

Interest as a Starting Place for Philosophy

Abstract: This paper discusses a puzzle about philosophical beliefs. Core philosophical beliefs that are widely shared among philosophers, such as the belief that skepticism is false, are often held with extreme confidence. However, this confidence is not justified if these beliefs are based on what are traditionally seen as the sources of philosophical evidence, such as intuitions or observation (or reasoning on these bases). Charity requires that we should look for some other basis for these beliefs. I argue that these beliefs are based on our knowledge of what we find interesting. Further, I argue that this is a good basis for belief. Knowing what we find interesting allows us to tune our inquiry in ways we could not otherwise.

How do philosophers know the things we know about philosophical topics – about, for example, justice, free will, knowledge, or reasons? A conventional answer is that some of what we know we know via observation or intuition, some we know via testimony or recollection, and some we know from reasoning about the data provided by these other sources. But there's a puzzle: it turns out that these routes to knowledge don't explain the confidence we have in some of our philosophical beliefs, often ones that are widely held and ground quite of bit of theorizing. This might suggest that many of us are not as good at philosophy as we'd like to think, but before we accept that we should first see if there is a more charitable explanation. This would be another type of data that philosophers can appeal to that justifies a high level of confidence in the puzzling beliefs. I think that that data is data about what we find philosophically interesting. Knowing what we find interesting gives us knowledge about the target of our inquiry: it has those qualities that we are interested in. While this may seem trivial, I'll argue that it not only

explains our puzzling confidence in core philosophical beliefs, but it also can be quite useful in tuning our inquiry appropriately.

I will begin by arguing that there is a puzzle to be solved, by discussing some examples of widely and very confidently held philosophical beliefs, where the apparent confidence we have in these beliefs is not supported by intuitions or reasoning that is based on intuitions. I focus on the inadequacy of intuitions because the other routes to knowledge – observation, testimony, and memory – aren't plausibly relevant to the sorts of beliefs I'll discuss. I will then give an account of how knowing what we find interesting can explain this confidence. This is to suggest that this knowledge is, in fact, a source of data that philosophers appeal to, perhaps tacitly. I will also argue that this is a *good* source of data – that there are advantages to founding philosophical inquiry on both intuitions and interestingness. The paper ends with considerations of some worries about the view, and questions we should go on to investigate.

The puzzle

The vast majority of philosophers today (roughly 96%) are not skeptics, and most (81%) are non-skeptical realists about the external world (PhilPapers.org, 2009). I take us non-skeptics to be (in general) extremely confident about our rejection of skepticism. This confidence is reflected in survey data, as a large majority of those endorsing non-skeptical realism in the 2009 PhilPapers.org survey gave the most confident possible answer (accepting the view rather than “leaning towards” it), which was not true for many of the other questions on the survey. This confidence is also reflected in the fact that a certain argumentative move is fairly standard in epistemology: pointing out that one's opponent's view leads to skepticism is often taken to refute the view. To be justified in such confidence (on most views of knowledge), we would need to be justified in equally high confidence that our perceptual beliefs are justified. This

confidence is hard to explain. To show why, I will discuss the non-skeptical claim that all or most of our perception-based beliefs are at least *prima facie* justified. Skepticism would be refuted if even one of these beliefs were justified and true, but I doubt that anyone who rejects skepticism thinks just that; certainly it is very uncommon for those who discuss skepticism or justification to argue as if they do. So I will attribute the stronger claim to non-skeptics and explore how we might be so confident in it.

We can confidently reject skepticism if we can confidently accept a theory of justification that shows that our perception-based beliefs are justified. Let's accept that we can justifiably accept some such theory (perhaps several such theories are currently justifiable). Even so, we can't be warranted in being extremely confident that our preferred theory is correct. For any theory of justification currently on offer, a worrisome and plausible objection has been made to it. Space doesn't permit my listing these, but the fact that these exist shouldn't be surprising, given that there is an ongoing debate among epistemologists in this area. Further, that there is wide-spread disagreement among epistemologists – the experts in this area – about which theory is correct suggests that epistemologists as a group are not all that reliable. Finally, basing our confidence that skepticism is false in our confidence that a specific theory of justification is true looks circular, since typically it is taken as a precondition for accepting a theory of justification that it does not entail skepticism. If our confidence that skepticism is false were based just on our acceptance of a specific theory of justification, then this confidence could be (justifiably) no higher than our confidence in the theory itself; thus, our confidence that skepticism is false cannot be justified in this way.

This does not mean that we can't be very confident that skepticism is false: all we need is justified very high confidence in believing there is *some* correct theory of justification, and that it will show that our beliefs about the external world are justified. What might underwrite that?

It might be intuitions. When I consider specific beliefs of mine about the external world, or specific perception-based beliefs, I do have the intuition that these are justified. But I don't think that these intuitions about specific beliefs make sense of most philosophers' confidence that our perceptual/external-world beliefs are *in general* justified. I would be very surprised if many of us have considered very many of our perception-based beliefs and checked to see if it is intuitive that each is justified; I know of no anti-skeptical arguments in the literature that make use of checking of this sort.¹ If we haven't checked to see if very many of our perception-based beliefs are intuitively justified, we might still have checked *some*, and we might be doing induction from this small group to our perceptual beliefs more generally. That induction would be too weak to justify the confidence we have that most or all of our perception-based beliefs will turn out to be justified. Further, arguments based on intuitions that perceptual beliefs are justified actually reverse how many philosophers argue: it is quite common to argue that intuitions justify beliefs on the grounds of their similarity to perception, so that if perception justifies belief, intuitions must as well (see e.g. Sosa, 1998, Huemer, 2001). Such arguments would make no sense if our trust in perception were based in our trust in intuitions. I conclude that intuitions about specific perceptual beliefs are not the basis of our confidence that skepticism is generally false.

Perhaps, though, we have a more general intuition about the justification of our perception-based beliefs. What would that intuition be, and how confident could we be that it is true? Obviously the intuition can't be, "All perception-based beliefs are justified," as this is

false; the justification of our perception based beliefs can be defeated. Perhaps the intuition is “All perception-based beliefs are *prima facie* justified,” or “All perception-based beliefs are defeasibly justified.” Here I have to wonder: how many of us really have intuitions with that content? Michael Huemer, for example, says that the claim that all perception-based beliefs are *prima facie* justified is “self-evident,” (Huemer, 2001, 103), which looks a lot like saying that it is intuitive. But then he goes on argue vigorously in its defense, which suggests to me that he didn’t mean that it is intuitive, but rather that the argument for it is extremely compelling. My own experience (and discussions with my colleagues suggest that such experience is typical) is that I rarely have intuitions with such clearly philosophical content – that is, intuitions whose content contains concepts and ideas that take a great deal of philosophical training to work with. What’s more, when I start to have such intuitions, they are almost always the product of a great deal of philosophical work on a topic, and tend to confirm whatever view I have been working out in this area. Even if such intuitions do justify our beliefs (and I’m a bit skeptical of this), we should at least somewhat dubious of them, as it’s quite plausible that they are too theory-laden to be good evidence.

The preceding point about intuitions with “philosophical content” is important to my later arguments and not obvious, so let’s give it a bit more attention. Why think that we should be somewhat dubious about intuitions with “philosophical content?” It should be plausible that non-philosophers do not typically have intuitions whose content “contains concepts and ideas that take a great deal of philosophical training to work with,” as this is almost true by definition. What’s more, when we reflect on how difficult it can be to motivate certain philosophical intuitions or views with our students, or when we look at how hard experimental philosophers work to frame the prompts or questions in their studies, we can see that it is rare for non-

philosophers to have intuitions with “philosophical content.” So, if we have such intuitions, we should worry that they are the product of our training and of our thought on these matters. But why worry that they are likely to be theory laden – that they reflect what our prior beliefs or the views we’d like to endorse? Aside from the *prima facie* plausibility of that claim, there is psychological research to back it up (this is not from studies of professional philosophers, unfortunately; psychologists have not spent much time on us yet). One is more likely to believe what one is motivated to believe (see Kunda, 1990, for a survey of the literature on this). This is partly due to changes in our reasoning strategies on these matters, but at least some of the psychological research on this topic points to unconscious mechanisms as giving rise to these beliefs. Such mechanisms either produce intuitions (on some definitions of “intuition”) or they produce mental states that *feel* like intuitions. In either case, we should expect that people (through no fault of their own) are more likely to find intuitive what is in their interest to find intuitive. It is in our interests to find our own theories intuitive. Second, it turns out that the familiarity of a claim can increase our perception of its truth (see Dechêne et al 2010 for a survey of this literature). This can be overridden in various ways – if we know that the claim is false, then familiarity will be overridden. But, for topics we are working on and that we have some prior inclination to believe, greater familiarity with central theses in these areas can lead to the feeling that they are true; these feelings either are intuitions or are easy to mistake for them. So, we have empirically based reasons to worry that any intuitions with “philosophical content” that we do have are the product of our theorizing. While this might not defeat our justification in believing the content of these intuitions, it should certainly reduce our confidence that these intuitions are trustworthy.

To further undermine the claim that general intuitions ground our confidence that skepticism is false, I will point out that intuitions with such general content are not the best sort of evidence. After all, even a single counterexample will show them to be false. It's often acknowledged that general intuitions are weaker evidence than intuitions about specific cases (see, e.g. Cohen, 1986). So, it's implausible that the general intuition that all perception-based beliefs are *prima facie* justified is what underwrites our general confidence in the falsity of skepticism, as it's implausible that all of us actually have that intuition, or that it warrants the kind of confidence that we have in rejecting skepticism.

We can make similar points about any intuition that looks like, "All beliefs with property P are [*prima facie*] justified." I think the best-case scenario here is that we have the intuition that *most* of our perception-based beliefs are justified (perhaps defeasibly), or that skepticism is false. There are still some worries to be had here, as I think these still look like the product of theorizing, but let's put those to the side. Can these intuitions underwrite our confidence that skepticism is false?

One does not have to look far to find someone who is dubious of intuitions as evidence, but one *does* have to look far to find an external-world skeptic. If confidence in the justification of our external-world beliefs were generally driven by intuitions, this would not be the case. Further, there are plenty of philosophers who, while they do not entirely doubt that intuitions are evidence, certainly would want to check closely into this before trusting them (reliabilists, for example). Yet I would bet that their confidence that their perceptual beliefs are justified outstrips the confidence that they would see themselves as being warranted in having in their intuitions.

I am sure there are arguments to be raised about each point I have made, and theories one could advance to undermine each. I do not have the space to address these. I will instead make a general point: given the number of worries I have raised about the grounds of our confidence that skepticism is false, a response to all of them will very likely have to involve the conjunction of a number of claims. For this to explain our confidence, we would have to be more confident in every claim in the given conjunction than we are that skepticism is false. I feel safe in assuming that, for most of us, the confidence we are warranted to have in the conjunction of things we'd need to believe to respond to what I have said is lower than the confidence we actually have that skepticism is false (given how high that confidence seems to be). And so the question remains: how can we be justifiedly so confident that skepticism is false?

This is not the only belief that is puzzling in this way. To find an example of your own, look for cases where philosophers seem extremely confident in some claim and where the rationality of that confidence is not well explained by any of the typically cited sources of philosophical evidence (e.g. intuition and observation, or inference from these). It will be easiest to find examples by looking at general claims that are believed with extremely high confidence. This sort of confidence is not easily supported by intuitions: we often don't have general intuitions, but rather do induction from cases, and we often have not considered enough cases for this to be strong enough to support our beliefs; when we have such general intuitions, if they have "philosophical content" then their justification is somewhat undermined by worries about theory-ladenness; and general intuitions don't warrant very strong confidence in any case as they are easily defeated by intuitions about individual scenarios. If there are plausible counterexamples or counterarguments to the general view that require additional claims to respond

to, you have an even better example. And if it is a belief that most everyone in an area shares, despite methodological disagreements, even better.

I'll quickly give two more illustrations from epistemology (the area I'm most familiar with). A large range of philosophers who share very different approaches to epistemology agree that knowing that p is better in general than merely truly believing that p , and that explaining this is a desideratum for a theory of knowledge.ⁱⁱ This is a general claim, so it must be supported either by induction from cases or general intuitions; neither gives terribly strong support. This support is further undermined by how notoriously difficult it is to account for the value of knowledge, and the number of apparent counter-examples (see, e.g. Kvanvig, 2003, Zagzebski, 2003). It is a puzzle that so many different philosophers feel so confident that this is true. For another example, consider the widely accepted view that knowing that p implies that p . Here we have a general claim that many, if not most, philosophers seem extremely confident in. A general intuition that this was true looks like something only a philosopher would intuit, and so general intuitions that support this claim are immediately suspect. We also have some intuitive counter-examples, as intelligent people often use "knowledge" non-factively (e.g. here's an assertion by a scholar in a book on ancient Greece: "People knew that the gods were in contact with them because the gods spoke to them..." (Erskine, 2003, 412)) One can respond to these counterexamples, but this requires a set of views the conjunction of which is probably less likely than we see the initial claim about knowledge to be.ⁱⁱⁱ No detailed discussion can be given here, but it should seem plausible that our confidence that knowledge is factive outstrips the confidence warranted by the obvious data

And so we have a puzzle: are many of us doing poorly in forming the fundamental beliefs upon which much of our philosophical inquiry is based, or is there some other source of evidence that justifies this confidence?

A solution to the puzzle

Why do we study justification? In part, because it is something that we think we can and do have – if we didn't think it were possible or very likely that our beliefs were justified, justification wouldn't be something that we would be interested in. Lawrence BonJour articulates something like this when he says that epistemology starts from the commonsense conviction that there are “*good reasons* within our cognitive grasp for thinking that our various beliefs about the world are true... If we did not have such a conviction, there would be nothing particularly implausible about skepticism...” (BonJour & Sosa, 2003, 40) I think there is more here than just a groundless conviction: we are convinced that we can be justified in our perceptual beliefs because the interest that found inquiry into justification is an interest in a property we can have, and so we know that the thing that we are trying to learn more about when we study justification is something within our reach. We see something similar when we ask why we study knowledge. We study knowledge in part because we take it to be more valuable than similar mental states that fall short of it, and because in part we take it to be factive. If it wasn't valuable in that way, or it wasn't factive, it would not be nearly as interesting as we take it to be. So we know that there is a thing we are interested in and trying to learn about when we study knowledge, and we know that that thing is a factive mental state and valuable in a certain way. This is in part how philosophy gets started – there is something or some type of thing we are interested in, and we set out to learn more about it.

How does this solve the puzzle? We know first that we find certain things interesting (before we go on too far, I should warn you that I will be using words like “things” and “is” in the loosest possible sense for a while, as the things we are interested in may be impossible). For example, we know that we are interested in a mental state that is factive and is more valuable than some alternative candidate mental states. So we set out to learn more about this interesting thing. This means that we know that we have some target of inquiry, and we know something about it: it is that thing we find interesting, or it has those attributes that make it interesting. Given that we can be extremely confident in our knowledge of what we find interesting (more on this below), we can be extremely confident in this knowledge of our target of inquiry.

It'll be easier to explain the next step if we first make some (possibly false) assumptions. Assume for a moment that knowledge is transparently the target we are aiming at when we investigate the interesting factive mental state that is more valuable than the relevant alternatives. Since we know about our targets of inquiry by knowing what interests us, we know something about knowledge by knowing something about our interests. And we can be extremely confident in this knowledge. Similarly, if we are interested in some a property that our beliefs can have that plays some important role in the norms of belief formation, then we know something about a target of our inquiry; if this target were transparently justification, we'd know something about justification with an extremely high level of confidence. And that's the solution to the puzzle.

However, it may not be so transparent that knowledge is the target aimed at by investigation in the factive valuable mental state, or that justification is the target aimed at by inquiry into the property that plays the role in norms of belief formation. After all, “knowledge” refers to knowledge, but it may not be transparent what “knowledge” refers to; we might not be able to know what “knowledge” refers to, even in our own mouths, just by knowing what we are

trying to investigate. So, can it really be said that knowledge of our interests allows us to be all that confident in our beliefs about knowledge? It seems that this confidence should be limited by our confidence in some set of semantic theories. While this is a problem, and it undermines the confidence we can have in certain sentences, it doesn't undermine the important information that we gain in this way. By knowing what we are interested in when we conduct inquiry into a certain domain, we know facts about the target of inquiry in that domain. We may not know for sure that certain descriptions are descriptions of these facts, but that's not as important as knowing what interests us about some target of inquiry. What we intend to study is the thing that interests us, so information about that thing is more important than information about sentences. And we *can* describe what we know in neutral terms: we know with extreme confidence that the thing (certain) epistemologists intend to study when they study knowledge is a valuable, factive mental state. We can't be *as* confident that this is what "knowledge" refers to. But this just means that what we take to be the study of knowledge might not in fact be the study of knowledge, if language turns out a certain way.

To really see how this solves the puzzle I started this paper with, we have to do some precisification. If we formulate the puzzle as, "How do we know that the sentence 'Our perceptual beliefs can be justified,' is true?" then I haven't given as strong a solution as I would like (although below I will talk about how we can use knowledge of what we find interesting to improve our ability to use intuitions as evidence about such sentences). But I think we should restate the puzzle. The puzzle should really be formulated as, "How do we know that the thing epistemologists are trying to study when they study justification is something we can have?" That puzzle is solved by appeal to knowledge of what interests us. Once we've solved it, then

we have to ask, “Is studying justification a good way of studying the thing epistemologists are trying to study when they study justification?” I will return to this sort of question below.

None of this is to say that people *aren't* sometimes interested in whatever it is that “knowledge” (or “justification” or some other term) refers to. But I’d bet that many of us are not. If the linguistic facts are sufficiently up in the air (and I think they are), then it is epistemically possible that “knowledge” refers to something that is neither factive nor particularly valuable. In that case, a number of us would say, “Well, then I guess I wasn’t really interested in knowledge,” or “It turns out that I was interested in knowledge* all along,” or something of that sort. Knowing that knowledge wasn’t our target of inquiry would be interesting, but that wouldn’t make knowledge itself interesting. And the same applies to justification: if it turned out that “justification” really did refer to a state that required absolute certainty to have, this would not necessarily show that skepticism is true. Instead, a range of us would say, “I guess what ‘justification’ refers to isn’t what we were really trying to study this whole time.” And similarly, I would bet that many of you reading this, whose research is on some topic R, are not interested in whatever “R” turns out to refer to, but rather interested in some thing or set of properties that you take R to be or have.^{iv}

Our target of inquiry will often turn out to be the thing we take to be our target of inquiry, even if we don’t know that merely by knowing what we are interested in. Even if our interest is not in whatever “knowledge” might turn out to refer to, and we are instead interested in a factive, valuable mental state that we take to be knowledge, much of our use of “to know” does make sense if “to know” does track the thing we are interested in. So we have good reason to say that “to know” has multiple meanings and one refers to the factive mental state we intend to study. This points out another route to philosophical knowledge that knowing our interests gives

us: we can compare what our intuitions tell us to what we know about what we are interested in. This can give us evidence that our intuitions are actually about that stuff. And this is very useful.

In fact, it's so useful that I will devote the next section to a discussion of this usefulness. I will argue that philosophy that uses knowledge of what we find interesting in addition to intuitions will go better than philosophy using intuitions alone. After that, I will discuss some limits of knowledge gained in this way, and some concerns about the view I am articulating.

Intuition and Interestingness

The discussion of knowledge gained via knowing what is interesting was motivated by considering some of the limits on intuition as a source of philosophical evidence. We often cannot be as confident as we'd like about beliefs based on intuition alone. This is especially true for general beliefs, rather than beliefs about specific cases, and for beliefs about propositions that involve "philosophical concepts." A final concern about intuitions is that it's difficult for us to assess how reliable they are – to calibrate them (see, e.g. Cummins, 1998). I'm not going to claim that any of these make it so that intuitions do not, in general, justify beliefs. Rather, I want to argue that knowing what we find interesting can shore up these weakness of intuition, so that we are better off doing philosophy consciously using our interests and intuitions (rather than just tacitly starting from knowledge of our interests, which we already do).

One advantage of beliefs about what we find interesting is that they seem more reliable than beliefs based on intuitions. For now I'll assume that beliefs about what we find interesting are either certainly true or very close (more on this below). The fallibility of intuition, on the other hand, is widely acknowledged, and examples of mistaken intuitions abound. Another advantage of interestingness is that we can be much more confident about general claims on the basis of knowledge of what we find interesting than we typically can on the basis of intuitions.

We've already seen why we can't be terribly confident about general claims based on intuitions alone. Knowledge based on our interests does not have these problems. We often do find very general types of things interesting, and the belief that a type of thing is interesting is not defeated by apparent counterexamples, because the feeling that something is *uninteresting* is not good evidence that it is (I'll discuss this below as well).

This points out another advantage of interestingness over intuition. If intuition is to be evidence for the truth of theory T, then intuitions should not be based on our acceptance of theory T. What's more, they cannot be based on reflection or much thought; the conventional view of intuitions is that they are not based on conscious reasoning (see [removed for blind review]). This is not a feature of introspection. We should expect what we find interesting to be shaped by reflection and consideration. This is to some extent desirable: as we mature and become wiser from experience and better understanding of the philosophical landscape, what we find interesting should change, and change for the better. That some type of thing T is interesting based on our reflection on it does not undermine the fact that it is T that is interesting. Further, if our theoretical commitments cause us to find T interesting, and knowing that T is interesting is the basis of our knowledge about a certain domain of inquiry (that it is inquiry into T), this does not evidence circular reasoning. The truth of the claim about T that we end up with (that it is the target of inquiry) is based on reflection about T, but not based on the prior claim that T is the target of inquiry. Of course, such a process can be problematic to some extent – we should worry that those who reflect too much on a particular domain might have their interests “warped” by this so that they are not shared by any other human being (I'll discuss a related concern below) – but this possibility does not automatically invalidate reflection based interests as evidence.

The last advantage of interestingness over intuition alone is that the knowledge we gain through knowing what we find interesting allows us to calibrate intuitions – to learn how reliable they are in various circumstances. We don't need to calibrate intuitions in order to use them, since they give us justified beliefs without being calibrated in the absence of any attempt at calibration. Even so, if we are really committed to the truth, we have strong reasons to try to calibrate our intuitions (calibration might be epistemically supererogatory).

Calibrating intuitions involves determining how reliable they are in various circumstances, and what sorts of errors they are prone to. The obvious way to calibrate a source of data requires that one already knows, independent of the source, some range of correct answers to questions that source of data is supposed to answer. One then checks the source of data against the given answers, sees when the source is correct and incorrect, and uses induction to assess the overall reliability of the source. The better a sample of source-independent answers one has, the better a job of calibration one can do. Knowing our interests allows us to calibrate intuitions in this obvious way to some limited extent.

To illustrate, imagine that we are interested in some epistemic property that, among other things, it is better to have than to not have (all else being equal); let's call that "knowledge." We wonder, "Are our intuitions that seem to be about knowledge really about what we are interested in?" To answer this, we look at some intuitions about cases. For example, we intuit that a person in Fake Barn Country, an area filled with many fake barns and one real barn, who looks at the real barn and believes "That is a barn," doesn't *know* that that is a barn, whereas a person in Real Barn Country, filled just with real barns, does know "That is a barn," when they look at a barn. Is the person in Real Barn Country also in a better state than the person in Fake Barn Country? We have reason to think so – they are likely to continue to form similar true beliefs

about barns, whereas the person in Fake Barn Country is likely to go on to form false beliefs about other (seeming) barns. So it seems that at least some intuitions about knowledge are intuitions about a state that is better to have than not (all else being equal). So this is some evidence that a) the intuitions about these Barn Country cases are really about what we are interested in, and b) that *other* intuitions about knowledge are also about what we are interested in. Of course, these two intuitions are only very weak evidence for the latter.

Typically there are just a few reasons why we are interested in some topic, so we won't know all that much about that topic in this way, and so can only check some of our intuitions directly against our interests. This allows only weak induction to some estimate of reliability. There are, however, other approaches to calibration that make better use of our knowledge about what we find interesting.

Checking our intuitions for consistency is sometimes suggested as a way to calibrate them (see e.g. Bealer 1998). If our intuitions are consistent, it is suggested, we have evidence that they are reliable. However, consistency is not by itself evidence that intuitions are reliable. We can see this by seeing that *coherence* of beliefs is not by itself evidence that they were reliably produced. Propositions are consistent when they don't contradict each other. They are coherent when the truth of each is evidence for the truth of the others. It has been shown that the coherence of a set of propositions is not evidence for the on-bulk truth of the propositions in that set (Shogenji, 1999), which entails that it is not evidence of the reliability of the process that forms them.^v From this it follows that consistency of beliefs is likewise not by itself evidence of a reliable basis. However, if we know that some members of a set of coherent beliefs are true, *then* we have evidence that most are true, and thus evidence that the process that produced the set is reliable. Knowledge of our interests gives us a way to know if some of our intuitions are

correct (using the method discussed above). If these correct intuitions are coherent or incoherent with others in a given domain (ethics, epistemology, etc), we then have even more evidence about their reliability.

Again, I will illustrate how this might work. Imagine that we know that we are interested in some property that is shared by the torture or murder of innocents; let's call this "moral wrongness." We also have the intuition that torturing or murdering innocents is morally wrong. This is some evidence that our intuitions about moral wrongness track the property we are interested in, but by itself is not very strong evidence. We go on to have an intuition that a person who scams the elderly out of their retirement savings for no good reason – call this person Bernie – has done something morally wrong. Should we trust this intuition? That is, should we see it as identifying an act that really has the property we are interested in? It would be nice to have more to go on here than just the fact that our intuitions about murder track our interests. Let's say also that we have the intuition that harming innocents is morally wrong. This is coherent with our intuition about murder – if one of them is true, then the other is more likely to be true. Further, it is also coherent with our intuition about Bernie's behavior. Since we know that our intuition about murder is correct given our interests, we have more evidence that the intuition about Bernie is correct. The more coherent intuitions we have about moral wrongness, the more evidence we can have that our intuitions track what we are interested in. And we can make inferences in the other direction when our intuitions are incoherent.

However, there are often clusters of intuitions on a topic that are coherent with themselves, but merely consistent with other intuitions on that topic in other coherent clusters. For example, many of our intuitions about perceptual knowledge are mutually supporting – if I have the intuition that I know that this table exists when I see it, and that you know the table

exists when you see it, these look coherent – but seem merely consistent with our other knowledge intuitions, such as intuitions about Gettier cases. If we have some cluster C of coherent intuitions on a topic that are merely consistent with other intuitions on that topic, and we have no interest-based knowledge that supports the truth of intuitions in C, then we have very little evidence that intuitions in C are reliable. Because they are merely consistent with the intuitions we can check, the reliability (and likely truth) of the intuitions we can check does not show that the ones we can't check are likely true. Induction from one group to the other is very weak at best, as they are different enough that one sample will not represent the other.

There is a way of calibrating even these intuitions; that is, intuitions that aren't directly checkable against our knowledge of what we find interesting, and are only consistent but not coherent with intuitions that definitely match what we find interesting. We start by investigating how intuitions are generated, looking into what information our intuitive faculties are sensitive to, what information they are not, how that information is used, what factors cause biases. This investigation is likely an empirical one. Once we learn how our intuitions are generated, we ask, "Should something that works in *this* way be a good source of evidence about such-and-such a philosophical topic?" This requires thinking about whether something that works in *this* way – whatever way intuitions turn out to work – has the right kind of access to the right kind of facts, and what sorts of errors we should expect of it given the demands of the domain in question. How do we know what the right kinds of fact are, or the demands of the domain in question? We figure out what intuitions would have to track to tell us about what we are interested in. For example, let's say you are interested in ethics because you are interested in what is good for us. Because of this, you might know that making accurate ethical judgments requires sensitivity to certain kinds of social facts, such as how humans are be affected by various actions. So, if our

intuitions are or are not sensitive to social facts, this will tell us something about them as a source of evidence about ethics. Similarly, if we find out that we are all biased to see the results of our own actions as better than they really are, we have found a source of error ethicists need to worry about. Since we know that our intuitions *are* sensitive to at least some social facts relevant to what benefits or harms others, we know that they meet at least one necessary condition for telling us what is right or wrong. At the same time, since we also know that we are biased to see ourselves as better than others, we also know that we should somewhat discount moral intuitions that either are about or obviously reflect on our own actions.

Let's briefly look at what these three approaches to calibrations might tell us when different philosophers have diverging intuitions. In some cases, such diverging intuitions are the product of something like merely verbal disagreement: philosophers A and B both use "justification" to refer to something each finds interesting, but the thing each of them is interested in that they call "justification" is quite different. In such cases, if their intuitions on justification diverge, this is not much of a problem. Put such cases to the side, and consider two philosophers who are both interested in some single specific thing they can (in some sense) point to – for example, they are both interested in the property paradigmatically had by the intentional murder of innocents, which we will call "moral wrongness." Both also intuit that murder is wrong, and so both know that at least that intuition about wrongness tracks what they are interested in. However, they have conflicting intuitions about the wrongness of other cases, such as trolley cases. Since they don't know, independent of their intuitions, that acts in these trolley cases do or do not possess the property they are interested in, how can this disagreement be resolved? The calibration I've discussed won't solve every such problem, but it often will. To give one example of how, if one of them has some incoherent intuitions about wrongness

(incoherent, that is, with their intuition about murder, which matches what they find interesting), but the other doesn't, then we have evidence that the coherent philosopher's intuitions about the trolley case are really about the sort of thing they are both interested in. Calibration can also help resolve conflicts of intuitions between philosophers who are both interested in whatever fits some abstract description. For example, take two philosophers who are both interested in whatever property is necessary for moral responsibility, which they call "free will." Imagine that each has very different intuitions about who has free will. How can we resolve this using the methods of calibration we have discussed? We know that, for intuitions about the interesting property to be reliable, they should be sensitive to whatever is needed to detect moral responsibility. To some extent we know what this is and can tell that one or the other philosopher's intuitions are insensitive to this. To give an extreme example, if one of the two were a sociopath, we would expect that she would not be able to properly detect moral responsibility, and thus that her intuitions about what they are both interested in can be discounted.

The upshot of all of this is that using introspection as a starting place for philosophy allows us to do work we could not have done otherwise. It gives us more confidence in many sorts of claims than intuitions alone can. And it allows us to calibrate our intuitions both by checking their coherence with what we know via introspection, and by checking to see if they are formed in a way conducive to learning about our topic of inquiry.

Some Limits of Interestingness-based Knowledge

I have argued that we can gain some very useful knowledge via determining what we find interesting, and that this route to knowledge does a better job of explaining our confidence about certain claims than do intuition and reasoning alone. But there are limits to what we can know in

this way. To see what they are, it will be helpful to start by asking, “What can we find interesting?”

This is a topic that requires further investigation, but it seems that we can be interested in anything in the class of things we can in some sense comprehend or be aware of. We seem to sometimes have what we might call “*de dicto* interests:” we might be interested in whatever it is that fits a certain description (e.g. “a factive propositional attitude that has some value”). Among the things that can be *de dicto* interesting are whatever it is that some word refers to, or whatever it is that someone else is interested in, or whatever it is that matters. We also seem to have *de re* interests: we can be interested in a specific thing, often demonstratively indicated (“this thing” or “the property shared by these objects”). These can be and often are combined – I might be interested in the factive mental state that I am in now as I truly see a glass of water in front of me, which has as its object the glass of water itself.

As already discussed, unless we are interested in whatever some word “W” refers to, our confidence that we are interested in W, and thus that W has the properties we are interested in, is limited in part by our confidence that we know enough to determine the referent of “W.”

If we have *de dicto* interests, or hybrid *de re / de dicto* interests, we can’t know that what we are interested in exists, or is “had” by any thing (in the case of properties), just in virtue of knowing what is interesting to us. Consider someone interested in that thing necessary for moral responsibility (who takes this to be free will). It could turn out that this is necessarily not possessed by anyone. Or, if one was instead interested in the property necessary for moral responsibility that is possessed by human beings, one might find out that such a property is impossible. Such a discovery does not show that what we thought was interesting in fact was not

interesting, however. Impossible things can still be fascinating (and we can certainly wish that they were possible), and it's worth learning that things we care about deeply don't exist.

Sometimes our interests are not very informative. I am deeply interested in ethics. But I find that when I reflect, all I know about these interests is that I'm interested in properties clearly had by a few paradigmatic cases of right and wrong. Close scrutiny of these cases and non-interest-based evidence, such as that from intuitions or observation or science, might teach me quite a bit, but my interests by themselves don't tell me much about ethics. And this will be true for many of us and many, perhaps most, topics. There is, however, something good about having under-developed interests. The more developed they are, the more likely what is interesting to us will not be interesting to others, or not exist, or not be possible.

There are three more possible limitations to what we can know by considering what we find interesting that require enough discussion to merit their own sections. One is that we might not be able to know that a target of inquiry is *worth studying* by knowing that we find it interesting; if that's right, then we have to ask whether starting inquiry with knowledge of what we find interesting is appropriate. We also have to wonder how confident we should be that we are really interested in what *seems* interesting. But, before I address these, I will discuss the question of whether know what we find interesting can really tell us about non-conceptual and non-linguistic stuff.

Interestingness and the extra-mental

Many philosophers want to study things that are outside our minds – things other than concepts or the meanings of words. I've suggested that philosophy does, and should, start partly from knowledge of what we find interesting. But this knowledge is knowledge about something

in our heads – our own feelings about what is interesting. How can this tell us about things *outside* our heads?

If we have some purely *de re* interest – e.g. interest in *this* thing, or in the property shared by *these* two things – then this by itself tells us something about the world – that this is interesting to us – but that’s not very much information. If our interests are *de dicto*, they can also sometimes tell us about the world outside our heads – e.g. if my interest in what I take to be epistemic justification is an interest in some property that my beliefs can possibly possess, then I can be confident that my beliefs can be justified. But as already discussed, the more detailed, and thus more informative, our *de dicto* interests are, the less confident we can be (based on these interests alone) that they are instantiated in the world. So it seems that knowledge of what we find interesting alone does not tell us much about the world outside our heads.

But knowing what we are interested in is just a starting place for philosophy. When we combine it with other sources of information, like intuitions, we can really learn about things outside of our heads. I will briefly show how this can work. I’ll use “toy” cases to keep the discussion simple. Things will of course be much more complex for real philosophical inquiry.

The clearest cases where knowing what is interesting, plus having intuitions, can tell us about the non-conceptual and non-linguistic world is when we have *de re* interests. We find that we are interested in *this* and things like it. We then see that we have the intuition that *this* is an instance of T. We go on to have intuitions about other things being T, and perhaps some intuitions about T-ness. If these are all coherent, we have some evidence that our intuitions about T are intuitions about a sort of thing we are interested in. We then study T things other than *this*, perhaps because they are easier to study. This allows us to learn about T, and learn about *this*. One can see this in my discussion of calibrating intuitions about moral wrongness

above. We might start with a *de re* interest in a property possessed by a few paradigmatic wrong acts, such as the torture of innocents. We find that we intuit that these acts are wrong, intuit that causing unnecessary harm is a sufficient condition for something being wrong, and intuit that unjustified theft is wrong. These are coherent, so we trust these intuitions. It is harder to study the paradigmatic cases of wrongness that we are *de re* interested in, such as torturing innocents for no reason, as they are very rare, and easier to study cases like theft that we only know are wrong via intuitions that are coherent with our interests. In doing so, we have learned something about the moral properties of the world, as we started with something in the world and built out from there.

How can *de dicto* interests plus intuitions help us learn about extra-mental things? A *de dicto* interest is an interest in stuff that fits a certain description – e.g. stuff that possesses traits X, Y, and Z. Our intuitive judgments will typically be categorization judgments – that an object belongs in category C. Let's say we've calibrated intuitions about C, and found that they are about what we are interested in – when something is intuitively C, it typically has traits X, Y, and Z. In many cases, there will be things that we cannot directly tell have X, Y, and Z. Having the intuition that something of this sort is C will then tell us that it also has traits X, Y, and Z. Further, by studying things in the world that intuitively are C, we may learn that they all possess some further traits beyond X, Y, and Z. This teaches us something about the stuff we are interested in. To illustrate, consider Sally, who is interested in a factive mental state that people are typically in when their beliefs are justified. Sally calibrates her intuitions and finds that intuitions that someone *knows* something track this *de dicto* interest. Sally then considers some people who are in factive mental states, but where it is unclear whether their beliefs are justified – e.g. people who have reliable ESP but don't know it. Sally finds that, intuitively, these people

don't get achieve knowledge via their ESP. Sally has learned something about the world – she's learned something about who is in an interesting mental state. Now Sally considers some Gettier cases, and learns that not everyone who is in a factive mental state where their beliefs are justified *knows*. Sally have learned by this something about the world – she's learned something about people in Gettier cases – and she has learned something about the mental state she's interested in – that it has a fourth condition that was not part of her *de dicto* interest. And this thing she has learned will then help her learn more about who (in the world) does and does not know.

Information about what we find interesting is information about ourselves, and partly about the world, that we can be very confident in. It is also information that we can use to calibrate our intuitions. This information, and what we learn using calibrated intuitions, can tell us quite a bit about things that are neither concepts nor the meanings of words.

Importance versus interesting-ness

Even if I'm correct that knowledge of what we find interesting is an unacknowledged starting point for philosophy, and that using this knowledge has advantages over not using it, we should still wonder if what we learn in this way will be anything *important*. Perhaps philosophical investigation that starts with what one finds interesting is a sort of objectionable navel gazing.

This worry only gets off the ground if “important” is something beyond “seems important to me,” or “is something I care about.” Otherwise, knowing what is interesting just would be knowing what is important. So let's read “important” (somewhat loosely) to mean something like “what *really* matters,” (you may have to pound the table when you say “really” in order to fully comprehend what I have in mind).

Let's grant that what is interesting need not be what is important. Even so, we need a starting point for philosophical inquiry, and as of now we can't use our knowledge of what is important as a starting point. The question of what is important is a substantive philosophical question, one that we are far from answering. Nor is it clear that we are reliably capable of recognizing what is important. In order to investigate those things that are important, we first need to investigate importance itself. We certainly should do this, and many people are, but it's unreasonable to ask all of philosophy to wait until they finish.

Why not say that every target of inquiry *should* be something important? Since sometimes we aren't interested in important things, this would suggest that our interests should not fix our target of inquiry. But this is not called for. Sometimes an investigation is worthwhile even if it turns out to not be investigation into something that is important. Often we are interested *de re* in something without knowing if it is important, or we have *de dicto* interests where the description of what we find interesting does not include require that it is important.^{vi} On subsequent investigation, and after developing a better understanding of what is important, we might find out that the thing we were interested in is not really important. This is worth knowing, and the investigation we conducted was worth doing, if for no other reason than that we discovered that something we found fascinating is not all really worthwhile. For example, many years ago I discovered that I had strong libertarian and incompatibilist leanings about free will. In retrospect, I would say that what interested me about what I *took* to be free will was something that could not be had in a world where determinism is true. However, it might turn out that that thing doesn't really matter – it might turn out that the compatibilists are right, and moral responsibility and meaningfulness of choice are compatible with determinism. It would be worth knowing that what I was interested in didn't matter, and that might require investigation in

the very thing I was interested in. Stipulating that all targets of inquiry are important things would prevent us from learning this sort of thing.

As an aside, my own view is that things are important because, and only because, they matter (in some sense) to people and other sentient creatures. There are things that matter even though they don't promote obvious goods like health, survival, or happiness. What is interesting to me, a sentient creature, is at least a decent candidate for being important, even if investigating that thing is irrelevant to my well-being. If others find this interesting, then it is even a better candidate for being important. And so starting philosophy from what is interesting is a decent way of getting to what matters; as it has other advantages that cannot be had by starting entirely from the pursuit of what matters, it makes a good starting place for philosophy.

To close this section, I want respond to worries voiced recently by Daniel Dennett, but shared by a number of philosophers (especially very naturalistically inclined ones), about analytic philosophy. Dennett argues that the best analytic metaphysicians can do is "sophisticated naïve anthropology," which carefully charts out the terrain of what seems true to us, where "us" is taken widely to include non-philosophers (Dennett, forthcoming). Ultimately, though, Dennett says "To me, [this] looks more and more like professional make-work, an artifact of our reasonable but ultimately optional desire for systematicity rather than a deeper mystery in need of solving." (Dennett, forthcoming, 8) His worries are *prima facie* also worries about the approach I've been discussing.

What does count as a deep mystery worth solving? Should we see, for example, the nature of any old natural kind as a deep mystery? I hope not. In the absence of giving us a fully developed theory of what does and does not matter, Dennett's counsel seem to be to look to what matters to people beyond the immediate circle of people who work in your area, and especially

what matters to smart laypeople (Dennett, forthcoming, Dennett, 2006). But if that's good advice, then the worry doesn't seem to be about looking into what is interesting. Rather, it seems to be that what is matters to any single philosopher might not really matter. And I think that is a reasonable concern. The answer to this concern is not to ignore your own interests, however, at least not from the start. Nor is it just to go talk to others. After all, it can be very difficult to express our interests in plain language. For that reason, when we talk to others, especially non-philosophers, about what they find interesting, it will be extremely helpful to have tried to map out what *we* find interesting, and how to express that. This will make it much easier to communicate with others and navigate their conceptual landscape. Note that this is not *just* anthropology. What we are interested in matters: if something can be a deep mystery without being interesting to anyone, then I have to admit that the (potential) fact that philosophers aren't investigating the deep mysteries doesn't seem a particularly significant concern to me.

Interest, Introspection, and Confidence

One of the selling points of knowledge of what is interesting as a starting point for philosophy is how reliable it is. But how reliable is it really?

First we should get straight what I am claiming. Here are some options:

- a) If we are interested in X, then we feel/believe we are interested in X and are (close to) certain that we are interested in X.
- b) If we feel interested in X, then we are (very likely) interested in X and we are (close to) certain that we are interested in X.

I don't advocate (a). I imagine that there are things we are disposed to be interested in but where that interest has not manifested itself (e.g. we might not have considered the topic yet). I endorse some version of (b), but further investigation is required into the question of how likely we are to

be right and how certain we can be. (B) talks about our feelings of interest rather than our beliefs about what is interesting because it is much harder to be mistaken about the feeling of interest than it is to have mistaken beliefs about what we are interested in. We might, for example, interpret some of our behavior as signaling interest in a topic (see below) and thus convince ourselves that we are interested in something despite absence of a feeling of interest.

Note that (b) is only about feelings of interest, and not feelings that something is uninteresting. This is because we often cannot be very confident that when we feel uninterested in something we actually are. For example, if I showed you Clark Kent, you might sincerely say, “I am not interested in that guy.” However, when you see Superman, you might feel an interest in that person – not just the external appearance, or the powers, or the cape, but the person himself. If that’s right, then you are really interested in Clark Kent as well, as they are identical. It seems that anything can in principle be interesting when presented one way, but uninteresting when presented another; if I claim that what feels interesting is interesting, then I cannot claim that we can be certain that what feels uninteresting is.^{vii}

Our knowledge of what we find interesting depends in part on how trustworthy introspection is, as we learn about our feelings of interest through introspection. Introspection has historically been seen as infallible, but that’s not a widely accepted view today. One reason to doubt introspection is that much of our behavior and cognition is generated automatically or unconsciously, and so we cannot reliably introspect as to why we did or thought these things. Since I’m only saying that introspection is close to infallible about occurrent feelings, which are the sorts of things that are inherently conscious, this is not much of a worry for me. However, we also often make things up about our own inner life; this is called “confabulation.” If we can

confabulate about our interests, then it seems that we cannot entirely trust our experience of something as being interesting.

When one looks at the data on confabulation, it is typically data of two sorts. One sort is data from severely impaired people (e.g. those whose left and right hemispheres have been severed). I don't know how seriously we should take data from the severely impaired as casting doubt on *our* ability to introspect. This data seems at best to give somewhat Cartesian reasons to doubt: it may just reduce us from *absolute* certainty to the next best thing. Another sort of data on confabulation is data about ordinary people unknowingly inventing false explanations for why they behave as they behave, or make the judgments they do. People in these studies invent feelings that they never had, but they report these as explanations for their *prior* acts. So they aren't making up feelings they have right now, nor are they making mistakes about their own occurrent feelings. To the extent that our knowledge of what we find interesting is inferred from our previous behavior and judgment, this data is worrisome. But as long as we can check to see if we actually have a feeling of interest now, we need not take it too seriously. Furthermore, Fiala and Nichols (2009) report that those who confabulate tend to be less confident in their confabulated explanations that we expect from people with direct access to their inner states. They suggest that confabulated feelings and mental states are subjectively distinguishable from directly experienced ones, although they don't make claims to conclusive proof. Put together, data on confabulation and automaticity at best undermine our claims to *absolute* certainty about our own feelings of interest, but only as long as we see something like Cartesian-style doubts as appropriate reducers of credences. It still seems that we can be very confident in our introspection about our occurrent feelings of interest.

Note also that at least some philosophers who take data on automaticity and confabulation very seriously, and are quite skeptical about our self-knowledge, still do not take this data as undermining the sort of self-knowledge that I am claiming certainty about. John Doris is a good example of a philosopher in this camp. The sort of introspection we are discussing is a lot like immediate conscious perception, and Doris does think that we have privileged access to this, even though we lack such access to much of the rest of our mental lives (so much so that Doris discusses his skepticism about our personhood on this basis) (see Doris, 2009, esp fn 21). Peter Carruthers (2009) gives a recent argument for a relatively skeptical view of introspection. In the course of this, he surveys four main (empirically based) theories of how we get self-knowledge; three of them require that we have some sort of reliable and privileged access to our own mental states. Carruthers defends a fourth type of view, but the version of this that he claims is best supported by the available data still requires that we have privileged access to some of our mental states: “allowing [our] mindreading system to have access to visual imagery, proprioceptive data, and emotional feelings is pretty much mandatory once we buy into a global broadcasting architecture...” (p8) (and Carruthers does buy into such an architecture). On this account, we have direct access to our own feelings, but only inferential access to our propositional attitudes. We also have direct access to our internal speech. This allows that we can have extremely reliable knowledge of what we find interesting: we can tell that we are considering some list or another of traits and whether or not we feel that that list is a list of interesting traits. This suggests that those who are skeptical about our access to our own mental states tend to not be skeptical about our access to our own experiential states.

Now, I cannot deny that there are some reasons to worry about our ability to discern what we find interesting. There are still empirical questions to ask here. It may be that our feelings of

interest can be swayed by irrelevant situational factors. And we may be able to tell that we are interested in something, but not quite what it is.

This last worry is raised by considering arguments for externalism about content. I don't have the space here to address the large literature on externalism and self-knowledge. I would suggest that that most versions of externalism are not a concern for my view in a fairly wide range of cases. As long as the externalism in question allows for privileged access to our own phenomenal states, then if we have a *de re* interest in some object, or if we are interested in whatever some word refers to, then we can know this for certain. *De dicto* interests might be knowable on such a view as well. Imagine that I think about whatever type of thing has the properties P, Q, and R, and this thought is expressed by some sentence "spoken" in my head. I feel that this type of thing is interesting. I have extremely reliable access to the feeling of interest, and the "hearing" of the thought. I can't in this way know what P, Q, and R *are*, but why can't I know that I am interested in whatever has whatever they are? And in other cases we might know exactly what properties we have in mind (perhaps by imagining them in some way not mediated by language). We might not know what properties those were, because we wouldn't automatically know what words referred to those properties. But we'd still know that we were interested in whatever thing has whatever those properties end up being. There is much more to be said here than I have room to explore, but I do think that these sorts of arguments can address at least the milder (but most widespread) versions of externalism. This does require more investigation on my part, and on the part of those who find the rest of my view plausible.

Conclusion

Look around the philosophical landscape: there are core, widely held beliefs, often about very general claims, that many of us with many different approaches to philosophy are very

confident in. This cannot be explained, I've argued, as justified by reasoning with the sort of evidence philosophers normally appeal to. The charitable interpretation is that there is another sort of evidence philosophers use, even if they only use it tacitly. This is knowledge of what we find interesting. Such knowledge allows us to know about our targets of inquiry with a great deal of confidence. What's more, it allows us to check our other sources of evidence, to see whether they are in fact good sources of evidence about our targets of inquiry. Significant questions remain about this source of evidence and the role it can play in philosophy; I hope that I have given some of you reasons to try to answer these questions.

Acknowledgements

[removed for blind review]

References

Alston, W.P. (2005) *Beyond "Justification:." Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Bealer, G.(1998) *Intuition and the Autonomy of Philosophy,*" in DePaul, M.R. & Ramsey, W. (eds) *Rethinking Intuitions*. Lanham, MA: Rowman & Littlefield Publishers.

Bishop, J.D. & Trout, M. (2005) *Epistemology and the Psychology of Human Judgment*. Oxford University Press.

BonJour, L. & Sosa, E. (2003) *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Wiley-Blackwell.

Carruthers, P. (2009). How We Know Our Own Minds: The Relationship Between Mindreading and Metacognition. *Behavioral and Brain Sciences* 32, 121-138.

Cohen, L.J (1986) *The Dialogue of Reason*. New York: Oxford University Press.

Cummins, R. (1998) Reflections on reflective equilibrium. In DePaul, M. & Ramsey, W. (eds) *Rethinking Intuition*. Rowman & Littlefield Publishers.

Dennett, D. (forthcoming) Kinds of things - Towards a bestiary of the manifest image. In Ross, D., Ladyman, J., Kincaid, H. (eds) *Scientific Philosophy and Metaphysics*.

Dennett, D. (2006) Higher order truths about chess. *Topoi*, 25, 39-41.

Dechêne, A., Stahl, C., Hansen, J., Wänke, M. (2010) The truth about the truth: A meta-analytic review of the truth effect. *Personality and Social Psychology Review*, 14, 238-257.

Doris, J. (2009) Skepticism about persons. *Philosophical Issues*, 19, 57-91.

Erskine, A. (2003) *A Companion to the Hellenistic World*, Malden, MA: Blackwell.

Fiala, B. & Nichols, S. (2009) Confabulation, confidence, and introspection. *Behavioral and Brain Science*, 32, 144-145.

Goldman, A. & Olsson, E. (2009) Reliabilism and the value of knowledge. In Haddock, A., Millar, A., Pritchard, D. (eds) *Epistemic Value*. Oxford University Press.

Hazlett, A. (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research* 80, 497-522.

Huemer, M. (2001) *Skepticism and the Veil of Perception*. Rowman & Littlefield Publishers.

Kunda, Z. (1990) The case for motivated reasoning. *Psychology Bulletin*. 108, 480-489.

Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.

PhilPapers.org (2009) Preliminary survey results. PhilPapers.org, <http://philpapers.org/surveys/results.pl>, last accessed September 29, 2011

Shogenji, T. (1999) Is coherence truth conducive? *Analysis*, 59, 338-345.

Sosa, E. (1998) Minimal intuition, in DePaul, M.R. & Ramsey, W. (eds) *Rethinking Intuitions*. Lanham, MA: Rowman & Littlefield Publishers.

Zagzebski, L. (2003) The search for the source of epistemic good. *Metaphilosophy*, 34, 12-27.

ⁱ Making such an argument (well) would be quite difficult, as the sample argued from should be random and large, and developing a method for randomly choosing among one's actual beliefs about the external world is not a trivial task.

ⁱⁱ For example, Alvin Goldman and Linda Zagzebski apparently agree on this, despite having radically different views of epistemology in other ways (Goldman & Olsson, 2009, Zagzebski, 2003).

ⁱⁱⁱ A typical response to such data is that "to know" is ambiguous, but there are alternative explanations available: e.g. if justifiedly believing p almost always meant that p were true, we'd typically infer that p from x knows that p , and that x does not know that p from $\sim p$, whether or not knowledge were factive. This would explain the data without "to know" having multiple meanings, or being factive. Even if this is not a great explanation, it undermines our confidence in the "multiple meanings" explanation. For more debate on this, see Hazlett, 2010.

^{iv} One only has to look to certain revisionist moves in epistemology to see this in action. Alston (2005) and Bishop and Trout (2005) both call for abandoning the study of justification, and Kvanvig (2003) for abandoning the study of knowledge, yet all see themselves as engaged in epistemology. This, and their apparent expectation that their arguments might sway other epistemologists, makes the most sense if they see the target of inquiry of epistemology not as whatever knowledge or justification are, but as aimed at some interesting thing or set of things (that we have traditionally but mistakenly taken knowledge and justification to be).

^v If a set of propositions is coherent, it is very unlikely that half are true and half are false, because they mutually support each other. But, it is just as likely that most are false as it is that most are true. This is because their coherence means that, if one is true, the others are likely to be true as well; but it also means that if one is false, the others are more likely to be false as well.

^{vi} Sometimes we *are* explicitly interested in what is important; for example, epistemologists might not just want to study any old factive mental state, but rather the factive mental state that is important for inquiry. In this way we can know that our target of inquiry, if it exists, is important.

^{vii} However, we *can* recognize that something is not the target of one of some specific interest of ours: if we are interested in things with property *p*, then anything that clearly lacks *p* can be recognized by us as not the target of that interest (although it might be the target of some other interest).