

For: *Epistemic Evaluation: Point and Purpose in Epistemology*, eds. John Greco and David Henderson, New York: Oxford University Press

Stephen R. Grimm
Fordham University

Knowledge, Practical Interests, and Rising Tides

The idea that our judgments about knowledge are sensitive to what is at stake, from a practical point of view, can be found in a variety of classic works in epistemology. Locke, for instance, seems to hold that more is required for knowledge when the topic at hand is “weighty” than when it is not,¹ and one of the upshots of W.C.K. Clifford’s famous example of the ship owner seems to be that evidence which might under normal circumstances count as “sufficient” or “good enough” for knowledge will no longer count as such when something significant is at stake.² When a mistake would lead to something like the loss of human life, the thought seems to be, only an exceptionally well-supported belief will count as good enough for knowledge.

Looking back further into history, we can also find indirect support in Descartes for the idea that our everyday judgments about knowledge are sensitive to practical stakes. Thus Descartes explicitly tells us, in both the *Discourse* and the *Meditations*, that for the purposes of his project he will be bracketing consideration of his practical goals and concerns and will instead be focusing only on what we might think of as his “purely epistemic goal” of finding the truth.³ The implication seems to be, then, that on his normal or everyday way of thinking—the way of thinking that he learned from those around him—considerations of practical stakes naturally influenced his judgments about knowledge. It therefore took a special act of will on

¹ According to Locke, for instance, “it is very wrong and irrational way of proceeding, to venture a greater good for a less, upon uncertain guesses and before a due examination be made, proportionable to the weightiness of the matter, and the concernment it is to us not to mistake” (2.21.66); and again, “where the assent one way or the other is of no importance to the interest of anyone ... there 'tis not strange that the mind should give itself up to the common opinion, or render itself to the first comer” (1975: 717). Owens (2000, esp. ch. 2) cites further passages from Locke along these lines. (Thanks to Josh Orozco for the pointer.)

² In Clifford’s example, the ship owner takes his boat to be seaworthy on the basis of (what seems like) good inductive evidence which nonetheless falls short of being “good enough”—and apparently, in light of the dramatic cost of being wrong. See Clifford ([1877] 1999).

³ See, for example, *Discourse* (4.1) and *Meditations* (1.11). For discussion of the importance of this move, see Williams (1978, ch. 2).

Descartes's part to exclude practical concerns from his judgments about knowledge, and to focus instead on a new, purified way of approaching the issue.

What is the point of dipping into history like this? For one thing, it is to try to undercut some of the more territorial of the recent concerns about "pragmatic encroachment" in epistemology, where the language of "encroachment" is obviously meant to suggest something that is not only unwelcome, but also new. If these brief historical claims are to the point, the idea that practical stakes can influence what it takes to know is not only far from new, but has a plausible claim to be the commonsense view of things. What should count as a departure from the norm, from this perspective, is something like the Cartesian approach—an approach that attempts, in an artificial, revisionist way, to set practical concerns to one side when it comes to our judgments about knowledge. Risking a little paradox, we might say that the Cartesian approach amounts to the encroachment of the "purely epistemic" on epistemology. At the very least, it suggests that a special effort is required to exclude practical concerns from our thoughts about what it takes to know.

That said, what counts as the commonsense or traditional view of things might still end up being the mistaken view of things, and there are two recent arguments that suggest that viewing knowledge in this way yields the wrong results. The first is that if we suppose that knowledge is tied to our practical interests, then it seems that what it takes to know should change radically from case to case as our practical interests change. But what it takes to know does not change radically from case to case; instead, it seems to remain remarkably stable across contexts. Even with respect to questions that we could not care less about, or relative to which the practical costs of being mistaken are basically nil, knowledge still requires us to have good evidence, or to have beliefs grounded in reliable sources, and so on. In short, and as several people have recently pointed out, it decidedly does *not* seem to be the case that one can know

more by caring less.⁴ And yet this is precisely what the idea that knowledge is tied to our practical interests seems to predict.

Second, there seems to be no fully satisfying way of explaining *whose* practical interests matter. To say, in a vague way, that knowledge is tied to “our” practical interests and concerns is one thing, but recent attempts to be more precise about the extent of this “our” have all met with serious problems. Thus John Hawthorne and Jason Stanley, for instance, are clear that their “subject-sensitive” view needs to explain away our judgments in some central cases, and they are quick to claim that “attributor-sensitive” views are in the same boat.⁵ But if all of the familiar ways of sharpening this basic idea are forced to explain away our judgments at some point, then perhaps it would be best to make this move at the outset and claim that our sense that practical stakes can bear on knowledge is just mistaken—perhaps to be explained more by facts about conversational appropriateness (or the like) than about the real nature of knowledge.

We can think of the first of these problems for the idea that knowledge is sensitive to practical interests as *the stability problem* and the second as *the whose stakes? problem*. In what follows I will argue that both problems can, in fact, be addressed in roughly the same way. More exactly, I will suggest that by first clarifying *the whose stakes? problem* an answer to *the stability problem* naturally falls out. Looking ahead, the idea I want to defend is that neither of the usual ways of sharpening the basic idea that knowledge is sensitive to “our” practical interests and concerns is quite right: thus neither the *subject’s* interests nor the *evaluator’s* uniquely determine the thresholds relevant to knowledge. Instead, I will argue that the thresholds relevant to knowledge are sensitive *both* to the interests of the subject *as well as* to the interests of the evaluator.⁶ What is more, I will argue that the thresholds are sensitive to the practical stakes of certain third parties as well.

⁴ See, for example, Schaffer (2006, pp. 96-97), Cappelen and Lepore (2006, pp. 1044-45), Schiffer (2007, pp. 194-95), Russell and Doris (2008), Zagzebski (2009, ch. 2), and Reed (forthcoming).

⁵ See Hawthorne (2004, ch. 4) and Stanley (2005, p. 11 and ch. 5; cf. 2007, p. 210).

⁶ I assume for simplicity here and throughout that the subject and the attributor are not one and the same person.

With some qualifications to be introduced later on, I will therefore defend what I will call a *rising tides account* of the way in which practical interests bear on knowledge—as a first approximation, an account on which rising costs *either* for the subject *or* for the evaluator *or* for certain third parties can raise the thresholds relevant to knowledge.⁷ Whatever one might think of “rising tides” views as accounts of economic welfare, as a view about the sorts of factors that help to determine what it takes to know, it seems just right.

1. Some Clarifications

To make this case I will first try to clarify some central terms, and then offer a few preliminary reasons for thinking that what it takes to know can depend on practical stakes.

Following Stanley’s label and basic idea, and in keeping with similar thoughts by Fantl and McGrath, let us think of *intellectualism* in the following way:

Intellectualism: whether a true belief amounts to knowledge depends exclusively on truth-connected factors.⁸

We can then try to capture the idea that knowledge is sensitive to practical stakes under the label *practicalism* as follows:⁹

Practicalism: whether a true belief amounts to knowledge depends on certain *non*-truth-connected factors—in particular (it seems), it depends on whether the belief is appropriately responsive to the practical costs of being wrong.

Although these positions might seem mutually exclusive, the reality, I think, is more complicated. Imagine, e.g., a sprinter trying to psych himself up for a race against an overwhelming favorite. Despite the long odds, suppose he manages to convince himself that

⁷ From the well-known metaphor first employed (according to most sources) by John F. Kennedy, that “a rising tide lifts all boats.” Kennedy’s general idea was that when the economy was doing well, all people would benefit.

⁸ In Fantl and McGrath’s words: “According to received tradition in epistemology, whether a true belief amounts to knowledge depends only on purely epistemic factors—factors that are appropriately truth related” (2007, p. 558). In Stanley’s words: “I will call the thesis that knowledge does not depend upon practical facts *intellectualism*” (2005, p. 6); putting this thought together with Stanley’s claim that “Someone’s practical investment in the truth or falsity of her belief is completely irrelevant to truth-conduciveness in any sense” (2005, p. 2), seems to give us the same thesis.

⁹ I opt for “practicalism” both because it captures the basic idea—that knowledge can depend on practical interests—as well as because the alternatives have various problems. “Pragmatism,” for instance, seems to have too many historical associations built in, and “pragmatic encroachment” lacks the right sort of grammar.

(this time!) he is going to win. What intellectualism tells us is that even though having this belief might be necessary for the sprinter to give his best effort, and will therefore score well with respect to this practical goal, if it is not appropriately related to the truth then it is a poor candidate for knowledge. When it comes to whether a belief amounts to knowledge, it seems, the only thing that matters is how well the belief fares with respect to traditional truth-connected factors—factors such as evidence, or reliability, or what have you.

But while I take it this is both a standard and plausible way of looking at things, it still leaves a good deal unresolved. For notice, even if we agree that only truth-connected factors such as evidence or reliability matter for knowledge, intellectualism—at least as defined above—still leaves us with no guidance about how much of a particular truth-connected factor is required for knowledge. Put another way, what the above view is silent about is how much evidence or reliability or what-have-you you have you is “enough” evidence or reliability or what-have-you for a belief to amount to knowledge.¹⁰ In order to round out the intellectualist view, we therefore need to say more. And there would seem to be two main ways in which this “more” might be understood.

On the one hand, and in keeping with a broadly “intellectualist” (perhaps even Cartesian) approach to things, we might say that it is only our “purely epistemic goals and concerns”—rather than our practical goals and concerns—that determine the thresholds relevant to knowledge; in other words, that it is our “purely epistemic goals and concerns” that determine how much evidence or reliability or the like is needed, in order for a belief to amount to knowledge.¹¹ On the other hand, one might say that our “purely epistemic goals and concerns”

¹⁰ As Ernest Sosa writes: “Our concept of knowledge involves various dimensions each admitting a threshold: (a) “belief”: how sure must one be? (b) “justification”: how much rational support is required for one’s belief? (c) “reliability”: how reliable are one’s operative sources or faculties? (d) “safety”: how easily might one have been wrong; how remote is any belief/fact mismatch?” (Sosa 2000, p. 6). Owens (2000) and Pace (forthcoming) highlight these questions as well.

¹¹ This seems to be the sort of position Jonathan Kvanvig has in mind when he talks about the “purely theoretical point of view”—in his words, “a point of view abstracting away from other purposes such as purposes which are practical, moral, aesthetical, religious, or political in nature” (Kvanvig 2009, p. 346). Kvanvig also appeals to this idea in his (2008).

do not exclusively determine where the thresholds relevant to knowledge are set, but that practical interests and concerns can bear on these thresholds as well.¹² We might think of the first of these views as an “extended” intellectualist view about how the thresholds relevant to knowledge are set, and the second as the original practicalist thesis, but now understood more specifically as a view about the sorts of interests and concerns that help to determine the thresholds relevant to knowledge.¹³

Suppose that one accepts this second view, on which our practical interests and concerns help to determine how much of a given truth-connected factor is needed for knowledge. Is this position compatible with the original (non-extended) intellectualist thesis as described above? Properly understood, I think the answer is Yes. Again, all the original intellectualist thesis told us was that only truth-connected factors mattered for knowledge. The claim that only our “purely epistemic goals and concerns” determine how and where the thresholds relevant to knowledge are set (or “how much” of a given factor is needed) was no part of that thesis, but was instead a separate add-on. As I am understanding these terms, it is therefore no contradiction to think of oneself as an “intellectualist practicalist” (or the like). What such a person would believe is that only truth-connected factors matter for knowledge, and that (in addition) how much of a particular truth-connected factor is needed is sensitive to our practical interests and concerns.¹⁴

¹² According to some philosophers, such as Pace (forthcoming), our moral concerns can bear on these thresholds as well. I think there is some truth in this, as I will suggest at the end of Section 5 below.

¹³ On this way of looking at things, the original practicalist thesis can be seen as ambiguous: it can be taken either (a) as the view that how one fares with respect to non-truth connected (such as physical well-being, as in the mountaineer case), rather than truth-connected factors, matters for knowledge; or (b) as the view that practical concerns help to determine *how much* of a given truth-connected factor is needed for knowledge. In what follows, I am therefore taking the view in this second sense. For more on this distinction, see my (2011).

¹⁴ If one wanted a more exhaustive picture of the various options, it would look as follows (thanks to Allan Hazlett for this idea): ‘

	Only truth-connected factors matter for knowledge	Non-truth-connected factors matter for knowledge
Thresholds are set exclusively by purely epistemic goals and concerns	Extended intellectualism (Descartes?)	(not plausible)

Clarifications aside, and supposing that the original intellectualist thesis is compatible with both views about how the thresholds relevant to knowledge are set, which of these further views should we adopt?

2. Against Extended Intellectualism

A number of strong considerations seem both to tell against the extended intellectualist view about thresholds and in favor of the practicalist one.

To begin with, we can note that the notion of what matters from a “purely epistemic point of view” (or the like) is considerably less clear than the notion of what matters from a practical point of view.¹⁵ Thus while I take it that virtually everyone would agree that things like making our mortgage payments on time, catching the correct train for an important job interview, and so on are things that we care about from a practical point of view, what it is that we care about from a “purely epistemic point of view” is far from evident. Prima facie it is therefore unlikely that such an obscure notion is actually guiding our judgments about knowledge, especially given the high degree of consensus we usually find on those judgments.

What’s more, the most obvious ways of trying to clarify or refine this idea seem to be of no help. For instance, suppose we opt for the view that what we care about from a purely epistemic point of view is simply this: to achieve the truth, and to avoid error. As Wayne Riggs (2003) points out, even if this is an accurate picture of our purely epistemic concerns, we are still left wondering how these different epistemic goals should be “balanced”: that is, how to balance our desire to achieve the truth on a particular topic as against our desire not to fall into

Thresholds are set by a variety of our goals and concerns, including our practical goals and concerns	practicalist intellectualism	anti-intellectualism
---	------------------------------	----------------------

For the purposes of this paper, I am simply assuming that the alternatives in the right column are implausible. But it is worth noting that something like anti-intellectualism has had influential takers in figures such as Foucault and Rorty, and has been defended more recently by Hazlett (forthcoming).

¹⁵ See my (2009) for more on this point.

error. Put in terms of our earlier point, the question of “how much” error-avoidance is needed, or how reliably our beliefs need to cleave to the truth, still remains unanswered.

Alternatively, suppose we claim that what we care about from a purely epistemic point of view is something richer than the simple goals of achieving the truth and avoiding falsity: to adopt Riggs’s own suggestion, perhaps what we care about from a purely epistemic point of view is to acquire a deep understanding of how the world works, or perhaps even to achieve wisdom. The basic problem with richer views of this sort, however, is that it is implausible that such rarified concerns have anything to do with the vast majority of our judgments about knowledge. For example, my judgment about whether the person next to me knows that the departing train is an express (rather than a local) hardly seem guided by such concerns. Instead, it seems much more reasonable to suppose that my guiding concerns here are practical ones—in particular, the practical costs (harms, losses) that would stem from acting on a mistaken belief. Were I mistaken in this belief, what I would miss out on is not understanding or wisdom; rather, what I would miss out on are the various things I intended to do at home (watching my favorite reality TV show, say), and it is how much the potential loss of these things matters to me that guides my judgment.

An appeal to practical interests therefore often seems needed to answer the “How much?” question that the plain (unextended) intellectualism thesis left unanswered.¹⁶ In addition to these concerns, however, a further, and more common, reason for thinking that our practical interests can bear on our judgments about knowledge comes from looking at cases. Although these cases will likely be familiar, it is worth having a concrete example in front of us,

¹⁶ To be clear, in offering these points against extended intellectualism I do not mean to deny that on some occasions our purely epistemic goals and concerns (assuming we have any) might be guiding our judgments about knowledge; perhaps this happens in our judgments about fundamental scientific principles, for example. I only mean to deny that the thresholds relevant to knowledge are determined *exclusively*—and, as it were, across the board—by our purely epistemic goals and concerns. Given that in many cases our beliefs concern topics that do not seem to advance any “purely epistemic goal or concern” of ours, it is this view that should be given up.

if for no other reason than that the details of these cases often change, and the details often matter.¹⁷

Suppose then that Hannah believes that her local bank will be open on Saturday, and this on the basis of her knowledge that the bank was open when she deposited her check on a Saturday two months prior.¹⁸ Suppose as well that it is true that the bank will be open on Saturday. Does her belief amount to knowledge? According to practicalism, in order to answer this question we need to bear in mind what is riding on the belief.

To take a Low Stakes version of the case, suppose that there are no urgent practical reasons for *us*, as evaluators, to make it to the bank on Saturday, and that Hannah reports her belief to us, along with her evidence for the belief. Under these circumstances, it seems that most of us would judge that Hannah does (in fact) know; or, at least, we would not be quick to judge that she fails to know. Intuitively, her evidence seems quite good, good enough to know.

Compare this with a High Stakes version of the case, where whether or not the bank is open matters a great deal to us. (Perhaps, if the bank is closed we will not be able to deposit a vitally important check.) Hannah again reports her belief, along with her evidence. Does she know? I take it that the most common reaction here, and the one that I share, is that in this case she does not know; or, at least, our inclination to say that she does not know seems considerably more powerful than in the previous case. Given how much hangs on the question, even though we might still regard her evidence as good, it no longer seems good enough to know.

I noted above that these Low Stakes/High Stakes cases are presented in various ways in the literature. Sometimes, for example, the subject being evaluated actually uses the word “know” and sometimes she does not; sometimes the practical stakes are the subject’s while sometimes (as in this case) they are the evaluator’s; and so on. I will come back to these differences later, but for now the only point I want to emphasize is that as I have described

¹⁷ In my (2011) I likewise describe the details of the High Stakes/Low Stakes case in this way.

¹⁸ “Bank cases” of this sort were first introduced by DeRose (1992). Other examples of High Stakes/Low Stakes cases of this sort include Stewart Cohen’s (1999) airport case and Fantl and McGrath’s (2002) train case. Stanley (2005) follows DeRose in focusing on bank cases.

things our judgments in the two cases are examples of what we might call *first-order epistemological judgments*. In other words, as third-party evaluators, what we are doing is judging whether or not Hannah's belief amounts to knowledge, first in the one case and then in the other; we are thus not making judgments, for example, about the truth value of sentences that use epistemic predicates to attribute or deny epistemic properties.¹⁹ In our case above, Hannah never said anything about "knows"—she simply reported her belief, a belief which in one instance seemed to amount to knowledge and in another instance did not.

While we therefore seem to have several good reasons to accept practicalism (based both on theoretical considerations about thresholds, as well as particular judgments about cases), as we noted earlier there also seem to be powerful reasons to reject it. It is therefore time to address the *whose stakes? problem* and the *stability problem* that we mentioned at the outset.

3. Whose Stakes?

Again, according to the *whose stakes? problem* there is no entirely satisfying way of determining whose stakes, exactly, matter when it comes to determining the thresholds relevant to knowledge.²⁰

According to Hawthorne and Stanley, for example, it is the interests and concerns of the subject that matter.²¹ But they readily grant that their view gets things wrong in certain key cases: namely, in cases where there is very little at stake for the subject but a great deal at stake for us as evaluators, their view predicts—incorrectly—that we would judge the subject to be a knower. They are therefore left with the task of trying to explain why our judgments that the

¹⁹ For example, it is not a claim that someone's utterance of the sentence "S knows that p" is false. These are more naturally thought of as meta-linguistic judgments than first-order epistemological judgments. Since it seems possible for these two judgments to come apart (see, for example, Sosa 2000), for simplicity I will focus on the first-order judgments.

²⁰ Or, according to some of the philosophers to be discussed in a moment, of the semantic value of "knows." For simplicity, I will continue to present the issue as a first-order one however.

²¹ It is worth adding that on Fantl and McGrath's view, it is not entirely clear whose stakes, exactly, are calling the shots. For example, although they explicitly accept the "subject sensitive" label for their view (2007, fn. 5), this doesn't seem quite right. Thus in their original (2002) Foxboro train case they were willing to deny knowledge to someone (the fellow on the train platform) with very good evidence, even though his stakes were quite low (in their words: "what's it to him?"). Their own flagship case therefore suggests a view where it is the attributor's stakes that matter more than (or at least as much as) the subject's.

subject fails to know in the low stakes cases are mistaken, even though our judgments are accurate in the other cases.²²

Similarly, attributor contextualists such as DeRose and Cohen are accused of getting things wrong when the stakes are low for the attributor but high for the subject. Thus Hawthorne claims that because on attributor contextualism “It is always the ascriber’s standards that call the epistemic shots, as it were” (2004, p. 59), it should be the case that when the stakes are low for the attributor, the attributor would judge that the subject knows. But this is not what we find. Even if getting the bank’s hours right is not important to us, as attributors, when there is a great deal at stake for the subject our tendency is likewise to judge that her belief does not amount to knowledge.²³

All told, then, there seems to be no single account about whose interests or concerns matter that gets all the cases right. And again, since all of these views are forced to try to explain why our judgments go astray on certain key cases, it might be thought that we would be better off remaining loyal to a more “orthodox” view from the outset—a view on which a mere difference in practical stakes cannot make a difference to whether someone knows, and on which our temptation to judge otherwise should be explained in some other way.²⁴

²² For discussion, see Hawthorne (2004, pp. 160–64) and Stanley (2005, pp. 77–79). For further criticism, see Schaffer (2006) and DeRose (2004; 2005). Although I will take for granted in one follows that a simple, well-motivated theory that gets all the cases right is preferable to one that has to turn error-theoretic at some points, I believe that some of the error-theoretic strategies that Hawthorne and Stanley adopt are questionable in their own right. For instance, the idea that we “project” our own stakes onto the subject, when the stakes are high for us but low for the subject, fails to account for the fact (to be discussed below) that thresholds rise when third party stakes are at issue as well. Perhaps Hawthorne and Stanley could reply that we project third party stakes onto the subject as well, but this would seem to add an additional complication into the account, making the rising tides view seem even simpler (in a positive way) by comparison.

²³ DeRose (2005, pp. 189–90) notes that there is nothing about attributor contextualism that requires us to deny that the attributor’s ascriptions might be sensitive to the subject’s stakes. While this seems right, DeRose does not seem to provide an explanation for why this should be the case; he also does not seem to see that our standards can rise when it is the stakes of third parties at issue. Henderson (2009; this volume) offers a very nice account—grounded in similar sorts of Craigian considerations to the ones I will consider in Section 5—of why third party interests should matter to our knowledge evaluations, and he describes his view as a kind of “motivated contextualism.” I would describe the view on offer here, however, not as a version of contextualism but rather as a version of stakes-sensitive invariantism—one where, again, the thresholds are sensitive to the concerns not just of the evaluator and the subject, but of certain third parties as well.

²⁴ For example, Duncan Pritchard (2006), in his review of Stanley’s book, suggests that these cases should be explained in terms of pragmatic rather than (broadly speaking) semantic influences. Other appeals to

So much for the problem—how should we try to solve it? Perhaps a bit surprisingly, I think that a way forward can be found by adding another case into the mix, one where the interests are attached to a different source than we usually find discussed in the literature.²⁵

Thus suppose that it matters little to Hannah whether the bank will be open on Saturday, and that it matters little to us (as evaluators) as well. But it matters a great deal to one of our colleagues, Allan, whether the bank will be open. Indeed, we can imagine that Allan has just told us that he will be evicted from his apartment and thrown out on the street if he doesn't manage to deposit his paycheck on time. Suppose moreover that after describing his situation Allan has just asked us whether we knew if the bank would be open on Saturday, and we honestly replied that we did not know (suppose we had no evidence either way; we were just ignorant). But now suppose again that we hear Hannah report her belief, along with her evidence (that she had been two months prior, etc.). Would we take her to be a knower?

Here again I take it that our natural inclination would be to say No. Given that a great deal is at stake for one of our colleagues, it again seems natural to judge that even if Hannah's belief turns out to be true, her evidence, while good, is not good enough for knowledge. At the very least, it seems implausible to accept the earlier judgments about the bank cases but to claim that Hannah knows in this case, where it is Allan's interests at stake, rather than our own.

If this thought is correct, then it seems to show that our judgments about knowledge are sensitive not just to the practical concerns of the subject and not just to our own concerns as evaluators; in addition, it seems, they are sensitive to the concerns of certain third parties as well. It therefore seems that the best way—perhaps the only way—to take all of these judgments at face value is to claim that our judgments about knowledge are sensitive to the highest of these stakes, whether the stakes are those of the subject, or of the evaluator, or of some third party. In other words, to accept something along the lines of the “rising tides” account sketched earlier—

conversational propriety in order to explain these cases include Rysiew (2001), Hazlett (2009), and Turri (forthcoming).

²⁵ Greco (2008; 2010) and Henderson (2009; this volume) are exceptions to this claim. Both explicitly suggest that third party stakes can raise the thresholds relevant to knowledge.

again, an account on which rising costs *either* for the subject *or* for the evaluator *or* for certain third parties can raise the thresholds relevant to knowledge.

Although this view apparently has the singular virtue of getting all the cases right, it needs to be refined in order to meet a variety of natural objections. I will attempt to do this in Section 6. The view would also be bolstered if we could offer a principled, independent reason for thinking that the thresholds relevant to knowledge can rise in this way—a reason that goes beyond our judgments about cases. This is something that I will attempt in Section 5. First, however, I will try to show how the beginnings of a solution to our second problem—the stability problem—emerges from the rising tides account.

4. The Stability Problem

Recall that, according to the stability problem, the thresholds relevant to knowledge cannot be sensitive to practical stakes because this would imply, absurdly, not just that knowledge can come and go with ease, but in particular that one can know more by caring less. We can see variations on this objection in the following passages from Herman Cappelen and Ernest Lepore, from Gillian Russell and John Doris, and from Linda Zagzebski:

[Interest-dependent views have] the following peculiar implications.... [If] you couldn't care less about what penguins eat, then epistemic standards are low, and it is easier to fall in the extension of 'x knows that penguins eat fish'. This aspect of Sensitive Moderate Invariantism opens up a strategy for increasing knowledge, a strategy not really available to humans. If you don't know whether penguins eat fish, but want to know, you might think that the only way to become more informed is through study; you have to gather evidence, try to learn more about penguins. If Sensitive Moderate Invariantism were correct, though, you have another option: You could take a drink or shoot heroin. If as a result you care less about penguins and their eating habits... you would know more (of course, on the assumption that p is true). But this is not how to improve your epistemic standing! (Cappelen and Lepore 2006, pp. 1044-1045)

Not giving a damn, however enviable in other respects, should not be knowledge-making.... As the various stakes cases seem to show, interest destroys knowledge and indifference creates it [on the practicalist account]. (Russell and Doris 2008, p. 433)

[Interest-dependent views] make the things the subject cares about relevant to what she knows. Just as Jim does not know the bank closes at 5:00 based on ordinary memory if it is highly important to him to get there before closing, so too, Jim does not know most of the ordinary things he thinks he knows if it is highly important to him that there is no Evil

Genius [deceiving him]. Standards are lower the less you care, higher as your caring increases, so according to this version of contextualism, the way to avoid Evil Genius skepticism is not to care. (Zagzebski 2009, p. 53)

The thresholds relevant to knowledge do not simply bounce around, then, as one might expect on a view that ties knowledge to practical interests. Instead, what it takes to know seems to involve much the same thing across contexts: namely, a high level of reliability, or good evidence, and so on.

How does the rising tides account help to deal with this problem? In short, because if our judgments about knowledge are sensitive to the stakes of third parties, then the thresholds relevant to knowledge will plausibly gravitate towards a level high enough to respect these stakes; that is, to respect the stakes of those third parties who might have an interest in the topic. Knowledge will therefore not come and go on this view, as our fickle interests as subjects or as evaluators change; if third party interests matter, this will instead help to stabilize the thresholds relevant to knowledge.

Of course this sketch invites a variety of questions. For one thing, why think that there might be third parties with an interest in whatever random topic is at issue? For another, why think that third parties might have access to the beliefs in question, so that their stakes might influence our evaluation of those beliefs?

To the first question, I think the right answer is that, given how plastic and unpredictable our practical concerns can be, there is always a story one might tell about why a topic might be of interest to someone, no matter how trivial or insignificant it might seem on the whole. Thus even if getting to the truth about the 323rd number of the Wichita, Kansas phone directory²⁶ might be (and presumably is) something that you and I could not care less about, it is easy enough to imagine someone who might care about this, if only because he or she wants to phone the person up. More generally, the very idea that there are certain topics that are necessarily trivial or unimportant, from a practical point of view, as opposed to just contingently trivial or

²⁶ One of Alvin Goldman's examples of a "trivial" belief; for this and others see Goldman (1999, p. 88; 2002, p. 61).

unimportant, seems like a mistake. Our practical needs are quite various, and our practical needs often change; if our judgments about knowledge are sensitive to the practical concerns of others, then this would plausibly push the thresholds relevant to knowledge quite high, high enough so that these concerns might be properly respected.²⁷

That said (and to return to our second question), why think that the relevant third parties might appeal to the beliefs in question, so that their interests might reasonably bear on them? Here I think there are two things to say.

The first is that, as information-sharing and information-dependent creatures, we are accustomed to the fact that others often turn to us for our beliefs about various topics, sometimes out of the blue.²⁸ And second, even though it might be unlikely that anyone with a real interest in the topic will approach us directly, we are familiar with the fact that as information-sharing creatures our beliefs, once expressed, can spread in unpredictable ways.

Clifford once again seems relevant here:

Our words, our phrases, our forms and processes and modes of thought, are common property... an heirloom which every succeeding generation inherits as a precious deposit and a sacred trust to be handled on to the next one, not unchanged but enlarged and purified, with some clear marks of its proper handiwork. It is not only the leader of men, statesmen, philosopher, or poet, that owes this bounden duty to mankind. Every rustic who delivers in the village alehouse his slow, infrequent sentences, may help to kill or keep alive the fatal superstitions which clog his race. Every hard-worked wife of an artisan may transmit to her children beliefs which shall knit society together, or rend it in pieces. ([1877] 1999, pp.)

Clifford's dramatic prose aside, the basic point to note is that since it is hard to say in advance who will come to rely on our judgments "downstream," it is not surprising that the thresholds

²⁷ Indeed, once one takes on board the plasticity and unpredictability of our practical concerns, another interesting possibility is that it is the plasticity and unpredictability of our own first-person practical concerns that keeps the thresholds high. Thus even though we might not, at the moment, care at all about a certain topic, we could usually imagine circumstances where we might come to have such an interest—and a prudent person would plausibly be sensitive to these possible concerns in his evaluations of belief as well. Although I will not pursue this point further here, I think there is a lot to be said for this line of thought: sensitivity to our possible first-person concerns would likewise seem to raise thresholds even with respect to those topics that we are currently indifferent about. My basic reason for emphasizing third-person interests and concerns instead is that I think it is only by appealing to these concerns that we capture the deontological character of many of our epistemic evaluations. I touch on this idea in later sections, but for more on this point see my (2009).

²⁸ Of course, in some unusual circumstances this is not true (say, with monks sworn to silence, or with social outcasts), but these exceptions basically prove the rule.

relevant to knowledge would gravitate towards a level high enough to respect the “typical” or “normal” stakes of others who might appeal to those judgments at some point later on.²⁹

The main idea, then, is that if third-party interests matter, then the fact that we as evaluators might not have a stake in a given topic, or that the subject of the belief might not have a stake, will not push the thresholds relevant to knowledge into the cellar, so long as we have reason to think that others might have a stake in the topic. But we do seem to have such a reason, given how plastic and unpredictable the practical concerns of others can be. By hypothesis the thresholds relevant to knowledge will therefore naturally gravitate, at least as a default, to a level high enough to respect the typical or normal stakes of others.

Of course in claiming that the typical or normal stakes of others on a given topic set the “default” threshold required for knowledge, I am supposing that these default thresholds can be overridden when we think that someone with elevated stakes might appeal to the belief. This is the dynamic that plays out in DeRose’s bank case, Fantl and McGrath’s train case, and so on. But why should we think that there actually *are* anything like “typical” or “normal” stakes, relative to a given topic? One reason is that if we can recognize when the stakes are “high” or “elevated,” relative to a given topic (and the cases suggest that we can do this quite easily), then this presupposes that we can recognize when the stakes are “normal.” It is worth noting, moreover, that what counts as a “normal” cost of being wrong can apparently differ from topic to topic, and perhaps for the same topic over time. Thus the cost of being wrong about whether a given beach has strong riptides, for instance, will “normally” or “typically” be quite high (injury, death, etc.), while the cost of being mistaken about my neighbor’s phone number will “normally” or “typically” be something less significant. On our theory, when the normal costs of being wrong are elevated in this way, what it takes to know on these topics will be elevated or higher

²⁹ I am supposing, moreover, that when one *tells* another that p, one represents oneself as a knower that p. (For more on telling, and how it differs from the broader speech act of assertion, see Moran 2006 and Millar forthcoming). I am not denying that we sometimes tell others less than we know, or that we sometimes represent ourselves as knowing less than we do. I do think, however, that conversational rules such as “assert the stronger” put pressure on us to tell others outright that p when we take ourselves to know that p, rather than to qualify our claims in some way (by prefacing our claim with “I think that...” or the like).

than usual as well—which is, I think, just what we find. Thus I take it that it characteristically takes more to know something like whether the beach has (or lacks) strong riptides, or whether the safety features of the space shuttle are in good working order, than it does to know topics where the normal costs of being wrong are less significant.

All that said, one issue that remains to be clarified is how we should think about the notion of “might” in the claim above that the “default thresholds can be overridden when we think that someone with elevated stakes might appeal to the belief.” For if we take this notion too broadly, then a new problem emerges, viz., not that knowledge will be too easy to come by (if one fails to care about the topic), but rather that it will be too hard. For instance, if we took the “might” to indicate something like logical possibility, then because it is fairly easy to imagine third parties for whom getting to the truth about a topic would not just be a concern, but a grave concern, then the thresholds relevant to knowledge would be stratospherically high. Thus it seems easy to imagine a third party who might be interested the 323rd entry in the Wichita phone directory because lives are at stake, and so on. Scepticism looms.

This is a problem I will return to in Section 6 of the paper. I think the problem is best addressed, however, if we first turn to what I earlier referred to as a “principled” account of why third-party interests should matter when it comes to our judgments about knowledge. This will help not only to provide independent support for the rising tides account, but also help to show how the account might be refined and clarified to address the concern just mentioned.

5. Craig’s Account

The principled account that I have in mind is Edward Craig’s “state of nature” account of our concept of knowledge.³⁰ On Craig’s view, a helpful way to “explicate” our concept of knowledge is by trying to identify the role that the concept plays in our lives. And the best way to identify that role, he claims, is by pointing to certain basic human needs and interests and

³⁰ See Craig (1990). Other philosophers who likewise appeal to Craig’s account include Fricker (1998), Weinberg (2006), Neta (2008), Greco (2008), Henderson (2009), and Kusch (2009). For critical discussion of Craig’s approach, see Kappel (2011) and Kelp (2011).

then asking how our concept of knowledge might have grown out of, or been a natural response to, those needs and interests.

Two basic human interests play a particularly important role in Craig's account. The first is our interest in appealing to others to supplement our information about the world: to gain more information about which mushrooms are fit to eat, for example, or which paths are tiger-free. The second is our interest in distinguishing the good sources of information from the bad (or the less than good)—that is to say, our interest in distinguishing those sources that are reliable on the question at issue from those that are less than reliable.

Craig then plausibly suggests that for creatures with this pair of interests, some method, some concept, would have been identified to help mark or “flag” the good sources of information. We might think of this as the knowledge* concept, where the star at the end is meant to signify that this concept is not yet our concept of knowledge, but rather (by hypothesis) an ancestor of our concept (cf. Pritchard 2009, ch. 4). Knowers*, on this view, would thus have been reliable sources of information about a given topic, and hence good sources about the topic.

But just how reliable would a source need to be in order to qualify as a good source? Here is where Craig's proposal is especially relevant to the rising tides account introduced earlier.

According to Craig, when the notion of a good source of information was first taking shape, the “how good? or “how reliable?” question would have had an essentially subjective answer; that is, the source would need to be as good or reliable as my first-person needs as an evaluator required (p. 85). Thus on this view,

if being wrong won't matter much, but being right will be very advantageous, I may be satisfied with an informant of lower reliability, one whose views have a lower probability of being true, than I would be if the situation were reversed, so that being wrong would turn out to be very damaging. (Craig 1990, p. 86)

Craig goes on to claim, however, that while the notion of a good source of information might have begun life with the subjective concerns of the evaluator in mind, it would naturally have evolved to take into account the concerns of others as well. Why? Because, as information-sharing and information-dependent creatures, just as I would like others to take my interests and concerns into account when evaluating sources of information, so too others would presumably hope that I would take their interests and concerns into account as well (pp. 87-88).

According to Craig's hypothesis, then, over time the concept of knowledge* would have become "objectified" so that our judgments about knowledge (or who counts as a knower) would be responsive not just to our own interests and concerns—that is, our own interests and concerns as evaluators—but to the interests and concerns of others as well. Or, perhaps more accurately, on this view the concept of knowledge* would gradually have given way to another concept—our contemporary concept of knowledge—in which sensitivity to the interests and concerns of others would play a central role. If Craig is right, then as modern-day knowledge evaluators when we judge that a belief amounts to knowledge we are judging that the belief is reliable enough not only given our first-person interests and concerns, but given the interests and concerns of third parties who might come to depend upon the belief as well.

The power of Craig's view therefore derives not just from the simplicity of its starting points—who could deny, after all, that we have the basic concerns Craig uses to develop his account?—but also from its ability to provide an independently plausible answer to the "How much?" question considered earlier. Thus to the question of "How much reliability or evidence (or what-have-you) is needed for a belief to amount to knowledge?" the Craigian answer is: as much as our interests and concerns require—where the "our" here would over time come to include not just the interests and concerns of the person doing the evaluating but also the interests and concerns of the third parties who might come rely on our evaluation.

One final point is particularly worth noting: if Craig's account of the gradual objectivization of the concept of knowledge is on the right track, then our whole practice of

knowledge evaluation would seem to be morally freighted. Why? Because if, while expecting others to bear in mind my interests while evaluating whether a belief amounts to knowledge, I all the while fail to respect theirs, then there would seem to be a failure of fairness here, or perhaps even a failure of justice. This point—about the sorts of quasi-moral expectations that inform our role as knowledge evaluators—will be significant as we look to refine the rising tides account.

6. Refinements

Recall that the main remaining issue for the rising tides account was how we should think about the “might” in the claim that our knowledge thresholds are sensitive to the elevated stakes of others who *might* appeal to the belief. In this section I want to propose that in our judgments about knowledge the “might” here is tracking the notion of what we may call a “real possibility”—that is, a possibility that we think “might actually” happen, as opposed to a possibility that we think “merely might” happen. After trying to clarify this idea, I will then try to show that appealing to the distinction to solve our problem is not ad hoc, because a distinction along these lines seems to play an important role in our wider, “non-epistemic” evaluations of others.

A few examples should illustrate the distinction I have in mind. Thus even though I believe it is broadly possible that someone with life-or-death stakes might appeal to my judgment about the 323rd number in the Wichita directory, I do not believe that this is a real possibility. I do not believe, that is to say, that there is a nearby world where this will happen, or that it “might actually” happen. Or again, even though I believe it is broadly possible that someone who bet his house on whether Derek Jeter got a hit in last night’s game might approach my table at Starbucks and appeal to my judgment about this topic (“I’ve gotta know! Did Jeter get a hit last night or not?!”), I do not believe that this “might actually” happen. By my lights,

although a possibility like this “might” obtain in the broad sense that it is conceivable, it is not a real possibility, a possibility that I think will occur in a nearby world.

Our knowledge evaluations, moreover, seem to track this distinction accordingly. Thus if you tell me that someone with life-or-death circumstances “might” appeal to my belief about my neighbor’s phone number—in the broad sense that such a possibility is not ruled out by the evidence—but that you do not think that this is a real possibility, or that it “might actually” happen, then I take it that I will not feel the same concern about the status of my belief as when I think that someone with these stakes “might actually” come to rely on my belief. More to the point (and here we can see some of the moral overtones that emerged from our discussion of Craig), not only will I not, as a matter of fact, feel the same sort of concern, but I will not think that I should be expected to show this concern, or that such an expectation is a legitimate one. And I think that similar things could be said about the bank cases, the train cases, and so on: the prospect that someone with elevated stakes “merely might” come to rely on the belief does not seem to trigger a rise in thresholds in the way that the prospect that someone with elevated stakes “might actually” appeal to the belief triggers the rise. It does not seem, for instance, that I could be legitimately criticized or blamed for failing to take these stakes into account in making my evaluations, in the same way that I could plausibly be criticized or blamed for failing to take into account the stakes of someone whom I acknowledge “might actually” rely on the belief.

As I suggested a moment ago, one further reason to think that our knowledge evaluations are sensitive to these different senses of “might” is that the distinction seems to play a role in many of our “non-epistemic” evaluations of others. Thus suppose you are visiting my house one day, and that we are eating peanuts from a bowl on my coffee table. You then tell me that I really should not leave the bowl on the table because someone with peanut allergies might be stopping by my house sometime in the next few days. In thinking about whether you are right about this—that is, in thinking about whether I really should be expected to remove the bowl—one basic factor that I will obviously want to consider is just how “real” the possibility is

that someone with a severe peanut allergy might stop by. If it turns out, for example, that you think that I should remove the bowl not because you think that someone with these allergies “might actually” stop by, but rather because you think that someone with these allergies “might” stop by in the sense that this possibility is conceivable or not ruled out by the evidence (after all, crazy things sometimes happen), then I will think your judgment is clearly wrong. Indeed, the idea that you might think that I should take into account these sorts of “merely possible” elevated costs seems so peculiar that it is hard to make the case sound realistic.

Suppose on the other hand that you say that I should remove the bowl because you have reason to think that one of your friends with severe peanut allergies *might actually* be stopping by sometime tomorrow. Perhaps, for instance, your friend has told you he will be in the area, and will stop by if his schedule allows it. In this case, I take it that even if I do not agree with you that I should remove the bowl (more on this in a moment), I will at least begin to feel pressure in this direction. And the reason seems to be that when it looks like someone with elevated stakes “might actually” be affected by my behavior, then there is some sort of responsibility on my part to take the person’s stakes into account; alternatively, that there would be grounds for criticizing me if I failed to take his stakes into account.

Of course, it should be clear these sorts of considerations (about how “real” a possibility is, for example) are not the only ones that will factor into my evaluations. Thus even if I think that someone with an extreme peanut allergy “might actually” stop by—indeed, even if I think such a person undoubtedly will stop by—it does not follow that I should automatically remove the bowl. For instance, it might be the case that this person has no right to enter my house, or perhaps has been warned *not* to come by my house, precisely because I have the nuts. The point here is only that these considerations will be *among* the factors that will bear on our judgments, when we are making evaluations of this sort. And just as with the nuts, so too, I want to say, with believers: thus while we do not think that to count as a knower someone’s position needs to be strong enough to respect the “merely possible” elevated stakes of others, we do seem to think

that his position needs to be strong enough to respect the elevated costs of those who “might actually” appeal to the belief.

One further point is worth emphasizing before turning to sum up: namely, that our judgments about knowledge can be mistaken, and this in more than one way. For instance, we can be mistaken about whether someone with elevated stakes “might actually” appeal to the belief in question, or about whether something that we take to be a cost is in fact a “real” cost,³¹ and so on. You might therefore innocently take Hannah to know because you think no one with elevated stakes might actually appeal to her belief, but someone else who is actually shouldering such stakes might be listening in all the while and take her not to know. On the view here, you would be mistaken in your judgment, and the eavesdropper would be right. It is in this sense I take it that knowledge is sensitive to the actual/objective facts about the elevated costs which “might actually” be in play, and it is in this sense (and as I noted earlier in fn 23) that I take the view to be a version of invariantism, albeit of an interest-relative sort.

It should be clear that there is nothing unusual about this result, however, because I take it that on virtually any account of knowledge we can be mistaken about whether the conditions for knowing are met. If I find myself in the middle of fake barn country, for example, I might take myself to know that this thing in front of me is a barn, but you (who know about the proximity of the fake barns) will judge my belief to be mistaken. In the usual case, I will make this mistake non-culpably, in which case you will withhold your blame or criticism. But you will still of course take me to fall short in the way relevant to knowledge. The same holds true, *mutatis mutandis*, for our account.

7. Conclusion

Summing up, I began by suggesting not just that our judgments about knowledge are naturally guided by practical concerns, but that the alternative intellectualist picture—on which

³¹ Thus I take it that a paranoid person might take something to be a cost that we do not recognize as such, and similarly for the depraved person. I am thus supposing that whether something counts as a “real” cost will be a matter for substantive debate, and not just up to individual predilections.

the thresholds relevant to knowledge are determined exclusively by our “purely epistemic goals and concerns”—runs aground for a variety of reasons. Support for the practicalist view therefore comes not just from looking at cases, but also from the fact that alternative views about how the thresholds relevant to knowledge are set seem like non-starters.

Of course, practicalism too was not without its problems. For one thing, there seemed to be no account of who stakes mattered that got all the cases right (the whose stakes? problem). For another, it seemed to predict, absurdly, that the thresholds relevant to knowledge would vary wildly as the practical interests and concerns of the subject or the evaluator varied (the stability problem). The stability problem in particular seemed to be a major reason why many contemporary philosophers have taken a dim view of practicalism.

As I suggested, however, recognizing that our judgments about knowledge are sensitive to a variety of concerns—not just to the concerns of the subject or the evaluator but of relevant third parties as well—seemed not simply to get the cases right but moreover to explain why the thresholds relevant to knowledge did not crater in response to changing interests.

The problem was then one of explaining why the thresholds did not skyrocket in response to the elevated concerns of those who “might” come to rely on the belief in question, and I suggested that the way around this concern was to note that our knowledge evaluations recognize a distinction between the elevated concerns of those who “might actually” rely on the belief in question as opposed to those who “merely might” rely on the belief. Far from being isolated or ad hoc, moreover, something like this distinction seemed to be in play in our wider, “non-epistemic” evaluations of others.

The practicalist account defended here is therefore a deeply social one. While on this view our judgments about knowledge are informed by our own concerns as evaluators—including our own practical concerns—they are also, and in some ways more profoundly, informed by the concerns of others who might come to rely on the belief in question. A respect for these concerns therefore seems to dictate what, at a minimum or as a default, is required for

knowledge, and also seems to dictate when more than the minimum will be needed. Not only does this view seem to make the best sense of our judgments about cases, it also points to ways in which our knowledge evaluations are connected to our other evaluative practices.³²

WORKS CITED

- Cappelen, Herman and Ernie Lepore. 2006. "Shared Content." In *The Oxford Handbook of Philosophy of Language*. Eds. E. Lepore and B. Smith. New York: Oxford University Press.
- Clifford, W. C. K. [1877] 1999. "The Ethics of Belief." In his *The Ethics of Belief and Other Essays*. Amherst, NY: Prometheus Books.
- Cohen, Stewart. 1999. "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives* 13: 57-89.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford: Clarendon Press.
- DeRose, Keith. 2004. "The Problem with Subject-Sensitive Invariantism." *Philosophy and Phenomenological Research* 68: 346-50.
- , 2005. "The Ordinary Language Basis for Contextualism, and the New Invariantism." *Philosophical Quarterly* 55: 172-98.
- Descartes, Rene. [1637/1641] 1998. *Discourse on Method and Meditations on First Philosophy*. 4th edn. Trans. Donald Cress. Indianapolis: Hackett.
- Fantl, Jeremy, and McGrath, Matthew. 2002. "Evidence, Pragmatics, and Justification." *Philosophical Review* 111: 67-94.
- , 2007. "On Pragmatic Encroachment in Epistemology." *Philosophy and Phenomenological Research* 75: 558-89.
- , 2010. *Knowledge in an Uncertain World*. New York: Oxford University Press.
- Fricker, Miranda. 1998. "Rational Authority and Social Power: Towards a Truly Social Epistemology." *Proceedings of the Aristotelian Society* 98: 159-77.
- Goldman, Alvin. 1999. *Knowledge in a Social World*. New York: Oxford University Press.
- Goldman, Alvin. 2002. "The Unity of the Epistemic Virtues" In his *Pathways to Knowledge*. New York: Oxford University Press.

³² Thanks to Jason Baehr, Nathan Ballantyne, Anne Beril, Michael Brady, Juan Comesaña, Sandy Goldberg, John Greco, Allan Hazlett, David Henderson, Chris Kelp, Fiona MacPherson, Josh Orozco, Michael Pace, Duncan Pritchard, Blake Roeber, Declan Smithies, Ernest Sosa, John Turri, and Sarah Wright for helpful feedback on earlier versions of this paper.

- Greco, John. 2008. "What's Wrong with Contextualism?" *Philosophical Quarterly* 58: 416-36.
- , 2010. *Achieving Knowledge*. New York: Cambridge University Press.
- Grimm, Stephen. 2009. "Epistemic Normativity." In *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. New York: Oxford University Press.
- , 2011. "On Intellectualism in Epistemology." *Mind*.
- Hazlett, Allan. 2009. "Knowledge and Conversation." *Philosophy and Phenomenological Research* 78: 591-620.
- , Forthcoming. "The Myth of Factive Verbs." *Philosophy and Phenomenological Research*.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. New York: Oxford University Press.
- Hawthorne, John, and Jason Stanley. 2008. "Knowledge and Action." *Journal of Philosophy* 105: 571-90.
- Henderson, David. 2009. "Motivated Contextualism." *Philosophical Studies* 142: 119-131.
- Locke, John. [1689] 1975. *An Essay Concerning Human Understanding*. Ed. Paul Nidditch. New York: Oxford University Press.
- James, William. [1896] 1956. "Is Life Worth Living?." In his *The Will to Believe*. New York: Dover.
- Kappel, Klemens. 2011. "On Saying that Someone Knows: Themes from Craig." In *Social Epistemology*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. New York: Oxford University Press.
- Kelp, Christoph. 2011. "What's the Point of 'Knowledge' Anyway?" *Episteme* 8: 53-66.
- Kvanvig, Jonathan. 2008. "Pointless Truth." *Midwest Studies in Philosophy* 32: 199-212.
- , 2009. "Responses to Critics." In *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. New York: Oxford University Press.
- Kusch, Martin. 2009. "Testimony and the Value of Knowledge." In *Epistemic Value*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. New York: Oxford University Press.
- Millar, Alan. Forthcoming. "Knowing from Being Told." In *Social Epistemology*. Eds. Adrian Haddock, Alan Millar, and Duncan Pritchard. New York: Oxford University Press.
- Moran, Richard. 2006. "Getting Told and Being Believed." In *The Epistemology of Testimony*. Eds. Jennifer Lackey and Ernest Sosa. New York: Oxford University Press.

- Neta, Ram. 2008. "How to Naturalize Epistemology." In *New Waves in Epistemology*. Ed. Stephen Pritchard and Vincent Hendricks. New York: Palgrave.
- Owens, David. 2000. *Reason without Freedom: The Problem of Epistemic Normativity*. New York: Routledge.
- Pace, Michael. Forthcoming. "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James's 'Will to Believe.'" *Noûs*.
- Pritchard, Duncan. 2006. Review of Jason Stanley: *Knowledge and Practical Interests*. *Notre Dame Philosophical Reviews*.
- , 2009. *Knowledge*. New York: Palgrave MacMillan.
- Reed, Baron. Forthcoming. "A Defense of Stable Invariantism." *Noûs*.
- Russell, Gillian, and John Doris. 2008. "Knowledge by Indifference." *Australasian Journal of Philosophy* 86: 429-37.
- Rysiew, Patrick. 2001. "The Context-Sensitivity of Knowledge Attributions." *Noûs* 35: 477-514.
- Schaffer, Jonathan. 2006. "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism." *Philosophical Studies* 127: 87-107.
- Schiffer, Stephen. 2007. "Interest-Relative Invariantism." *Philosophy and Phenomenological Research* 75: 188-95.
- Sosa, Ernest. 2000. "Skepticism and Contextualism." *Philosophical Issues* 10: 1-18.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. New York: Oxford University Press.
- , 2007. "Replies to Gilbert Harman, Ram Neta, and Stephen Schiffer." *Philosophy and Phenomenological Research* 75: 196-210.
- Turri, John. Forthcoming. "Epistemic Invariantism and Speech Act Contextualism." *The Philosophical Review*.
- Weinberg, Jonathan. 2006. "What's Epistemology For? The Case for Neopragmatism in Normative Metaepistemology." In *Epistemology Futures*. Ed. Stephen Hetherington. New York: Oxford University Press.
- Williams, Bernard. 1978. *Descartes: The Project of Pure Enquiry*. New York: Viking Penguin.
- Zagzebski, Linda. 2009. *On Epistemology*. New York: Wadsworth.