capricious agent. We then used SRP to argue that a capricious agent is not a skeptic. So *S* is a skeptic, and *S* is not a skeptic—we have our contradiction. Before exploring just how strong an anti-skeptical result the direct pragmatic picture provides, I'll consider two objections to the arguments I've offered so far.

1.1.1 Objection 1: Agnosticism as Equiprobability

A skeptical epistemologist might object that I've mischaracterized the doxastic state she recommends. She might hold that being agnostic concerning a set of

¹⁶I'm assuming that insofar as one is attracted to the direct pragmatic picture of belief, one will also be sympathetic to a directly pragmatic account of concept possession—one will hold, perhaps, that what concepts a creature has is somehow a matter of which features of the world it is capable of making its behavior depend on.

hypotheses involves regarding each of the hypotheses as equally probable. This idea is closely related to the “principle of indifference”—the claim that when one has no more reason to believe one hypothesis than another, one should regard those hypotheses as equally probable. If the skeptical epistemologist can characterize the skeptical doxastic state as one obtained by following an indifference principle, then perhaps there needn’t be any conflict between the direct pragmatic picture and the possibility of skepticism. At least in simple cases, treating each of a set of hypotheses as equally probable typically *will* rationalize certain actions. For instance, if I think that each ticket in a large lottery is equally likely to win, I’ll be happy to take certain bets concerning the outcome—for any particular ticket, I’ll be willing to bet against that ticket’s winning at quite steep odds. If we can understand the skeptical state as one in which the subject has opinions that can be represented by a precise probability distribution that assigns each member of some set of hypotheses equal probability, rather than as a state of some more extreme form of doxastic neutrality, then perhaps the skeptical state *will* turn out to rationalize some actions over others.^{17,18}

There are a number of problems with the proposal that we reconcile skepticism with the direct pragmatic picture by understanding agnosticism in terms of equiprobability. First, even in the simple cases in which it’s clear *how* to apply the principle of indifference (e.g., lottery cases), it isn’t obviously well motivated.

¹⁷Thanks to an anonymous referee for raising this point.

¹⁸Clarificatory note: I am not claiming that probability distributions can only rationalize actions if they are precise. Rather, I am claiming that if they are precise, then they’re well placed to rationalize actions, since they will entail that actions have precise expected utilities (at least when combined with precise utility functions); this is a sufficiency claim, rather than a necessity claim.

If I know that an n -ticket lottery is fair, then I should certainly think that each ticket is equally likely to win: for each ticket t I should have a degree of belief of $\frac{1}{n}$ that t will win. But suppose I don't know that the lottery is fair—suppose that my only information about the lottery is that one of the tickets will win. In this case, having a precise degree of belief of $\frac{1}{n}$ concerning each ticket seems to involve acting as if I know more than I really do; it will involve being quite confident, for each ticket t , that t will lose. But it's hard to see what could justify that confidence, especially if we require justification to meet the skeptical epistemologist's standards. Treating agnosticism concerning which ticket will win as involving an assignment of equal probability to each ticket's winning seems to gloss over an important distinction; there's a crucial difference between believing that a lottery is fair—this clearly *does* involve equiprobability—and being completely agnostic about whether or not it is fair.

Some hold that considerations like the ones above provide conclusive reasons to reject the principle of indifference, and the picture of agnosticism as equiprobability along with it.¹⁹ Even if we're not convinced, however, serious problems remain. In slightly more complicated cases, naïve applications of the principle of indifference lead to contradiction.²⁰ When there are multiple ways of partitioning a space of possibilities, it's not possible to assign equal probability to each member of each cell of each partition. There are sophisticated attempts to save the spirit of principles of indifference from contradiction—the maximum entropy approach advocated by E.T. Jaynes (1968, 1973) being the most notable such attempt—and a discussion of such attempts would be well beyond the scope of the present paper. Luckily,

¹⁹See Joyce (2005, 2010).

²⁰See van Fraassen (1989, chap. 12).

it is unnecessary. While it is highly controversial whether strategies like Jaynes' provide satisfactory treatments of the cases they are designed to handle, it should be uncontroversial that even if such strategies are successful, they don't generalize in the ways they would need to in order to serve the skeptical epistemologist's purposes.²¹

I conclude that the skeptical epistemologist cannot reconcile her position with the direct pragmatic picture by appeal to the idea that the appropriately skeptical state is one that can be represented by a precise probability function obeying some principle of indifference. In the next subsection I'll consider a response on behalf of the skeptical epistemologist that appeals to the distinction between belief and acceptance.

1.1.2 Objection 2: Belief and Acceptance

Many philosophers have drawn distinctions between belief and acceptance.²² The skeptical epistemologist might appeal to this distinction and protest that while the direct pragmatic picture is attractive as an account of *acceptance*, it fails as

²¹Jaynes' treatment of Bertrand's paradox relies crucially on physical symmetries in the setup of the problem, and on the idea that the correct answer should be invariant with respect to translations along these dimensions of symmetry. No analogous strategy is applicable in the more general case that the skeptical epistemologist needs to deal with.

²²See, e.g., van Fraassen (1980), Stalnaker (1984), Harman (1986), Cohen (1989), Bratman (1992), and Wright (2004) for a sampling, though not all (in fact, probably not any two) of these authors draw the distinction in exactly the same way. Wright's discussion is of particular relevance, as it concerns the relation between acceptance and skepticism. His way of drawing the belief/acceptance distinction is different from the one I focus on in this section, but like the one I discuss, it is also uncongenial to the skeptical epistemologist.

an account of belief. She might then concede that my arguments so far show that a neutral state of *acceptance* is impossible, but insist that it is merely a neutral state of *belief* that she advocates.²³

Some ways of drawing the distinction between belief and acceptance will allow for this escape route. As I read L.J. Cohen (1989), he essentially adopts a version of the direct pragmatic picture concerning acceptance and a version of the phenomenological picture—a view I’ll discuss in §1.3—concerning belief. As I argue in §1.3, the phenomenological picture of belief *does* straightforwardly allow for the possibility of skepticism (though I argue that it has other flaws), so Cohen’s account of the belief/acceptance distinction could provide comfort to the skeptical epistemologist. In this subsection, however, I’ll argue that merely drawing a belief/acceptance distinction isn’t enough to save the skeptical epistemologist. On at least one attractive way of drawing the distinction between belief and acceptance, the doxastic state that the skeptical epistemologist recommends will still turn out to be impossible.

On the way of drawing the belief/acceptance distinction I prefer, acceptance is the more general notion.²⁴ To accept that P for purposes ϕ is to treat P as true for purposes ϕ . For instance, if one accepts for the sake of argument that the monarchy should be abolished (perhaps in a discussion about how best to achieve this end), then one will argue as if the monarchy should be abolished: one won’t object when one’s conversational partners presuppose that the monarchy should be abolished, one will assert claims that one takes to follow from the undesirability of the monarchy, and so on. Along similar lines, if one accepts some scientific

²³Thanks to an anonymous referee for raising this point.

²⁴My discussion of acceptance follows that of Stalnaker (1984, chap. 5).

theory as a working hypothesis, one will design one's experiments in ways that would yield the most insight if the theory really were true, and one will ignore alternative theories when conducting one's research.

Belief, on this picture, is global (or near-global) acceptance—acceptance for all (or at least a very wide range of) purposes. If one accepts that the monarchy should be abolished for the sake of an argument, but does not *believe* that the monarchy should be abolished, then there must be *other* purposes for which one does *not* accept that the monarchy should be abolished. For example, one will not vote in favor of a government that will aim to eliminate the monarchy. If one accepts a theory as a working hypothesis, but does not *believe* the theory, then there must be some purpose for which one does *not* accept that the theory is true—perhaps one designs one's *own* experiments as if the theory is true, but one also supports *other* people's doing research that is based on alternative, incompatible theories.

Part of what makes this picture attractive is that it makes intelligible the close relationship between belief and acceptance, while also explaining some of the key differences between belief and mere acceptance. On this picture, because belief is a limiting case of acceptance, belief and acceptance aren't treated as two wholly unrelated propositional attitudes, like fear and amusement, or hope and regret. But we can also use this picture to explain differences between belief and acceptance, such as why having inconsistent beliefs is a problem, but accepting inconsistent propositions can be fine. The explanation is that inconsistent beliefs will lead to self-defeating behavior, but one can accept inconsistent propositions without sabotaging oneself, so long as one accepts them for different purposes.²⁵

²⁵Of course, this picture faces its own challenges. In particular, it must explain what's involved

If this is how we think about the relationship between belief and acceptance, what should we make of the skeptical epistemologist's suggestion that we should suspend *belief* concerning empirical propositions, while nevertheless *accepting* such propositions? We won't be able to make much sense of it. Given the direct pragmatic picture, it's hard to see how one could even count as accepting that P for purposes ϕ , without holding some empirical *beliefs* concerning the conduciveness of accepting that P to the furtherance of one's aims—if the direct pragmatic picture is right, then there doesn't seem to be room for a distinction between accepting that P for purposes ϕ and believing that accepting that P for purposes ϕ is the best course of action available. It's hard to see how one could have a systematic pattern of acceptances without also thereby counting—according to the direct pragmatic picture—as having beliefs concerning the potential fruitfulness of various policies of acceptance. But such means-end beliefs are incompatible with skepticism, for reasons already discussed.

Ultimately, we can agree that there is a distinction between belief and acceptance without thereby undermining the argument for the incompatibility of skepticism with the direct pragmatic picture. Before moving on to discussion of the indirect pragmatic picture, however, I want to address the question of which forms of skepticism are impossible if the direct pragmatic picture is true, and to in believing that P , while nevertheless accepting that $\sim P$ for at least *some* purposes, as this is clearly possible. For instance, I might believe that Jones is the woman for the job, but accept for the sake of argument that she's not (in a discussion with someone who disagrees with me, perhaps). I think examples like this can be handled without giving up the idea of belief as a limiting case of acceptance, but it takes some finesse, and is beyond the scope of the present paper.

distinguish the impossibility result we get if we adopt the direct pragmatic picture from some similar anti-skeptical positions in the literature.

1.1.3 The Extent of the Anti-Skeptical Result

In this subsection I'll address two related questions. First, which forms of skepticism are impossible if the direct pragmatic picture is correct? Second, how might an anti-skeptical argument built on the direct pragmatic picture (together with an epistemic ought implies can principle, which I will defend in §2) differ from similar anti-skeptical arguments that have been defended by Hilary Putnam (1981)?

We saw in the previous section that local agnosticism (like Sue's agnosticism about who will win an election) is possible if the direct pragmatic picture is true; it is for this reason that the direct pragmatic picture provides no response to various local skeptical arguments.

Some forms of local skepticism are of relatively little epistemological interest. For instance, consider a skeptical epistemologist who holds that while one can rationally dismiss the hypothesis that the world was created five minutes ago, or that one is being deceived by an evil demon, or that the laws of nature are about to change, one ought to be agnostic between a commonsense view of the world and a particular skeptical hypothesis. Let's say it's a brain-in-vat hypothesis according to which while one is a brain in a vat, the software the vat is running ensures that life will appear to go on as it usually does. This sort of agnosticism isn't ruled out by the direct pragmatic picture; certain actions (e.g., eating delicious cakes) will still be rationalized by the attitudes of this sort of local skeptic, since those actions would satisfy one's desires for gustatory pleasure regardless of whether the commonsense hypothesis or the vat hypothesis were true.

While this sort of skepticism isn't ruled out by the direct pragmatic picture, I take it that it's not a very epistemologically interesting form of skepticism. It's hard to imagine any attractive argument that would motivate doxastic neutrality between common sense and the particular brain-in-vat hypothesis just discussed that wouldn't motivate a much more general skepticism. The most compelling arguments for agnosticism concerning whether or not we're brains in vats, in my mind, are ones along the lines sketched at the beginning of this paper—arguments that appeal to quite general principles that require agnosticism between *all* phenomenally indistinguishable hypotheses. If there are plausible skeptical arguments that would motivate agnosticism concerning some particular brain-in-vat hypothesis without motivating agnosticism concerning the other sorts of skeptical hypotheses I've been discussing, I'm not aware of them, and they're not often discussed in the epistemological literature.

There are other forms of local skepticism, however, that *are* epistemologically interesting, and which don't seem to be ruled out by the direct pragmatic picture. Agnosticism about unobservable entities (van Fraassen, 1980), or moral facts (Sinnott-Armstrong, 2006), or other minds, all seem to be possible even if the direct pragmatic picture is true—for each of these local skepticisms, it seems at least *prima facie* that there *are* bodies of behavior that can be rationalized by various desires together with doxastic states that include skepticism about such matters.²⁶ If this is right, then the direct pragmatic picture allows that such skep-

²⁶In fact, none of these cases are straightforward. To take just the case of agnosticism about unobservable entities, Melchert (1985), Horwich (1991), and others have argued that if *accepting* a theory about such entities has the same behavioral consequences as *believing* it, then the distinction van Fraassen tries to draw between accepting such theories and believing them is

ticisms are possible, and provides no method for responding to skeptical arguments that attempt to establish that they are normatively required.

This isn't to say, however, that the only epistemologically interesting forms of skepticism ruled out by the direct pragmatic picture are *completely* global forms of skepticism. Hilary Putnam (1981) famously argued that, contrary to the skeptical epistemologist, considerations from the theory of reference establish that we are not brains in vats (and so we shouldn't be agnostic about whether we are brains in vats). Imagine a reformed skeptical epistemologist who accepts Putnam's arguments, and recommends that we be agnostic only between all scenarios that are consistent with broadly Putnamian views about reference. So while we may rule out the possibility that we've always been brains in vats, we must treat hypotheses according to which we have been recently envatted as on all fours with more commonsensical ones.²⁷ As many writers have pointed out, even if we rule out hypotheses according to which we have never had any causal contact with vats, if we have to be agnostic between all the hypotheses that remain, we're still left with an extremely strong form of skepticism—this is one of the reasons why Putnam's arguments are often not taken to provide us with as much as we'd like in responding to the skeptic.

But it is far from clear that even the fallback attitude recommended by the reformed skeptical epistemologist is possible, if the direct pragmatic picture is true. Suppose you have set for yourself the goal of climbing Mount Everest. Given

 untenable; it may turn out to be impossible to both accept a theory about unobservables, while remaining agnostic about its literal truth. Such an argument will likely be attractive to those sympathetic to the direct pragmatic picture.

²⁷I don't know where this response to Putnam on behalf of the skeptic was first offered, but see Brueckner (1986) for one example.

a commonsense view of the world, it is obvious (at least in broad strokes) what sorts of actions would be appropriate ways of trying to achieve this goal—it would make sense to practice by climbing smaller mountains, to spend a lot of time in high-altitude, low-oxygen environments to prepare your lungs for the climb, etc.

But suppose, instead of holding a commonsense view of the world, you are completely doxastically neutral between all hypotheses consistent with causal theories of reference. Now it is far less clear which actions, if any, will seem like appropriate means to try to achieve your goal. After all, you might be a recently-envatted brain in a world where Everest doesn't exist (perhaps the world will have to contain some mountains, or mountain-like entities, but it certainly needn't contain Everest). Or alternatively, you might be in a world where Everest exists and there's a chain of reference leading from it to you, but the laws of nature are about to change in radical ways such that the actions that in the past would have been reliable means of ensuring that one eventually reached the summit of Everest will in the future have disastrous consequences. If you are completely neutral between *all* such hypotheses, no actions will look better than any others as means to reaching the summit of Everest. And of course, the desire to reach the summit of Everest isn't unusual in this respect—we can run the same argument *mutatis mutandis* for any other normal desire one might have.

Just as the direct pragmatic picture is inconsistent with the possibility of the near global skepticism discussed in §1.1, it is also inconsistent with the possibility of the more subtle, reconstructed skepticism that accommodates Putnamian arguments about reference; according to the direct pragmatic picture, since there is no body of behavior that could be rationalized by the doxastic attitudes recommended by the reconstructed Putnamian skeptical epistemologist, such attitudes

are impossible to have. Given an ought implies can principle, the direct pragmatic picture implies not only that it's not the case that we ought to be skeptics, but also that it's not the case that we ought to be reconstructed Putnamian skeptics.

There are also less concessive responses to Putnam's anti-skeptical argument that don't threaten the anti-skeptical result obtained on the direct pragmatic picture. One might hold a version of the view that some mental contents are *narrow*, in that whether we hold them does not depend at all on our external environment.²⁸ A defender of narrow contents might dismiss Putnam's argument by holding that the skeptical epistemologist ought to be understood as making recommendations about what our narrow doxastic attitudes ought to be. Even if it is a fact that the mental state of (wide) belief that one is a brain in a vat couldn't turn out to be true, on this response, that fact is irrelevant to the central and interesting questions concerning skepticism.

Of course, it's controversial whether there is any viable notion of narrow mental content. But at least some influential ways of developing the idea are consistent with the direct pragmatic picture and still lead to anti-skeptical results. In "Beyond Belief," Dennett (1987) argues that a version of the direct pragmatic picture of belief can ground not only attributions of wide beliefs to agents, but also attributions of narrow beliefs. How might it apply in a particular case? Suppose some agent *S* has the following dispositions—if she were in an environment containing a glass of water, she would drink the water, if she were in an environment containing a glass of XYZ, she would drink the XYZ, and more generally, if she were in an environment containing a glass of clear, odorless liquid indistinguishable from water, she would drink it. If the rest of her dispositions are as we might expect,

²⁸See Brown (2002) for a survey.

we might use the direct pragmatic picture to conclude that she desires to quench her thirst and (narrowly) believes that having a drink of water would satisfy this desire.

There are strong reasons for skepticism that a strategy like the one Dennett pursues really could provide the foundation for a general theory of narrow content.²⁹ But even granting that such a strategy can succeed, we can see quite quickly that it won't provide any comfort to the skeptical epistemologist. Near-global agnosticism, whether it is understood as a narrow or broad mental state, will not rationalize any actions (whether they are characterized narrowly or broadly), and so on the direct pragmatic picture, is not a possible mental state. The direct pragmatic picture, then, provides a much more robust anti-skeptical argument than Putnam's, since it doesn't require rejecting the possibility of narrow content.

The lesson of this section is that the direct pragmatic picture is powerfully anti-skeptical, at least when combined with the epistemic ought implies can principle I'll defend in §2. In the next section, I will introduce the indirect pragmatic picture, and examine the obstacles to being a skeptic when we assume that it correctly characterizes the nature of belief.

1.2 The Indirect Pragmatic Picture

On the direct pragmatic picture, the connection between beliefs and the rationalization of action is direct—whatever beliefs would best rationalize an agent's actions are the ones the agent has. On the indirect picture, there's still a constitutive link between beliefs and the rationalization of action, but it comes at

²⁹Stalnaker (1999) criticizes Dennett's approach in "On What's in the Head."

a one step remove, in the following sense. According to the indirect pragmatic picture, a proposition P counts as believed by an agent if a representation with the content that P is produced by a belief-producing psychological mechanism of that agent. However, for a system that produces representations to count a *belief*-producing system, it must have the function of producing representations that play the action-guiding, rationalizing role that the advocate of the direct pragmatic picture thinks beliefs must play.

The indirect pragmatic picture is a big tent—it is really a family of views, rather than a particular, precisely defined view like the direct pragmatic picture, which is just the conjunction of WRP and SRP. Particular versions of the indirect pragmatic picture can be obtained by spelling out the notion of what it is for a psychological mechanism to have the function of producing representations that play the belief-role. For instance, a biosemanticist (e.g., Millikan 1989) might appeal to considerations about the evolutionary history of the mechanism. I don’t intend much to rest on the idea of “function” in the above characterization—all that’s really necessary for a view to count as a version of the indirect pragmatic picture is that it holds that a representational system only counts as a belief system because of some connection the system has to rationalizing action, and that beliefs are the outputs of such a system.

A paradigmatic version of the indirect pragmatic picture is the language of thought hypothesis, (LOT) understood as a conceptual thesis about the nature of belief.³⁰ According to the LOT theorist, a proposition gets to be believed by an

³⁰LOT is sometimes treated as an empirical hypothesis, sometimes as a conceptual one. Field (1978) offers purely conceptual arguments for the LOT hypothesis, but Fodor (1975) offers both empirical and conceptual arguments.

agent because a sentence that expresses that proposition appears in the agent’s “belief box.” However, a representational system only counts as an agent’s belief box if its function is to produce outputs that combine with desires to rationalize action in the way that the direct pragmatic theorist accepts—in line with the remarks about function above, we will get more specific versions of the LOT view when we add a specific account of function.

In this section, while I won’t try to find a single anti-skeptical result common to all ways of developing the indirect pragmatic picture, I will argue that on various natural ways of developing this picture we can derive particular antiskeptical results, albeit weaker ones than those supported by the direct pragmatic picture. My general strategy will be to argue that while various versions of the indirect pragmatic picture allow for the possibility of skeptics, such skeptics cannot consistently believe that they are skeptics while accepting the specified version of the indirect pragmatic picture. To establish this, we will first have to consider different ways we might flesh out the schematic notion of “function” in the indirect pragmatic picture to get views about belief on which skepticism is possible.

One way an agent could be a skeptic, consistent with the indirect pragmatic picture, is similar to the case of “mad pain” discussed by Lewis (1983a). We might say that a representation-producing system counts as the belief system of an agent S just in case other members of S ’s species have physically similar representation systems that play the relevant role in guiding their behavior. In such a case, even though S ’s representation system doesn’t guide S ’s behavior, S ’s representation system would count as a belief box by proxy. In this scenario, S might be a skeptic because S ’s belief box could contain only sentences that describe the contents of her own mind, perhaps along with other sentences with agnosticism markers; without

the requirement that S 's own beliefs must rationalize S 's behavior, there is no obvious obstacle to allowing this possibility.

Another version of this example might involve S herself at some past time; S might have a subsystem that counts as a belief box, but only because of the role it used to play in guiding S 's behavior. Perhaps now, the only sentences that appear in it marked as beliefs are sentences that concern the present contents of S 's mind—the others have agnosticism markers.

Earlier in this section I alluded to biosemanticist views of function. We can develop a version of the indirect pragmatic picture that allows for the possibility of skepticism by appealing to the following biosemantics inspired claim: a representational system counts as a belief system for an agent just in case the structure of that system was shaped by natural selection for the purpose of producing outputs that play the action-guiding role typical of beliefs. That is, if an agent S 's ancestors had more reproductive success than their conspecifics due to their having progressively more sophisticated representational subsystems that produced outputs that played the action-guiding role typical of beliefs, then if S has some similar such representational subsystem herself, it counts as a belief-producing system, even if it doesn't itself produce outputs that guide S 's actions in the manner typical of beliefs. Again, because this picture imposes no requirement that S 's own beliefs must rationalize her behavior, there is no obvious obstacle to S 's being a skeptic consistent with this version of the indirect pragmatic picture.

We have seen a number of ways how, consistent with some version or other of the indirect pragmatic picture, an agent might be a skeptic. In light of these possibilities, is the indirect pragmatic picture of belief a congenial one for the skeptical epistemologist? It is certainly more congenial than the direct pragmatic picture,

but ultimately it still does not represent an attractive option for the skeptical epistemologist. While skepticism is possible on the indirect pragmatic picture, open-eyed skepticism—skepticism in which one realizes that one is a skeptic, and understands what being a skeptic entails—is not. Skeptical epistemology may be consistent with the indirect pragmatic picture, but they make for an awkward combination.

Why don't the views discussed so far allowed for open-eyed skepticism? If we opt for the “mad pain” inspired version of the indirect pragmatic picture, then someone who believes herself to be a skeptic must also believe herself to be a member of a species, most members of which are not skeptics. If we opt for the “past self” version of the indirect pragmatic picture, then someone who believes herself to be a skeptic must also believe herself to have existed in the past, and to have not been a skeptic in the past. If we opt for the biosemanticist version of the indirect pragmatic picture, then someone who believes herself to be a skeptic must believe herself to be a member of a species that evolved by natural selection. But none of these beliefs that one is committed to if one believes oneself to be a skeptic on the various versions of the indirect pragmatic picture—that one is a member of a species with certain characteristics, or that one existed in the past and acted in various ways—can be held by a skeptic.

More generally, natural versions of the indirect pragmatic picture entail that some psychological system's being a *belief* system is a highly extrinsic, contingent fact about that system, and is not the sort of thing that can be believed by a skeptic. Should this worry the skeptical epistemologist? Perhaps not—maybe the skeptical epistemologist should merely advocate skepticism, while acknowledging that we cannot consistently believe as we ought, understand what this entails, *and*

believe that we are believing as we ought. In the remainder of this section I'll argue that the resulting position would be an unattractive one.

One way of defending this position would be to argue that while some version of the indirect pragmatic picture of belief is correct, agents are rationally mandated to be agnostic about this fact: the skeptical epistemologist might hold that agents can consistently be skeptics and believe that they are skeptics, so long as they don't hold true beliefs about what their being skeptics entails. The first objection to such a view is that it is harder to motivate than standard skeptical positions—the skeptical strategy I considered at the beginning of this paper only motivates skepticism concerning contingent matters. Recall, this strategy involved claiming that in some epistemically important sense, scenarios in which P is true are just like those in which P is false for a wide range of propositions P . The strategy went on to use this claim to motivate agnosticism concerning the truth of the relevant propositions. For necessary truths, this strategy doesn't get off the ground, since there won't be any scenarios in which such truths fail to hold. But, plausibly, the true theory of belief is necessarily true; it is a theory about what features psychological states must have to *qualify* as beliefs, not about which features beliefs have as a matter of fact in the actual world.

Things are trickier if we take the true theory of belief to be necessary but *a posteriori*. While a full discussion of this matter would take us far afield, I'll simply note that the claim that the indirect pragmatic picture is true seems importantly different from the paradigm examples of the necessary *a posteriori* identified by Kripke (1980). Arguments for the indirect pragmatic picture, unlike arguments that water = H₂O, are of a broadly armchair sort; they rely on considering thought experiments and judging that someone would (or wouldn't) count as having various

beliefs in this or that hypothetical scenario.

The second, more flatfooted objection to such a position is the following. Suppose some version of the indirect pragmatic picture is true, and facts about what agents believe at a time *in fact* depend on facts about the agent's evolutionary history, or her conspecifics, or some other subject matter which, according to the skeptical epistemologist, agents ought not have beliefs about. In such a scenario, if an agent has beliefs about her beliefs, then she is plausibly violating the requirement of the skeptical epistemologist, *even if said agent does not believe that what she believes depends on her evolutionary history, conspecifics, and the like*. Her position would be somewhat like that of an agent who believed that she had hands, but also (strangely) believed that this was a necessary truth, and so not threatened by skeptical arguments. It is natural to say that such an agent does *not* count as a skeptic—she has a belief that she oughtn't have, according to the skeptical epistemologist; the fact that she (falsely) believes such a belief to be in a category of beliefs that is immune from skeptical challenges doesn't mean that her beliefs are actually legitimate by skeptical standards. According to this objection, someone who believes she is a skeptic is thereby *not* a skeptic, regardless of what she believes about which theory of belief is correct.

There is another, more natural strategy the skeptical epistemologist might use to defend the position that we ought to be skeptics without being open-eyed. The skeptical epistemologist might hold that not only ought we be agnostic about the external world, past, future, and various other matters, but also *about whether we are agnostic*. While this move avoids some of the pitfalls of the strategy discussed above, it still puts the skeptical epistemologist in a difficult dialectical position. The skeptical epistemologist must hold that a kind of reflective equilibrium in

epistemology is necessarily unattainable, at least for the skeptic—if you are a skeptic, then you cannot be in a position of simultaneously (a) believing a set of epistemological principles about which doxastic attitudes are appropriate in which circumstances, (b) holding some first order doxastic attitudes, and (c) recognizing a harmony between the epistemological principles referred to by (a) and the attitudes referred to by (b).³¹ It is natural to think that this is part of what we aim at in epistemology—we try to bring our first-order beliefs into harmony with our beliefs about what we ought to believe, and we hope that it will be possible to recognize success.

Treating such a harmony as one of the aims of epistemology doesn't immediately beg the question against the skeptical epistemologist, since skepticism doesn't obviously preclude this—we might have thought that the skeptical epistemologist just as much as her opponent can share this aim. The skeptical epistemologist, we might have thought, could hold that agnosticism is the appropriate first-order attitude, and also recognize that she herself is appropriately agnostic. Once we adopt views about belief like the indirect pragmatic picture, however, such a position becomes unavailable to the skeptic.

The skeptical epistemologist might embrace this result. She might argue that once we accept the indirect pragmatic picture, we should realize that beliefs about our beliefs are just as epistemically unfounded as beliefs about tables, chairs, and the like.³² She'll then reject the quest for reflective equilibrium, as achieving reflective equilibrium would require having unjustified beliefs—namely, beliefs about

³¹The idea that this sort of harmony is part of what we're aiming at when we do epistemology is of course inspired by the picture of the methodology of ethics defended by Rawls (1971).

³²Thanks to an anonymous referee for raising this point.

one's own beliefs. This position is certainly available, and nothing I've said so far tells against it. While it strikes me as a bullet-biting response, perhaps it's not a *much* less attractive position than the traditional package of skeptical views according to which we *can* achieve reflective equilibrium, but only by embracing skepticism.

I've argued that the indirect pragmatic picture is not much more hospitable to the skeptical epistemologist than the direct pragmatic picture. In the next section I will consider what features a view about belief would need to have to make it congenial to the skeptical epistemologist. While I won't attempt to offer knock-down objections to the idea that the correct theory of belief has such features, I will argue that theories with these features face hurdles that we needn't be anti-skeptical epistemologists to worry about.

1.3 The Phenomenological Picture

The preceding section suggests some features we should expect from an account of belief congenial to open-eyed skepticism. It should entail that the fact that some agent *S* has a particular set of doxastic attitudes at some place and time puts no restrictions on what the world is like at other places and times. In particular, an agent's being a skeptic at some place and time shouldn't entail anything about what the world is like at spatiotemporally remote locations. This way, it can allow that an agent can believe that she is a skeptic and understand what this entails, while also maintaining her skepticism. A similar view has been defended by a number of writers, not concerning doxastic attitudes but concerning qualia, or the

“introspectively accessible, phenomenal aspects of our mental lives” (Tye, 1997).³³ Suppose I am currently enjoying a piece of chocolate. Call the sensation I am experiencing *E*. According to some writers, reflection on thought experiments about various members of the philosopher’s bestiary—zombies and spectrum inverters, mostly—should convince us that my experiencing *E* is metaphysically independent of all physical facts along with any facts (mental ones too) about any other places and times; according to them it is metaphysically possible that the world could’ve been physically just like it actually is, and mentally just like it actually is except for the fact that I never experience *E*.³⁴ Consistent with everything else being the same, I might’ve had some other experience, or no experience, when I actually had *E*. Also, the world could’ve been as different as you like, consistent with my experiencing *E*. Create me and destroy me in the blink of an eye—I might still experience the taste of chocolate during my brief moment of existence.

Might a skeptical epistemologist try to develop a theory of doxastic states along the model of the account of qualia sketched above? In fact, there’s a growing literature on “phenomenal intentionality,” much of which attempts to provide just such an account (or at least, much of which holds that such an account is in principle available).³⁵ Engaging in depth with this body of literature is well beyond the scope of this paper. In the remainder of this section, I will make a *prima facie* case for the position that there are serious obstacles to developing a view of doxastic states on which they are similar to qualia in the respects the skeptical

³³Some writers use “qualia” to refer only to non-intentional states, but I intend no such restriction.

³⁴The view I have in mind is inspired by Chalmers (1996), though I’m not sure he’d accept all of the claims in the text.

³⁵See Lycan (2008) and Kriegel and Horgan (Forthcoming) for some recent surveys.

epistemologist needs. In particular, I will focus on two main obstacles to carrying out such a project. The first is that such a view would face difficulty in providing a satisfactory solution to the problem of intentionality, and the second is that doxastic states seem to be, in relevant respects, quite different from qualia. I will discuss these problems in turn.

The problem of intentionality is the problem of explaining how it is possible that “some things in the world—for example, pictures, names, maps, utterances, certain mental states—*represent*, or *stand for*, or are *about* other things—for example, people, towns, states of affairs.” (Stalnaker, 1984, p. 6) Intentionality is the property of representing or being about something. Doxastic states are paradigm instances of the sorts of things that are supposed to have intentionality—if I believe that trees have leaves, then my belief is *about* trees. The first two pictures about doxastic states I discussed—the direct and indirect pragmatic pictures—both suggest ways of tackling the problem of intentionality. According to the direct pragmatic theorist, I count as having beliefs about the weather because propositions about the weather play a certain role in a theory that predicts my behavior. For instance, I might count as believing that it will rain—being in an intentional state—because the proposition that it will rain plays an explanatory role in a psychological theory that rationalizes my actions; we may suppose that I put on my galoshes, carry an umbrella, etc.

The solution to the problem of intentionality on the indirect pragmatic picture is less straightforward, and will depend on the particular version of the indirect pragmatic picture one accepts. To take one example, Hartry Field (1978) outlines a two-part strategy for solving the problem of intentionality for mental states if we adopt the LOT theory of belief. The first step involves using a causal theory

of reference to explain how certain mental symbols manage to represent parts of the world. The second step involves explaining how we count as having beliefs when those mental symbols play a certain role in our cognitive economy. To take another example, on the “mad pain” version of the indirect pragmatic picture, we might start by identifying paradigm cases of belief, where such cases are exactly the cases identified by the direct pragmatic picture. We could then go on to claim that a creature would count as believing that P even if the direct pragmatic picture would not deliver this result, so long as the creature is in an internal state that, for most members of its species, would amount to a paradigm belief that P (i.e., would amount to a belief that P by the lights of the direct pragmatic picture).

While I do not have an argument that an informative answer to the problem of intentionality is *impossible* on the phenomenological picture of doxastic states, the prospects look dim. The phenomenological picture entails that an agent’s having doxastic states places few to no restrictions on what the world is like with respect to things other than the agent, so the advocate of the phenomenological picture has very little to work with. Advocates of the pragmatic pictures can appeal to relations between the agent and her environment—both facts about what objects causally impact the agent, as well as facts about the states of affairs the agent tends to bring about; facts of this sort play key roles in the pragmatic pictures in explaining how an agent manages to have beliefs that could be true or false *about* his environment.³⁶ It’s true that advocates of narrow content like Dennett who ac-

³⁶See Field (1978) to get an idea of the role played by the causal theory of reference in one version of the indirect pragmatic picture. Also, Stalnaker (1984, p. 15) defends a view of belief according to which “to believe that P is to be disposed to act in ways that would tend to satisfy one’s desires, whatever they are, in a world in which P (together with one’s other beliefs) were

cept the direct pragmatic picture can appeal to intrinsic but relational facts about an agent—e.g., facts about how she would act in various specific environments—to explain why her mental states have the content they do. But as argued earlier, Dennett’s view also leads to anti-skeptical results, and it’s hard to see how a similar strategy could lead to a view that allowed for the possibility of skepticism. The burden is on a defender of the phenomenological picture to explain how agents can have doxastic states that can be about their environments, without appealing to any of the facts of the sort that play key roles in the pragmatic solutions to the problem of intentionality. Perhaps the defender of the phenomenological picture will take the intentionality of doxastic states to be a brute, inexplicable fact.³⁷ If so, this is a serious drawback of the phenomenological picture, compared to the pragmatic pictures.³⁸

However, there is a second obstacle to developing the phenomenological picture, which is that doxastic states do not have much in common with qualia. In *The Conscious Mind*, David Chalmers draws a distinction between two classes of mental phenomena:

At the root of all this lie two quite distinct concepts of mind. The first is the *phenomenal* concept of mind. This is the concept of mind as conscious experience, and of a mental state as a consciously experienced mental state. . . The second is the *psychological* concept of mind. This is the concept of mind as the causal or explanatory basis for behavior.

true.” This is clearly a view according to which facts about which states of affairs that an agent tends to bring about are critical in fixing facts about what the agent represents.

³⁷In many cases, this is the strategy that advocates of phenomenal intentionality pursue.

³⁸Of course, advocates of the phenomenological picture are likely to think that the pragmatic pictures are hopelessly optimistic, and that taking intentionality as brute is unavoidable.

A state is mental in this sense if it plays the right sort of causal role in the production of behavior, or at least plays an appropriate role in the explanation of behavior. According to the psychological concept, it matters little whether a mental state has a conscious quality or not. What matters is the role it plays in a cognitive economy.

On the phenomenal concept, mind is characterized by the way it *feels*; on the psychological concept, mind is characterized by what it *does*. There should be no question of competition between these two notions of mind. Neither of them is *the* correct analysis of mind. They cover different phenomena, both of which are quite real. (1996, p. 11)

Defenders of the pragmatic pictures place doxastic states on the psychological side of Chalmers' distinction, while defenders of the phenomenological picture place doxastic states on the phenomenal side. However, it is not clear that beliefs have any distinctive phenomenology. While there may be something that it feels like to take a cool shower after a hot summer day, it is a strain to say that there is any similar sense in which there is something that it feels like to believe that the legislative branch of the New York State government is bicameral. One way to make this clear is to note that while beliefs can vary in strength, this variance is not accompanied by any concomitant variance in the intensity of some feeling. Frank Ramsey makes this point nicely:

We can, in the first place, suppose that the degree of a belief is something perceptible by its owner; for instance that beliefs differ in the intensity of a feeling by which they are accompanied, which might be called a belief-feeling or feeling of conviction, and that by the degree

of belief we mean the intensity of this feeling. This view would be very inconvenient, for it is not easy to ascribe numbers to the intensities of feeling; but apart from this it seems to me observably false, for the beliefs which we hold most strongly are often accompanied by practically no feeling at all; no one feels strongly about the things he takes for granted.³⁹ (1931, p.169)

The challenge facing the skeptical epistemologist is this. She needs an account of doxastic states according to which it is possible to be a skeptic (so the direct pragmatic picture will not do), and according to which an agent's being a skeptic places no restrictions on what the rest of the world apart from the agent is like (so the indirect pragmatic picture will not do). While some writers have thought that qualia have the features that the skeptical epistemologist needs doxastic states to have, the above considerations show that the analogy between qualia and doxastic states is weak; doxastic states have intentionality, but do not obviously have phenomenology.⁴⁰ So the phenomenological picture does not look promising. The skeptical epistemologist needs an account of doxastic states according to which they're characterized neither by their feel nor by their role in an agent's cognitive

³⁹I take it that this point also tells against views on which, while standing beliefs have no distinctive phenomenology, occurrent judgments *do*, and standing beliefs should be understood in terms of dispositions to have the phenomenology of occurrent judgment. Even when I explicitly consider the question of what my name is—a question about which I am quite confident—I don't have any powerful or intense phenomenology.

⁴⁰I don't mean to imply that it is obvious that qualia lack intentionality. Many writers defend the view that phenomenal states have representational properties (see Byrne (2001), for example). But such views don't help the skeptical epistemologist until they're combined with further views about the nature of phenomenal states.

economy, and it is not clear that one is to be had.

To be fair, the three-way distinction between the pragmatic pictures and the phenomenological picture does not exhaust logical space.⁴¹ One strategy that doesn't fall into any of the categories discussed so far would involve rejecting the pragmatic and phenomenological pictures, and holding that facts about who believes what are brute facts, irreducible to and inexplicable in terms of facts of other sorts (whether behavioral, dispositional, phenomenal, or whatever). It strikes me that this view would inherit many of the vices of the phenomenological picture, with fewer of the virtues. It would still be left unable to solve the problem of intentionality—it would have to take intentionality as brute. But it would also be vulnerable to the further charge of lack of motivation. When dualists and type-B materialists hold that phenomenal facts are not explicable in terms of physical facts, they can appeal to thought experiments that seem to support the distinctions they draw. It's far less clear, however, that analogous thought experiments could support the existence of brute belief facts. It doesn't even seem conceivable that there could be some scenario, physically and phenomenally just like this one, but in which people's beliefs were different from their actual beliefs.

Even if we set aside views that take facts about belief as brute, it might seem that there are other types of views about belief that I haven't addressed. It's true that some views in the literature emphasize the relation between beliefs and specifically *linguistic* behavior, rather than behavior more generally.⁴² It might

⁴¹Thanks to a referee for pressing me to clarify my stance on this issue.

⁴²Stalnaker (1984, chap. 2) contrasts his pragmatic picture with what he calls the “linguistic picture,” which he associates with writers like Donald Davidson and Hartry Field. I suspect that Robert Brandom (1994) should be included in this category as well, though his approach differs in crucial respects from Davidson's and Field's.

seem that such views provide another option distinct from the pragmatic pictures and the phenomenal picture (though without taking facts about beliefs as brute), and might provide a way out for the skeptical epistemologist. This appearance would be misleading. A version of the direct pragmatic picture that focused on linguistic behavior would still lead to the result that skepticism is impossible—the skeptical belief state wouldn’t rationalize any linguistic acts for the same reasons that it wouldn’t rationalize acts more generally. The sorts of background beliefs necessary to rationalize linguistic actions—beliefs about the meanings of words, the interests and knowledge of one’s audience, etc.—are just as off-limits to the skeptic as empirical beliefs more generally. A version of the indirect pragmatic picture that focused on linguistic behavior—e.g., a version according to which a system’s counting as a *belief* system depends on the way it guides specifically linguistic behavior—would still be a view on which a creature’s having beliefs will be an extrinsic fact about that creature, and so such views will still lead to the impossibility of open-eyed skepticism. The earlier sections of this paper would have played out much the same had I focused on views that stress the centrality of the relation between belief and language use specifically, rather than belief and action more generally.

Ultimately, while the three-way distinction I’ve employed doesn’t exhaust logical space, it does seem to me to exhaust the space of views that have been taken seriously in the literature, and also to exhaust the space of views that deserve to be taken seriously.

In the next section I will defend a version of the claim that ought implies can in epistemology. Together with my arguments from this section, it will imply that it is not the case that we ought to be skeptics (if the direct pragmatic picture is

true), or at least that it is not the case that we ought to be open-eyed skeptics (if the indirect pragmatic picture is true). But since the position that we ought to be skeptics even though we cannot be open-eyed skeptics is unattractive, I take it that even if the indirect pragmatic picture is true and the direct one is false, my arguments in the next section will motivate rejecting the claim that we ought to be skeptics.

2 Ought Implies Can

David Hume (2000) famously argued that because our beliefs are formed by custom or habit, skepticism is in some sense not a possible doxastic state for us; because of the nature of human psychology, we are unable to refrain from forming certain beliefs about the unobserved on the basis of the observed. An anti-skeptic impressed by Hume's observation might try to appeal to an epistemological ought implies can principle to leverage Hume's point about habit into an anti-skeptical argument. Such a strategy strikes me as unpromising; a skeptical epistemologist could reasonably respond to such a strategy by granting that our beliefs are determined by habit, but insisting that this is an epistemic *defect*—she might hold that the beliefs of a *rational* agent would *not* be determined by habit, and that such an agent would instead be a skeptic. That is, the relevant ought implies can principle that we'd need to generate anti-skeptical conclusions from Hume's observation is implausible—it rules out the possibility of contingent psychological defects in humans that make it impossible for us to be fully rational, but such defects seem possible (and probably actual).⁴³

⁴³See Kahneman et al. (1982).

In this section, however, I'll argue that the ought implies can principle necessary to generate anti-skeptical conclusions when combined with the claims of §1 is considerably weaker and more plausible than the one we'd need to generate anti-skeptical conclusions from Hume's claims about habit.

A number of writers have discussed the principle that ought implies can specifically in the context of epistemology—some endorsing it, some attacking it. Fred Dretske (2000, p.598) endorses a strong epistemological ought implies can principle, and uses it to argue that we have a right to hold perceptual beliefs because we are psychologically incapable of failing to form them. William Lycan (1985, p.146) suggests that ought implies can is less plausible in epistemology than in ethics, and expresses doubt that beliefs can be warranted in virtue of our being psychologically incapable of abandoning them. My strategy in this section will be to distinguish the principle I am defending from other, stronger principles, and to show how some *prima facie* problems for those stronger principles do not arise for the one I need for my anti-skeptical argument. I will then go on to argue that the principle I need is independently plausible.

One motivation for an epistemological ought implies can principle—a motivation I will ultimately reject—involves the thought that people can only be obligated to be in some doxastic state if they can be blamed for failing to be in it, and they can only be blamed for failing to be in it if there is something they can do to get into it. The fate of an ought implies can principle motivated by this thought will be closely tied to the fate of doxastic voluntarism—the doctrine that people can exert voluntary control over their beliefs. But if *S*'s being obligated to believe that *P* implies that *S* is able to believe that *P*, and *S*'s being able to believe that *P* requires *S*'s being able to exert voluntary control to come to believe that

P , then problems loom. Doxastic voluntarism is far from obvious,⁴⁴ and even if some beliefs are subject to voluntary control,⁴⁵ it is hard to deny that some aren't. Nevertheless, those that aren't may still be irrational, unjustified, or such that we ought not to hold them. For example, a hopeless paranoid schizophrenic may be unable to give up the belief that he is the target of a Martian conspiracy. That he is psychologically incapable of giving up this belief does not mean that he's not irrational for holding it; in the relevant sense of "ought," it is still the case that he ought to give it up. Being unable to appreciate the force of the evidence that Martians neither exist nor have infiltrated Earthly governments does not change the force of that evidence. By the same token, a skeptical epistemologist might hold that our beliefs are determined by habit (as Hume said) and so we are unable to be skeptics, but that this makes us in relevant respects like the hopeless paranoid schizophrenic.

Examples like this might lead us to conclude that there is no interesting version of an ought implies can principle in epistemology. Alternatively, we might look for ways of weakening the principle to avoid the result that our schizophrenic has a clean bill of epistemic health. By substituting in more or less demanding senses of "can," we get stronger or weaker versions of the principle that ought implies can in epistemology. For instance, while we might understand "can" such that a subject can be in or fail to be in some doxastic state only if she can bring it about that she is or is not in that state by exercising voluntary control, we needn't. We

⁴⁴Alston (1988) argues that we shouldn't conceive of epistemology in deontological terms, roughly because he thinks that ought implies can in the sense described above, but beliefs aren't subject to the sort of voluntary control that would support obligation claims.

⁴⁵For instance, if Velleman (1989) is right about the nature of intentions, then intentions are a type of belief that is subject to voluntary control.

might instead understand “can” such that a subject can be in a doxastic state just in case it is metaphysically possible for a subject be in that state, regardless of whether she’s able to do anything to get herself into that state. Armed with this understanding of “can,” the principle that ought implies can is immune to counterexamples like the one involving the schizophrenic. While the schizophrenic is unable to get himself into the state of not believing in Martian conspiracies, such a state is one that it is metaphysically possible for one to be in. I am in it, and I suspect you are too. So our modified principle is consistent with the claim that the schizophrenic ought not to believe in Martian conspiracies.

Is there any motivation for an ought implies can principle once we reject the blame-based motivation that leads to a principle requiring voluntary control over one’s beliefs? I think there still is. It might be useful for us to evaluate beliefs with a set of epistemic criteria according to which certain belief states would be ideal, even though people are sometimes (or perhaps always) psychologically incapable of achieving these ideals, perhaps because of phobias, mental illnesses, or just computational limitations. This is why the strong principle fails. But it seems that even a (psychologically) unattainable ideal must represent a metaphysical possibility for it to provide a useful yardstick against which to compare the doxastic states of actual believers.

My inability to run faster than a cheetah (or really very fast at all) reflects a respect in which I fall short of the ideal when it comes to speed. There is some temptation to say that the ideal in speed would require instantaneous travel—and so is probably unattainable in any universe with our physical laws. However, there is no temptation to say that the ideal in speed requires being faster than oneself. To be faster than oneself is metaphysically impossible, and that “ideal” can’t even

be approached by possible runners—getting faster does not help—and so does not provide a useful or illuminating yardstick against which to compare actual or possible runners. As Aristotle said in his *Politics*, “In framing an ideal we may assume what we wish, but should avoid impossibilities.” (1265a17-18)

I would like to say the same thing about skepticism. My inability to believe every logical truth plausibly reflects a respect in which I fall short of an epistemic ideal, even though believing all logical truths is humanly impossible.⁴⁶ The ideal of believing all logical truths, while humanly unattainable, is one that could be satisfied by some possible creatures, and we can see how humans can approximate this ideal without ever fully reaching it. But my inability to completely suspend judgment with respect to all possibilities in which things appear as they actually do (or, if we accept the indirect pragmatic picture, to knowingly do this while understanding what that entails) is not a respect in which I fall short of an ideal. If my arguments from §1 are sound, then this is a metaphysically impossible doxastic state, and it is not even clear what possible doxastic states would count as approximations to the impossible agnostic one—it does not make for a useful or illuminating benchmark against which to compare the doxastic states of actual or possible believers.⁴⁷

Suppose the skeptical epistemologist rejects the appeal to unsatisfiable epistemic obligations—might there be another way out? Here’s one possibility. It

⁴⁶See Christensen (2004) for the idea that logical omniscience is an epistemic ideal, albeit a psychologically unattainable one.

⁴⁷I believe that similar considerations tell against the skeptical fallback position according to which we face an *epistemic dilemma*, in that we have various epistemic obligations that can be fulfilled individually but not jointly. Discussing this, however, would make an already long paper longer.

is natural to understand the original skeptical argument as resting on two main claims. First, for any two hypotheses that are consistent with present appearances, we have no reason to think either is any more likely than the other. Second, when we have no reason to think one hypothesis is more likely than another, we ought to be agnostic as to which is true. Given these two claims, we ought to be agnostic about quite a lot. But even if we reject the second claim and reject the conclusion, the first on its own is enough to generate surprising and counterintuitive results. Not only are we normally inclined to deny that we ought to be agnostic between commonsense hypotheses and skeptical ones, we also think we have good reasons for believing in common sense.

We might imagine a fallback skeptical position according to which anything goes in the realm of belief—the skeptical epistemologist would hold that while we ordinarily take ourselves to have reason to believe in commonsense hypotheses, in fact these hypotheses are on an epistemic par with skeptical ones, and we would be just as well within our epistemic rights to believe in deceiving demons and clever vatmasters as we are to believe in tables and chairs. The skeptical epistemologist’s fallback position is a version of epistemic permissivism—the view that after we have fixed a body of evidence, there might be many permissible sets of beliefs one might adopt on the basis of the evidence, and no one set of beliefs need be epistemically mandatory.⁴⁸ This is an intriguing suggestion, but I won’t take it up here.

⁴⁸See White (2005).

3 Conclusions

I have argued against the skeptical view that we ought to be agnostic between all hypotheses that are consistent with how things appear to us. Plausible views about the nature of belief—the direct and indirect pragmatic pictures—imply that we cannot be agnostic in this way, or cannot knowingly be agnostic in this way. Since ought implies can in epistemology—at least in the weak sense that I defended in §2—this means that plausible views about the nature of belief conflict with skeptical epistemological views.

If my arguments in this paper are sound, skeptical views in epistemology can be refuted without drawing on considerations from within epistemology proper. While my arguments are consistent with most positive epistemological theories about what we ought to believe and why, no particular rationalist or empiricist arguments are necessary to show that the skeptical epistemologist is wrong about what we ought to believe; it is enough to note that she asks the impossible.

Still, we might wonder just how much epistemological significance the arguments of this paper really have. Skeptical arguments like the ones sketched at the beginning of this paper are sometimes presented as paradoxes.⁴⁹ Understood this way, the epistemologist's task is not that of determining whether or not their conclusions are correct—it's obvious at the outset of inquiry that they are not—rather, her task is to explain where they go wrong, and to hopefully reveal some surprising epistemological facts in the process.⁵⁰

But the arguments of this paper, it might seem, don't help with that task. Even

⁴⁹See Byrne (2004) for a paradigm example of this sort of presentation.

⁵⁰While this is a popular approach to thinking about skeptical arguments, it is not universally shared. See Unger (1975).

if they provide additional, novel reasons to believe that the skeptical epistemologist is wrong, so what? We already knew that the skeptical epistemologist was wrong, and it may seem as if the arguments of this paper don't point to any *further* epistemologically important upshots. In the remainder of this paper, however, I'll argue to the contrary that the anti-skeptical arguments I've presented do point to some surprising, epistemologically interesting conclusions.⁵¹

It is now common to distinguish between ambitious and modest anti-skeptical projects.⁵² The ambitious project is to refute the skeptical epistemologist on her own terms—to refute her without begging the question. The modest project involves establishing to our own satisfaction that the skeptic is wrong—in pursuing the modest project we may appeal to premises the skeptical epistemologist would reject, so long as such premises strike us as reasonable on reflection. The conventional wisdom is that refuting the skeptical epistemologist on her own terms is impossible, and that the best we can achieve in our anti-skeptical efforts is a satisfactory completion of the modest anti-skeptical project.⁵³

If the arguments of this paper are sound, however, the conventional wisdom is too cautious. Appealing to premises in the philosophy of mind—in particular, appealing to the direct or indirect pragmatic pictures—doesn't constitute begging the question against the skeptical epistemologist. Nor does appealing to the claim that epistemic oughts imply metaphysical cans. While it's a slippery matter exactly what constitutes begging the question, there is a clear difference between my anti-skeptical arguments and paradigm instances of the modest anti-skeptical

⁵¹Thanks to an anonymous referee for pressing me to expand on these issues.

⁵²The distinction I have in mind is made by Pryor (2000).

⁵³This is certainly Pryor's view. It is also, as far as I can tell, the view defended by Timothy Williamson (2000, 2007).

project.⁵⁴ Responding to skepticism by assuming that we have some empirical knowledge (Williamson, 2000), or by assuming that we are justified in taking our experiences at face value (Pryor, 2000), would certainly count as begging the question, and therefore would not constitute a successful completion of the ambitious anti-skeptical project; these positions come close to simply asserting what the skeptical epistemologist denies. But responding to the skeptical epistemologist by appeal to a combination of premises in the philosophy of mind and premises concerning the relation between epistemic obligation and metaphysical possibility does not amount (at least in the first place) to asserting anything all that close to what the skeptical epistemologist denies.

So one surprising epistemological consequence of the arguments of this paper is that not only is the skeptical epistemologist wrong, but she can be refuted with non-question-begging arguments.

The arguments of this paper also bear on debates outside of skepticism. If we accept these arguments, there's at least *some* pressure to accept a broadly coherentist methodology of inquiry. According to this methodology, the proper starting point of inquiry is whatever we start out believing; our beliefs aren't undermined or shown to be unjustified if it turns out that they could not have been arrived at via a process of belief revision that started from extreme skepticism.⁵⁵

⁵⁴We might (unhelpfully) say that any valid argument for P must beg the question against a believer in $\sim P$, since someone who believes that $\sim P$ is committed (implicitly at least) to rejecting the premises of any valid argument for P .

⁵⁵For recent defenses of the view that the proper starting point of inquiry is all of what we start out believing, rather than some privileged subset (e.g., facts about our present phenomenal states), see Harman (2001), Maddy (2007), and Stalnaker (2008). The idea goes back at least as far as Neurath.

If skepticism were a real possibility for us, then it would be tempting to look for a kind of foundationalist epistemology that could explain how one could start out as a skeptic, and then rationally reason one's way to a commonsense picture of the world. But if skepticism is impossible, then we might be happy to settle for an epistemological picture that only tells us how to proceed given the beliefs we have, and doesn't tell us how it *would* be rational to proceed from some impossibly impoverished doxastic position.

On the picture I'm sketching, the skeptical epistemologist and a certain kind of anti-skeptical foundationalist share a common assumption. They agree that we can identify what we should believe by determining what belief state would result if we started out as skeptics, and proceeded by revising our beliefs according to the correct epistemic rules. Their disagreement concerns only the nature of the correct epistemic rules—the skeptical epistemologist thinks that the correct rules won't let us escape our skeptical starting point, while the anti-skeptical foundationalist thinks that they will. But once we come to think that this starting point is an impossibility, the common assumption of the skeptical epistemologist and the anti-skeptical foundationalist starts to look questionable. It's not clear why we should care about how far we could get towards reconstructing a commonsense view of the world if we started out in an impossible doxastic state and followed the correct epistemic rules.⁵⁶

In this conclusion I have discussed two epistemological projects: the ambitious project of refuting the skeptical epistemologist on her own terms, and the foundationalist project of showing how belief revision can proceed from a starting point

⁵⁶It's not even clear whether the question makes sense—this will depend on our views about counterpossible conditionals.

of skepticism. If the arguments of this paper are sound, then contrary to conventional wisdom, the first project can be completed—we can not only come up with an internally consistent, defensible anti-skeptical position, but we can go further and refute the skeptical epistemologist without begging the question. The second project, however, would come to seem less interesting if we accepted the arguments of this paper. It's no surprise that the skeptical epistemologist is wrong. But if she turned out to be wrong for the reasons I've identified, that *would* be surprising, and would force us to rethink a number of debates in contemporary epistemology.

Bibliography

- Alston, W. P. (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives*, 2:257–299.
- Aristotle (2008). *Politics*. Translated by Benjamin Jowett and Henry William Carless Davis. Cosimo Inc.
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press.
- Bratman, M. (1992). Practical Reasoning and Acceptance in a Context. *Mind*, 101(401):1–16.
- Brown, C. (2002). Narrow Mental Content. In *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/content-narrow/>.
- Brueckner, A. (1986). Brains in a Vat. *Journal of Philosophy*, 83(3):148–167.

- Burnyeat, M. (1983). Can the Skeptic Live His Skepticism? In Burnyeat, M., editor, *The Skeptical Tradition*. University of California Press.
- Byrne, A. (2001). Intentionalism Defended. *The Philosophical Review*, 110(2):199–240.
- Byrne, A. (2004). How Hard Are the Sceptical Paradoxes? *Nous*, 38(2):299–325.
- Chalmers, D. J. (1996). *The Conscious Mind*. Oxford University Press, Oxford.
- Christensen, D. (2004). *Putting Logic in its Place*. Oxford University Press, Oxford.
- Cohen, L. J. (1989). Belief and Acceptance. *Mind*, 98(391):367–389.
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Clarendon Press, Oxford.
- Davidson, D. and Harman, G., editors (1972). *Semantics of Natural Language*. D. Reidel, Dordrecht.
- Dennett, D. (1987). *The Intentional Stance*. The MIT Press, Cambridge, MA.
- DeRose, K. (1995). Solving the Skeptical Problem. *The Philosophical Review*, 104(1):1–52.
- Dretske, F. (2000). Entitlement: Epistemic Rights without Epistemic Duties? *Philosophy and Phenomenological Research*, 60(3):591–606.
- Field, H. (1978). Mental Representation. *Erkenntnis*, 13(1):9–61.
- Fodor, J. (1975). *The Language of Thought*. Harvard University Press, Cambridge.

- Fraassen, B. C. V. (1989). *Laws and Symmetry*. Oxford University Press.
- Gibbard, A. (2003). *Thinking How to Live*. Harvard University Press, Cambridge, Massachussetts.
- Harman, G. (1986). *Change in View: Principles of Reasoning*. The MIT Press, Cambridge, MA.
- Harman, G. (2001). General Foundations Versus Rational Insight. *Philosophy and Phenomenological Research*, 63(3):657–663.
- Horwich, P. (1991). On the Nature and Norms of Theoretical Commitment. *Philosophy of Science*, 58(1):1–14.
- Hume, D. (2000). *An Enquiry Concerning Human Understanding: A Critical Edition*. Oxford University Press.
- Jaynes, E. T. (1968). Prior Probabilities. *IEEE Transactions On Systems Science and Cybernetics*, pages 227–241.
- Jaynes, E. T. (1973). The Well-Posed Problem. *Foundations of Physics*, pages 447–493.
- Joyce, J. (2005). How Probabilities Reflect Evidence. *Philosophical Issues*, 19(1):153–178.
- Joyce, J. (2010). A Defense of Imprecise Credences in Inference and Decision Making¹. *Philosophical Perspectives*, 24(1):281–323.
- Kahneman, D., Tversky, A., and Slovic, P. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.

- Kriegel, U. and Horgan, T. (Forthcoming). The Phenomenal Intentionality Research Program. In *Phenomenal Intentionality*. Oxford University Press.
- Kripke, S. A. (1972/1980). *Naming and Necessity*. Harvard University Press, Cambridge, MA. Reprinted from Davidson and Harman 1972, with a new preface.
- Lewis, D. (1974). Radical Interpretation. *Synthese*, 27:331–344.
- Lewis, D. (1983a). Mad Pain and Martian Pain. In *Philosophical Papers, Volume I*, pages 122–130. Harvard University Press.
- Lewis, D. (1983b). New Work for a Theory of Universals. *Australasian Journal of Philosophy*, 61:343–377.
- Lycan, W. (2008). Phenomenal Intentionalities. *American Philosophical Quarterly*, 45:233–252.
- Lycan, W. G. (1985). Epistemic Value. *Synthese*, 64(2).
- Maddy, P. (2007). *Second Philosophy: A Naturalistic Method*. Oxford University Press.
- Melchert, N. (1985). Why Constructive Empiricism Collapses Into Scientific Realism. *Australasian Journal of Philosophy*, 63(2):213–215.
- Millikan, R. G. (1989). Biosemantics. *The Journal of Philosophy*, 86(6):281–297.
- Norton, J. (2008). Ignorance and Indifference. *Philosophy of Science*, 75(1):45–68.
- Pryor, J. (2000). The Skeptic and the Dogmatist. *Nous*, 34.

- Putnam, H. (1981). *Reason, Truth, and History*. Cambridge University Press, Cambridge.
- Ramsey, F. (1931). Truth and Probability. In *The Foundations of Mathematics and Other Logical Essays*, pages 122–130. Harcourt, Brace, and Co.
- Rawls, J. (1971). *Theory of Justice*. Belknap Press, Cambridge Massachussetts.
- Ross, J. (2006). *Acceptance and Practical Reason*. Rutgers University, PhD Dissertation.
- Sinnott-Armstrong, W. (2006). *Moral Skepticisms*. Oxford University Press, New York.
- Stalnaker, R. (1984). *Inquiry*. The MIT Press, Cambridge, MA.
- Stalnaker, R. (2008). *Our Knowledge of the Internal World*. Oxford University Press.
- Stalnaker, R. C. (1999). *Context and Content*. Oxford University Press, Oxford.
- Strawson, G. (1994). *Mental Reality*. The MIT Press, Cambridge, MA.
- Tye, M. (1997). Qualia. In *Stanford Encyclopedia of Philosophy*.
<http://plato.stanford.edu/entries/qualia/>.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford University Press, Oxford.
- van Fraassen, B. (1980). *The Scientific Image*. Oxford University Press, Oxford.
- Velleman, D. (1989). *Practical Reflection*. Princeton University Press, Princeton.

- Vogel, J. (2004). Skeptical Arguments. *Philosophical Issues*, 14(1):426–455.
- White, R. (2005). Epistemic Permissiveness. *Philosophical Perspectives*, 19(1):445–459.
- White, R. (2010). Evidential Symmetry and Mushy Credence. In *Oxford Studies in Epistemology, Volume 3*. Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Williamson, T. (2007). *The Philosophy of Philosophy*. Blackwell, Oxford.
- Wright, C. (2004). Wittgensteinian Certainties. In *Wittgenstein and Scepticism*. Routledge.