Pragmatic Encroachment and Epistemically Responsible Action

Draft – Please Cite Final Version

1. Introduction

Pragmatic encroachment (PE)[1] is the view that whether one knows a proposition is, at least in part, a function of practical factors that surround the truth or falsity of that proposition (see Stanley, 2004; Hawthorne, 2005; Fantl and McGrath, 2009; Weatherson, 2011). "Practical factors" are typically spelled out in terms of the stakes given the truth or falsity of a proposition, and are contrasted with factors that epistemologists have traditionally considered important, like the amount of evidence one has for a proposition, whether one believes it, whether one believes it reliably, etc. Proponents of PE have typically provided two kinds of arguments for their view. One is based in PE's ability to accommodate people's propensity to deny knowledge to subjects when the stakes are high and ascribe knowledge to their low-stakes counterparts when all non-practical factors are kept constant. These propensities form the basis of what I will call the *direct argument* for PE:

> **Direct Argument for Pragmatic Encroachment**: People are more likely to deny
> knowledge to high-stakes subjects than they are to low-stakes subjects, all else

---

[1] PE goes by several different names in the literature, most notably "subject-sensitive invariantism" and "interest-relative invariantism". As the details of individual versions of the view differ, I will take "pragmatic encroachment" to be the general theory concerning the relevance of practical, "non-traditional" factors in determining whether one has knowledge.

being equal; the best explanation for this pattern of knowledge ascriptions is that

knowledge is, in fact, sensitive to the stakes involved.

The direct argument can be challenged on two fronts: the most common approach has been

to challenge whether the relevant pattern of knowledge ascriptions is *best* explained by the

view that whether one has knowledge is partly a function of the relevant stakes. For

instance, some have argued that the relevant knowledge ascriptions should be explained by

a *semantic* thesis about "knows"[2], whereas others have argued that the relevant ascriptions

can be explained by the effects of psychological factors that undermine belief or

confidence[3].

Another challenge comes from recent empirical work, which calls into question

whether people do, in fact, make knowledge ascriptions in the ways that the PE-theorist

thinks they should. Results of the empirical investigations are mixed. For instance, a recent

series of experiments from Feltz and Zalpentine (2010) and Buckwalter and Schaffer

(2013) claimed not to have found any evidence that the stakes for a vignette character had

an effect on how likely a participant was to ascribe or deny knowledge to that character,

while a different set of experiments from Sripada and Stanley (2012), Pinillos (2012),

Pinillos and Simpson (2014) and Shin (2014) claimed to have found the opposite. While

this is not a complete catalogue of all the studies that have been conducted, and while new

experiments continue to be conducted, it is safe to say that, at least at this point, the

experimental jury is out.

---

[2] This is the contextualist approach, which states that the semantic value of a knowledge-ascription sentence depends on the context in which it is uttered. See, for example, DeRose (2009).

[3] See, for example, Nagel (2008).

However, the PE-proponent has another line of argument, one that does not appeal *directly* to patterns of knowledge ascriptions in high- and low-stakes situations. According to what I will call the *indirect argument*, PE is entailed by a combination of a popular view concerning the relationship between knowledge and practical reasoning, and judgments that it is more difficult to reason well when the stakes go up. An example of the relationship between knowledge and practical reasoning I have in mind is the following:

> (RKP) Where one's choice is *p*-dependent, it is appropriate to treat the proposition that p as a reason for acting if and only if you know that *p*. (Stanley and Hawthorne, 2008: 578)[4]

RKP turns on a notion of treating a proposition as a reason for acting in an *appropriate* way, where this notion of propriety is epistemic in nature. I won't defend the details of RKP here; the underlying idea, however, is that our practical reasoning can be good or bad in a specifically epistemic way, and reasoning well on the basis of a proposition requires that we have a substantial epistemic relationship with that proposition. If we accept RKP, we can then formulate the *indirect argument* for PE:

Indirect Argument for Pragmatic Encroachment

1) (RKP) Where one's choice is *p*-dependent, it is epistemically appropriate to treat the proposition that p as a reason for acting if and only if you know that *p*[5].

---

[4] For my purposes here the notion of "p-dependency" is unimportant (see Stanley and Hawthorne (2008) for details if interested).

[5] While RKP states that knowledge is both necessary and sufficient for appropriately treating a proposition as a reason for acting (given the constraint of *p*-dependency), this principle is sometimes separated into discussions of the necessity and sufficiency relationships. Jessica Brown (2008), for instance, labels the view that knowledge is necessary for appropriate practical reasoning NEC, and the view that knowledge is sufficient for appropriate practical reasoning SUFF. For the indirect argument to go through, the PE-proponent needs to at least accept SUFF; since several experimental results discussed below refer to RKP

2) Whether it's epistemically appropriate to treat *p* as a reason for acting depends on the practical circumstances surrounding the truth of *p*.

C) Therefore, (PE): Whether one has knowledge that *p* depends on the practical circumstances surrounding the truth of *p*.

Versions of the indirect argument can be found, for example, in Hawthorne (2004), Stanley (2005), Fantl and McGrath (2009), and elsewhere. Indeed, in addition to the philosophical support, some recent empirical work has argued in support of the indirect argument (especially Pinillos (2012)); specifically, it is argued that there is empirical reason to accept premise (2), as experimental results indicate that as the stakes go up the epistemic demands for good reasoning go up, as well.

The indirect argument is valid; is it sound? I will argue that it is not. The first sign of trouble is that if people do, in fact, think that the epistemic demands for good reasoning go up when the stakes go up, and if people do, in fact, accept that there is a significant and important connection between what we know and the conditions under which we can reason well, then why do we not see much more robust results in empirical tests of the direct argument? One could, of course, posit an error theory – either at the level of folk judgment or at the level of experimental design – in those experiments that do not report the judgments anticipated by PE. One could also argue in the other direction, positing an error theory that says that those experiments that produced results in favor of PE were flawed, in which case one would have to also explain what is wrong with the indirect argument.

---

explicitly, however, I will continue to present the indirect argument as relying on it. Ultimately, however, my argument is neutral between RKP and SUFF.

I will argue here instead that judgments that appear to support the indirect argument for PE are better interpreted not as judgments that the epistemic demands for *knowing* increase as one's practical situation becomes more demanding, but instead as judgments reflecting a different kind of normative epistemic evaluation, namely whether one is acting in an *epistemically responsible* way. By "acting in an epistemically responsible way" I have in mind the following kind of evaluation:

Epistemically Responsible Action: One acts in an epistemically responsible way in a given situation just in case one is able to meet the demand for the provision of reasons to believe that one knows what one is doing, to a standard determined by that situation.

This is, then, a *kind* of error theory (although a different one from those mentioned above): when the stakes are high, people will sometimes report a judgement that a subject lacks knowledge, but what is *really* going on is that participants are judging that a subject is not acting in an epistemically responsible way. Importantly, on the view I am defending here, one can know that *p* but still act on *p* in an epistemically irresponsible way: the practical demands of one's circumstances might make it such that one ought to make sure that one knows what one is doing before one does it.

The general idea, then, is the following: when someone treats a proposition as a reason for acting we can evaluate them epistemically both in terms of whether they know that proposition, as well as in terms of whether they are acting on their knowledge in the right kind of way. In the context of the plausibility of the indirect argument, then, one can meet the conditions of the RKP by knowing the relevant proposition, but, because of the nature of one's situation, might nevertheless act on that proposition in an epistemically

irresponsible way[6]. My charge against the PE proponent, then, is that she is interpreting judgments that are indicative of whether we are adhering to certain normative epistemic requirements *generally* as being indicative of whether we have *knowledge* specifically. There are, however, normative epistemic requirements that make demands of us that are indicative of something other than our possession of knowledge.

Support for this view comes from what I argue is the most plausible way to interpret the experimental results both for and against PE, and is based on recent work in psychology and cognitive science that argues that people are *epistemically vigilant*, i.e. people look for more or fewer reasons to accept information depending on the practical circumstances under which that information is presented to them. A consequence of our epistemic vigilance is that practical circumstances can bring with them higher normative epistemic standards; these standards, however, need not be reflective of one's possession of knowledge. If this the case, though, we have significant reason to reconsider the plausibility of the indirect argument for PE[7].

My argument proceeds as follows. In section 2 I will survey some of the main evidence for PE, both in terms of the direct and indirect arguments. In section 3 I will show how work on epistemic vigilance supports the view that we hold subjects to a normative

---

[6] The claim that one can know that *p* but still fail to meet some epistemic standard when reasoning on the basis of *p* might strike some as prima facie counterintuitive. However, I think that once we have a notion of epistemically responsible action on board these intuitions largely go away. I address this and related objections in detail in section 5.

[7] A note about what I am *not* doing. While the PE proponent is interested in defending a kind of theory of knowledge, I am not: my interest is in showing that the evidence used to support PE via the indirect argument better supports the view that normative epistemic requirements as they pertain to practical reasoning are not exhausted by knowing. This is not an argument that some specific theory of knowledge is better than PE, but only that the support for PE is significantly undercut. I will not here propose a conception of knowledge to take its place. Nor will I argue for or against the RKP; I do, however, address the ramifications of my argument for the RKP in section 5.

epistemic standard in practical reasoning that does not pertain solely to whether one has knowledge. In section 4 I will show how this standard – what I call epistemically responsible action – can make sense of the ambivalent results from experimental philosophy. In section 5 I respond to various objections, specifically that the view I am arguing for here is inconsistent with the RKP, and in section 6 I very briefly consider the plausibility of PE in light of my arguments.

2. Pragmatic Encroachment In and Out of the Armchair

PE has been proposed as an explanation for why certain knowledge ascriptions seem to be affected by practical factors like stakes and interests. Consider, for example, the following pair of "bank cases":

Low Stakes: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.'

High Stakes: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was

7

at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah

points out, banks do change their hours. Hannah says, 'I guess you're right. I don't

know that the bank will be open tomorrow.' (Stanley, 2005: 3-4).

The story goes that, intuitively, Hannah's assertions that she does know in *Low Stakes* and

does not know in *High Stakes* that 'the bank will be open tomorrow' are both true. We are

also to assume the following: in both cases, Hannah has the same amount of evidence in

support of her belief about whether the bank is open, employs the same cognitive

mechanisms in reaching her beliefs, employs equally reliable belief-forming methods, and

so forth[8]. The traditional invariantist (who thinks that whether one has knowledge is not

directly a function of the practical consequences of being wrong) is, at first glance,

seemingly unable to accommodate both of these intuitions simultaneously; at least, she has

some explaining to do. The PE theorist, however, can accommodate the relevant judgments

with ease: the practical consequences surrounding the truth of the relevant proposition

makes it such that one can know it in *Low Stakes*, but not in *High Stakes*.

Armchair judgments, then, give us reason to take PE seriously. Out of the armchair, a

number of those working in experimental philosophy have conducted studies testing

whether the intuitions that Stanley and others think we should have in response to the

above cases and cases like them are actually widely held. As mentioned earlier, the results

vary: some studies claim to show that the knowledge judgments of lay-people are in line

with what PE would predict, whereas others do not. As we also saw above, the PE

proponent can defend her theory *indirectly* rather than by direct appeal to intuitive

---

[8] Whether all of these assumptions are true is a matter of debate: we might, for instance, think that Sarah's challenge in *High Stakes* changes some of these facts. I won't pursue this line of argument here.

knowledge ascriptions. The indirect argument relies on the RKP in combination with judgments that support a connection between practical demands and epistemic demands for good reasoning. (Examples of this kind of argument can be found throughout work that defends PE; see (Hawthorne 2004: 176), (Stanley, 2005: 10-11), (Fantl and McGrath 2009: 66), amongst others.)[9]

Why might we think that the epistemic demands for good reasoning go up as the stakes increase (in other words, why think that premise (2) in the indirect argument is true)? We again might make a general appeal to intuitions: we could appeal to cases in which it seems that when the stakes are high one needs to be in a better epistemic position in order to treat a proposition as a reason for acting. Fantl and McGrath (2009), for example, appeal to intuitions concerning a situation in which one is presented with an opportunity to cross a frozen lake. It would seem odd, they claim, to be willing to cross the ice but not to be willing to say that one knew that the ice was thick enough to support one's weight. If our judgments that it would be epistemically improper to act on a given proposition correlate with the practical consequences of that proposition turning out false, then this might be evidence for a principled connection between what one can know and what one can treat as a reason for acting.

Away from the armchair, Pinillos (2012) has argued that there is empirical evidence that supports premise (2) of the indirect argument. The motivation for his experiments is the following:

---

[9] Fantl and McGrath (2009), while they ultimately argue that RKP is false, appeal to a principle of a similar type, what they call the Knowledge-Justification Principle: (KJ) If you know that p, then p is warranted enough to justify you in φ-ing, for any φ (66). We could, then, derive PE if we accept KJ and that the demands for being warranted enough to justify one is φ-ing depends on the practical circumstances of the situation in which one φ's.

If it turns out that the folk implicitly accept RKP, then this will be further evidence that [PE][10] is true. Suppose we were to find that when ordinary people were asked if X knows P (in some carefully selected situation), then they would often answer a different question: whether it is appropriate for X to use P as a reason for acting. If we found that people did this, then this would be some evidence that they were behaving as if those questions were equivalent, and so this would be some evidence that they implicitly accept RKP. (28)

To test his view, Pinillos conducted an experiment in which participants read one of the following two scenarios:

Typo-Low: Peter, a good college student has just finished writing a two-page paper for an English class. The paper is due tomorrow. Even though Peter is a pretty good speller, he has a dictionary with him that he can use to check and make sure there are no typos. But very little is at stake. The teacher is just asking for a rough draft and it won't matter if there are a few typos. Nonetheless Peter would like to have no typos at all.

Typo-High: John, a good college student has just finished writing a two-page paper for an English class. The paper is due tomorrow. Even though John is a pretty good speller, he has a dictionary with him that he can use to check and make sure there are no typos. There is a lot at stake. The teacher is a stickler and guarantees that no one will get an A for the paper if it has a typo. He demands perfection. John,

---

[10] Pinillos does not use the acronym PE; instead, he uses "FS-IRI", or "folk-sensitive interest-relative invariantism." For my purposes I take the difference between these views to be solely terminological.

however, finds himself in an unusual circumstance. He needs an A for this paper to

get an A in the class. And he needs an A in the class to keep his scholarship. Without

the scholarship, he can't stay in school. Leaving college would be devastating for

John and his family who have sacrificed a lot to help John through school. So it turns

out that it is extremely important for John that there are no typos in this paper. And

he is well aware of this.

After reading one of the two above scenarios, participants were then given the following

prompt: "How many times do you think Peter (John) has to proofread his paper before he

knows that there are no typos? ___ times." Pinillos' results showed that, to a statistically

significant degree, the values reported for the number of times that John needed to

proofread were higher than they were for Peter.

The above type of experiment is what Pinillos calls an "evidence gathering

experiment": participants are asked to judge how much more evidence a character in a

vignette would need to gather before the participant judges that character as knowing or

acting appropriately. Many of the other experiments that survey intuitions about

knowledge that did not show a statistically significant effect of stakes on knowledge are not

of this type; rather, participants in such experiments are generally asked directly whether

they agree (or the degree to which they agree) with a statement that ascribes knowledge to

a given character in a given vignette[11]. The thought behind the evidence-gathering

---

[11] Others, like that from Sripada and Stanley (2012), tested judgments both about the amount of evidence that a vignette subject would need to gather in order to count as knowing, as well as judgments about whether a subject had knowledge directly. Interestingly, while all of the evidence-gathering judgments elicited by Sripada and Stanley were consistent with PE, not all of the knowledge judgments were. See their (2012) for details.

experiments, then, is that if participants judged that one needed to gather more evidence in order to qualify as knowing or acting appropriately when the stakes were high than when they were low, then participants thought of knowledge attributions as being sensitive to those practical concerns. If Pinillos is correct, we seem to have empirical support for the indirect argument for PE.

However, I will argue that the results from Pinillos' experiment and others like it can be explained in a better way; specifically, I will argue that they can be explained in a way that makes sense of why some experiments have shown results in favor of PE while others have failed to. The argument goes as follows: participants' judgments of experimental cases do not necessarily track whether the subject in the vignette has knowledge, but rather whether that subject is acting properly on their knowledge to an epistemic standard determined by their situation. This standard consists in an expectation that one is being *epistemically vigilant* to a degree that is dictated by one's situation. These kinds of judgments, therefore, track a different kind of normative epistemic status, one not encapsulated by the RKP, and that I have called *epistemically responsible action*. To support the premises of this argument we first need to look closer at the experimental data from Pinillos, as well as what it means to say that people are epistemically vigilant.

3. Epistemic Vigilance and Epistemic Responsibility

While Pinillos' evidence-gathering experiments showed that participants generally required that one gather more evidence regarding the truth of a proposition in order to have knowledge when the stakes were high than when they were low, when participants were asked to *explain why* they made the judgments they did, Pinillos found the following:

Out of the 66 subjects who gave asymmetric responses consistent with stakes

having the predicted effect [namely, that in the above vignettes, John needed to

check more times than Peter], 46 chose to explain their answers. What was found is

truly surprising. Although the prompts were clearly about knowledge, only one

person in this group actually used the word-form "know" in their answer…25 of

them explained their answers by claiming some variation on the theme that John is

required to/needs to/should check more times because the stakes are higher for

him…These answers reflect the simple idea that the normative requirements for

Peter and John differ because of the differences in stakes. (29)

Pinillos provides an explanation for why participants did not refer to Peter or John's

knowledge when they explained why they chose the number of times either Peter or John

needed to check their work: subjects "substituted" a complex question – the question of

whether the evidential requirements for knowledge increased in high-stakes situations –

with a simpler question – the question of what John should do given the circumstances[12].

That participants made this substitution, argues Pinillos, is evidence that they treat

questions about knowledge and appropriate action as equivalent, which, in turn, supports

the view that participants implicitly accept RKP. If participants do, in fact, accept RKP (or

---

[12] The notion of "substitution" that Pinillos appeals to here is from the work of Kahneman and Frederick (2002). Kahneman and Frederick argue that, as a result of employing representativeness and availability heuristics, people often answer a complex question as though it were a simpler question. The provide the following examples:

> A person who is asked "What proportion of long-distance relationships break up within a year?" may answer as if she had been asked "Do instances of swift breakups of long-distance relationships come readily to mind?" This would be an application of the availability heuristic. A professor who has heard a candidate's job talk and now considers the question "How likely is it that this candidate could be tenured in our department?" may answer the much easier question: "How impressive was the talk?". This would be an example of the representativeness heuristic. (53)

something like it), and their evaluations of the propriety of reasoning change depending on the stakes, then this is reason to think that the indirect argument is a good one.

It's not clear why, though, we should think that the participants in Pinillos' study subconsciously substituted a question about knowledge with a question about what John "needs to do", given that the first prompt was explicitly about what John *knows*, and that knowledge ascriptions do not appear to be complex in any particular way. There is, after all, good reason to think that we think of questions about knowledge as "simple" questions: knowledge ascriptions are something that we make frequently and with apparent expertise, "know" is an extremely common verb in the vocabulary of every major language, and we interpret the world in terms of knowledge from a very young age (see *****). The results of Pinillos' study indicate that participants hold vignette subjects to more stringent epistemic standards when the stakes go up, but we do not need to interpret these standards as ones that pertain directly to one's knowledge. Instead, we can interpret the results as follows: when we evaluate someone in a high-stakes situation as needing to acquire more evidence, this is not necessarily an evaluation that one needs to work harder to gain *knowledge*; indeed, if we take the results at face value, participants are *not* making this evaluation. Instead, I argue that these evaluations are indicative of whether someone is acting in an epistemically responsible way, where a failure to act in such a way is the result of someone failing to be epistemically vigilant enough, given their situation. To argue for this view, we need to look at what it means, exactly, to be epistemically vigilant.

Recent work in psychology and cognitive science on the functions of reasoning and communicating starts with a general observation: while human beings rely on communicating with one another, we want to minimize the chances of being misinformed.

As a result, we employ cognitive mechanisms that help reduce the risk of being, from our perspective, misled. Employing these mechanisms, however, is potentially cognitively costly: we don't want to have to actively monitor all of the information that we are presented with in detail to determine if we should accept it, so in order to minimize such costs we employ these mechanisms selectively, and more or less stringently, depending on the circumstances. Monitoring of information for trustworthiness has been called *epistemic vigilance* (Sperber and Mercier, 2011; Sperber et. al. 2010). Sperber and Mercier (2011) argue that "the task of epistemic vigilance is to evaluate communicators and the content of their message in order to filter communicated information" (60). Sperber et al. (2010) argue that epistemic vigilance is furthermore a process that occurs as information is being processed and comprehended, and that,

> Comprehension involves adopting a tentative and labile stance of trust; this will lead to acceptance only if epistemic vigilance, which is triggered by the same communicative acts that trigger comprehension, does not come up with reasons to doubt (368-9).

Work on epistemic vigilance suggests that, in general, when we are presented with information that is inconsistent with our expectations, our desires, and our beliefs; that comes from what we perceive to be an untrustworthy source of information; or that has significant relevance to us, we both take longer to process this information, and are more likely to generate reasons for thinking that this information is false, than when that information is consistent with our expectations, desires, or beliefs; or comes from what we perceive to be a trustworthy source; or has little significance to us. By generating reasons to think that information that is presented to us is false in the former situations we thus

apply more stringent epistemic standards towards the providers of said information: we need, in general, more or better reasons to accept that information.

Consider first instances in which we are presented with information that is inconsistent with our expectations, desires, or beliefs. In one set of experiments, Eagly et al. (2000) tested how easily one could remember information that was either proattitudinal (cohered with or supported the attitudes of the participants) or counterattitudinal (did not cohere with or challenged the attitudes of the participants). While both kinds of information were as easily recalled, they were recalled for different reasons: Eagly et al. argued that proattitudinal information was easily recalled because of its coherence with the attitudes of the subject, whereas counterattitudinal information was easily recalled because the subjects had thought about it more, either by trying to find flaws with the information they were presented with, or by generating reasons to disbelieve it.

Generating reasons to doubt some information increases the epistemic standards required for accepting that information: one will expect more or better reasons to accept information that is counterattitudinal than that which is proattitudinal, since one will have doubts that need to be satiated. For instance, in another study Ditto et al. (2003) presented participants with either favorable or unfavorable medical diagnoses. When diagnoses were favorable, participants were much more likely to simply accept the information that was given to them. When they were unfavorable, however, participants were much more likely to double-check the results they were given. That unfavorable and/or unexpected information elicits automatic double-checking suggests that people need more or better reasons to accept said information that in cases where information is favorable and/or expected.

The same results can be found in in cases in which people are presented with information that conflicts with one's beliefs. Richter et al. (2009), for example, argue that in cases in which information was presented to subjects that conflicted with information that was either "active in working memory or easily accessible in long-term memory", they "spontaneously activate knowledge that is incongruent with the content of a message" (540). Activating reasons to reject information again sets a higher epistemic standard: in order to accept a given bit of information that conflicts with one's beliefs, one's doubts created by the activation of incongruent knowledge has to be satiated.

Similar to instances in which one is distrustful of the content of a message, people will also take longer to process and will automatically bring to mind "alternative meanings" when information is received from what one takes to be an untrustworthy source. Schul et al. (2004), for example, argue that when one trusts the source of some information,

> receivers tend to concentrate on message-congruent associations. However if the
>
> source is suspected of being untrustworthy, receivers spontaneously activate
>
> message-incongruent associations, as they are considering what might happen if the
>
> message is invalid (677).[13]

---

[13] Our vigilance does not seem to be limited to any particular kind of informational content. Jonathan Haidt (2001), for example, argues that similar effects as those described above are observed when surveying people's *moral* judgments, in particular. According to Haidt's model (what he calls the "social intuitionist" model), moral judgments are made initially as quick, intuitive judgments, which are followed by slow moral reasoning when one is required to defend their judgments to others. We engage in moral reasoning in circumstances similar to those described in the epistemic vigilance literature: Haidt argues that we are more likely to scrutinize evidence that supports a view that conflicts with our moral judgments, that we are more likely to scrutinize moral judgments of those that we do not find trustworthy, and that we are likely to simply accept moral judgments if they are not relevant to us, or if it is made by someone we trust. Although Haidt's conclusions are drawn with an eye to supporting his social intuitionist model of moral judgment, we can also see how they are accounted for by the general mechanisms that underlie epistemic vigilance, with one simply exercising one's vigilance towards moral claims (for similar results concerning "justice inferences," see Ham and van den Bos (2007)).

The results from Schul et al., then, support the view that epistemic standards increase when one is presented with information from a source that is perceived as untrustworthy.

So far, then, we have seen a number of kinds of circumstances in which we might hold the provider of some information to higher epistemic standards as a product of factors that are independent of how likely that information is to be true, namely those cases in which information one is presented with is counterattitudinal, comes from a source one finds untrustworthy, or that one finds important. This last characteristic of epistemic vigilance is most important in terms of evaluating arguments for PE: we tend to be more epistemically vigilant when the perceived importance of a situation is high. When reasoning in high-stakes situations, then, we would expect to spontaneously generate reasons for doubt, and thus require more or better reasons to overcome that doubt and accept the relevant information.

Furthermore, the mechanisms that underlie our nature as epistemically vigilant are also those that create expectations in us that others be epistemically vigilant when the situation calls for it. For instance, Gloria Origgi (2012) discusses the notion of what has been called "epistemic trust." According to Origgi, epistemic trust comes in two varieties: *default trust* – "the minimal trust we need to allocate to our interlocutors in order for any act of communication to succeed" – and *vigilant trust* – "the complex of cognitive mechanisms, emotional dispositions, inherited norms, reputational cues we put at work while filtering the information we receive" (224). When there is nothing particularly at stake for us given the truth or falsity of some information we tend to take a stance of default trust towards the source of that information. When the truth of some matter *is* important to us, however, we tend to be epistemically vigilant:

We take the responsibility to check the reliability of our epistemic vigilance when the stakes are high. For example, in normal situations, the stakes related to the Moon landing for my life are quite limited. But if I was an American citizen and was called to vote on the possibility of re-opening the Apollo program, which would imply a huge public investment of money, then I would be more vigilant about the veracity of the facts and the feasibility of the program." (224)

The vigilance that we exhibit when accepting information for ourselves is furthermore that which we ought to exhibit not only when the truth of some matter is important for us, but when it is important for others when transmitting that information to them. As Orrigi notes, "if the stakes are higher in endorsing a belief, I should be vigilant about the mechanisms that make me endorse it and re-transmit it" (227).

There are, furthermore, many different kinds of mechanisms that can underlie our endorsing and retransmitting a belief. Origgi is primarily concerned with the mechanism of testimony: when we acquire a belief from others through testimony, we should be more epistemically vigilant with regards to trusting that person as a good source of evidence when the stakes go up. However, it seems that we are similarly vigilant concerning many other kinds of sources of evidence when the stakes go up, as well[14]. Boaz Miller (2014a), for instance, argues that the kinds of factors that impact the ways in which we trust others as sources of testimonial evidence are those that impact the ways in which we trust sources of evidence generally. For example, Miller argues that certain values influence "the evidential thresholds for justified epistemic judgments, such as belief formation or theory acceptance"

---

[14] Thanks to an anonymous referee for suggesting that I address this point.

(76), and that prejudice can affect how one interprets the plausibility of non-agential sources of evidence (Miller cites Fricker (2007) who provides the example of a journal referee who does not trust a newly implemented reviewing process because of bias). Philip Nickel (2013) also argues that it is appropriate to conceive not only other agents as trustworthy sources of evidence, but in addition cognitive processes like memory, faculties like perception, as well as pieces of technology. This is because, according to Nickel, we are entitled to our expectations that non-agents properly perform their function as a source of evidence, and that such entitlements are indicative of our trust in those non-agents (225).

If we trust not only other agents as sources of evidence but also processes and objects, it would stand to reason that when the stakes go up we exhibit vigilance with regards to the source of our evidence, regardless of what the source of that evidence should be. Common experience, I think, confirms this: in general, if it's very important that some information that we acquire as a result of one of our belief-forming mechanisms is true we are liable to make sure that that mechanism is calibrated in the right way. Furthermore, if you are to rely on me as a source of evidence, you would expect me to be vigilant towards the source of my evidence when the stakes are high. For instance, if it's very important that I be on time for a meeting, I might be expected to do more than, say, glance at the information on my calendar once: you might expect me to either supplement my memory with something more reliable (say, by writing it down), or by taking extra strides to make sure that I have, in fact, memorized it (say, by repeating the information over and over to myself, or by associating it with a mnemonic, etc.). Or, consider an instance in which I rely not on an individual agent as a source of information, but on the output of the efforts of a collection of agents: for example, I rely on a quick Wikipedia search when the truth of the

information I am looking up is not terribly important, but were I to rely on this information

when the truth of the information *was* important, say when you and I are working on a

project together, I might be expected to look at the sources the relevant article cites, or else

otherwise confirm the veracity of that information. The same seems to be the case when it

comes to relying on technology as a source of evidence: I rely on my watch to tell the time,

but when the stakes are high I might be expected to make sure that my watch is, in fact,

functioning correctly, again, especially when you are relying on me as a source of

evidence[15]. Finally, Origgi (2010) argues that we ought to be epistemically vigilant towards

our research and publication practices as processes that can be relied upon to produce

good evidence, and that we are right to expect others to exhibit this vigilance.

When we rely on others as sources of evidence, then, we expect them to have been

vigilant with regards to the source of their evidence, whatever that source should be. In

general, then, we expect others to be trustworthy sources of information, and one of the

reasons why we generally try to be trustworthy towards others is that we expect others to

be epistemically vigilant towards us. Sperber (2013), for instance, argues that:

> From an evolutionary point of view, the fact that humans can engage in
>
> communication the way they do, expect others to be rather honest and indeed are
>
> rather honest themselves begs for an explanation. Our explanation is that
>
> communicators are generally honest *because* they are addressing a vigilant
>
> audience. Quasi-universal vigilance makes dishonesty less likely to be beneficial in

---

[15] Goldberg (2012) argues that there are important differences between the ways in which we rely on other people and the ways in which we rely on instruments as sources of evidence, namely that the former but not the latter can be conceived of as an extension of one's belief-forming processes. This is not to say, however, that we cannot trust instruments as sources of evidence at all, just that the ways in which we rely on people and instruments are different.

the short run and more likely to be costly in the long run: falsehoods may be

disbelieved, and dishonesty may have reputational costs." (69)

Part of Sperber's argument is that human beings developed the capacity to reason in the

first place because we rely on communicating with others in order to navigate the world,

and doing so effectively requires that we trust each other. To illustrate this idea, Sperber et

al. (2010) present the following analogy:

When we walk down a street through a crowd of people, many at very close

quarters, there is a constant risk of inadvertent or even intentional collision. Still, we

trust people in the street, and have no hesitation about walking among them. Nor is

it just a matter of expecting others to take care while we ourselves walk carelessly.

We monitor the trajectory of others, and keep an eye out for the occasional

absentminded or aggressive individual, automatically adjusting our level of

vigilance to the surroundings. Most of the time, it is low enough to be unconscious

and not to detract, say, from the pleasure of a stroll, but it rises when the situation

requires. Our mutual trust in the street is largely based on our mutual vigilance.

Similarly, in communication, it is not that we can generally be trustful and therefore

need to be vigilant only in rare and special circumstances. We could not be mutually

trustful *unless* we were mutually vigilant." (364)

Here's how our nature as epistemically vigilant can explain our judgments in the

evidence-gathering cases presented by Pinillos: when we consider the case of high-stakes

John we will expect him to be more epistemically vigilant. As in the cases above when the

content of some information is important to us, we double-check, and thus we expect John

to double-check, as well. On the other hand, when considering the case of low-stakes Peter,

since the matter is of no importance to him we do not expect him to double-check or seek

out any more evidence. If John fails to be epistemically vigilant by failing to double-check or

gather more evidence, then we judge him to be failing to meet some normative epistemic

standard. As Sperber notes above, communicating with one another requires trusting each

other to do the epistemic work required of them, given their situation. If someone fails to

double-check or seek or more evidence when they should be epistemically vigilant, we are

liable to evaluate them negatively.

Now, one way of interpreting these results is that the epistemic standard that Peter

risks failing to meet is indicative of a failure to know. Miller (2014b), for example, argues

that experimental results of the kind that I cite above,

> nicely accord with PE. They show that when the stakes are high, e.g. when coffee
>
> lovers encounter evidence suggesting they should give up their precious hobby, they
>
> raise the evidential thresholds for knowledge, but they do not do so when the stakes
>
> are low or null. (261)

While some of these data are certainly in *accordance* with PE, they are not best interpreted

as being indicative of increased evidential requirements for knowledge. We have seen one

reason for this already in the results of Pinillos' experiments: when participants are

presented with a prompt asking them to evaluate how much evidence the subject in the

vignette needs to gather in order to have knowledge they report that John needs to gather

more evidence than Peter, but the explanations that the participants actually provided

indicate that their evaluations were not ones that they thought reflected possession of

knowledge. Another reason is that the experimental results are inconsistent: while it seems

that the effects of epistemic vigilance are predictable and robust, we do not always evaluate

vignette subjects as failing to meet an epistemic standard when they fail to gather evidence when stakes are high, and so we do not always hold others to higher epistemic standards when the truth of some information is more important to them.

The view I defend below is that participants' evaluations of vignette subjects in experiments testing PE are better interpreted as evaluations of epistemically responsible or irresponsible action. We evaluate someone as epistemically irresponsible when they fail to be epistemically vigilant in those situations in which they entrust themselves to others as a source of information. As I noted earlier, this explanation is a *kind* of error theory, namely one in which people will sometimes confuse judgments about whether one knows a proposition with judgments about whether one is acting in an epistemically responsible way. I elucidate this idea below.

4. How Epistemic Responsibility Judgments Explain the Data, Generally

So far I have argued that the PE theorist faces two significant problems when it comes to appealing to empirical support for her view. First, empirical results concerning the direct argument for PE do not consistently cohere with the expectations of PE; second, empirical results concerning the indirect argument indicate that participants hold vignette subjects to a higher epistemic standard as the stakes for those subjects increase, but that participants do not necessarily conceive of this standard in terms of a subject's knowledge. I have proposed that a standard that participants do hold vignette subjects to is an evaluation of whether the vignette subjects are being epistemically vigilant to the extent to that their situation demands. If they fail to be sufficiently epistemically vigilant, we are prone to evaluate them as acting in an epistemically irresponsible way: we evaluate them

as failing to be able to provide sufficient reason to believe that one knows what one is doing. The purpose of this section is to show how this interpretation of the above and other experimental results can explain why different experimental results seem to cohere with PE and why others do not. The view I argue for in this section is the following: we are more likely to evaluate one as epistemically responsible or irresponsible when one puts oneself in a position to be a trustworthy source of information. This is because our expectations of epistemic vigilance are a product of our expectations that one does enough epistemically to be able to provide others with reason to think that one does, in fact, know what one is doing. To continue Sperber's metaphor from earlier, we expect people to be more careful when they are walking down a crowded street than when they are walking by themselves.

Experimental results support this expectation. Consider first an experiment that shows results that are inconsistent with PE. Feltz and Zarpentine (2011) presented subjects with the following scenario:

> John is driving a truck along a dirt road in a caravan of trucks. He comes across what looks like a rickety wooden bridge over a [three foot/yawning thousand foot drop]. He radios ahead to find out whether other trucks have made it safely over. He is told that all 15 trucks in the caravan made it over without a problem. John reasons that if they made it over he will make it over as well. So, he thinks to himself, "I know that my truck will make it across the bridge." (690)

When prompted to rate how much they agreed with the statement that "John knows his truck will make it across the bridge" there were no significant differences in responses from participants when John was in the high-stakes situation (when there was a thousand

foot drop) than when he was in the low-stakes situation (when there was a three foot drop). This experiment, then, runs contrary to what we would expect were PE to be true.

The reason why Feltz and Zarpentine decided to test this particular kind of case is because it avoids running afoul of what they call an "attributor effect": according to this effect, "people are more reluctant to agree with third-person knowledge attributions than first-person attributions" (693). Now, it's not clear why, exactly, there should be an attributor effect at all: why, for example, would we be less likely to agree with fellow caravan-member Sarah if she were to make a third-person attribution that "John knows that the truck will make it across safely" than we are when John makes the first-person attribution that "I know that the truck will make it across safely"? Here is a possible answer: when John reasons to himself he does not put himself in a position to be treated as a trustworthy source of information, and we thus do not evaluate him as someone who we expect to be particularly epistemically vigilant. In the case of a third-person knowledge attribution, however, we *do* tend to evaluate a subject as someone who we expect to be vigilant, since the ascription of knowledge to someone at the same time implies that one is trusting that person as a knower. If we are going to be trusting John, then we expect that he has been vigilant in coming to acquire the evidence he has: we would expect him, perhaps, to make sure that his truck is in proper working order, in one way or another. On this interpretation, then, the attributor effect occurs because the act of making a third-person evaluation tends to bring along with it an evaluation of someone as a potentially trustworthy source of information.

This interpretation of the experimental results requires arguing that results that do show an effect of practical factors on knowledge evaluations are at least sometimes the

result of a conflation of judgments of epistemically responsible action with judgments of knowledge. Of course, participants in the above experiments do not typically respond that they find the subjects in the cases they are evaluating to be acting in an epistemically irresponsible way. I have suggested that one reason why this might be the case is that they do not have the option: when prompted to either evaluate a subject as *knowing* or not, or to state how much they agreed or disagreed with a *knowledge* attribution, participants are not likely to conscientiously consider the difference between having knowledge and acting in an epistemically responsible way. Again, as the Pinillos results show, once the option to explain one's judgments becomes available, appeals to knowledge are almost nowhere to be found. And although the theory that I have used to discuss evaluations of epistemically responsible action are complex, I do not think that evaluations of others as acting in epistemically responsible or irresponsible ways is an evaluation that is difficult, complex, or subtle. Instead, it is an evaluation that someone has failed to do everything that is expected of them epistemically in a given situation. My argument here is that this kind of evaluation is not necessarily indicative of one's knowledge[16].

   This interpretation can explain results from additional studies, as well. For instance, in another set of experiments, Joseph Shin (2014) tested whether a different kind or practical factor – the *time constraint* under which a subject had to make a decision –

---

[16] Although it is true that the vast majority of experiment participants are unfamiliar with the details of various normative epistemic judgments, this is not to say that they do not still make those judgments. After all, I have argued here that our responsibility judgments are based in our epistemic vigilance, which, it has been argued, is a basic and fundamental aspect of ourselves as reasoners. And, despite the majority of laypeople being unfamiliar with the notion of epistemic vigilance, we are nevertheless still vigilant. Thanks to an anonymous referee for raising this concern.

affected participants' evaluations of whether that subject had knowledge. Shin's hypothesis is the following:

> *Time-constraints-sensitivity*: With respect to some cases, all else being equal, people are more likely to ascribe knowledge to a subject with less time in her situation, than to a subject with comparatively more time. (160)

To test this hypothesis, participants were asked to evaluate one of two cases: one in which a subject had to make a decision under very short time constraints, or another in which a subject had to make a decision with plenty of time to decide, but which they had left to the very last minute. In each case the following facts obtain: vignette character nurse Sally has to choose between administering one of drugs A, B, or C to her patient; she seems to recall from her training that drug C is the best choice; she does, in fact, administer drug C; and she is correct that drug C is the best choice of treatment. In the low-time case, however, Sally has two minutes to decide which drug to administer, while in the more-time case, Sally has four months to decide which drug to administer, but puts the decision off, and has left herself two minutes to decide which drug to administer. When prompted to evaluate whether subjects agreed that Sally knew that drug C was the right choice of drug to administer, participants were, to a statistically significant degree, more likely to agree in the low-time case than in more-time case. Shin concludes that since time constraints are practical factors that do not have anything to do with traditional invariant factors that these results thereby support the PE thesis.

Shin is well aware that there are many ways that one could interpret the results of his experiment. One might, for instance, argue that since more-time Sally had four months to make her decision but decided to wait until the last two minutes, that she lacks

knowledge because she does not form her belief via a reliable process (she had, after all, plenty of opportunity to gather more evidence, but didn't). Or, one might argue that we are likely to judge that more-time Sally is epistemically *vicious* in some way, and that exhibiting epistemic vices in forming one's belief is inconsistent with that belief fulfilling the conditions to count as knowledge. These are certainly *possible* explanations, but Shin dos not find any particular reason to think that they are true: after all, when considering whether Sally is forming her belief reliably, time constraints do not seem to matter[17]; similarly, it's not clear, on what basis, we should find more-time Sally to be any more epistemically vicious than her low-time counterpart.

Although I think that Shin is correct that we should not think that nurse Sally is any less *reliable* in the more-time case, it certainly seems like more-time Sally is *epistemically worse*, in some way, than low-time Sally; that she is, in some ways worse, is borne out in the judgments of the participants in the study. But it's not obvious that the best explanation for why more-time Sally is epistemically worse-off than low-time Sally is because of a difference in *knowledge* between the two. The interpretation of the experimental results that I have presented thus far is that one can fail to meet a normative epistemic standard that does not pertain to a possession of knowledge. We might, then, say that low-time Sally and more-time Sally do not differ on the basis of their relevant knowledge, but rather on the basis of whether they are able to act responsibly on the knowledge they have.

What kind of reasons, though, might we be able to provide for the view that the difference between more-time Sally and less-time Sally is better explained by a difference

---

[17] If they were to matter, they would at most matter contingently, inasmuch as time constraints tend to lead to sloppiness with our belief-forming methods. There is, though, no inherent connection between being short on time and the reliability of a belief-forming process.

in epistemic responsibility and not a difference in knowledge? Appealing to traditional invariantist intuitions would beg the question: both Sallys have the same evidence for their respective beliefs, both seem equally reliable in the relevant senses, etc., and therefore, we might be tempted to say, they know the same things. This is obviously not a good enough argument for the PE-theorist. Explaining the experimental results in terms of PE, however, seems to overreach: without an antecedent commitment to some kind of interest-relativity, why should we think that *mere* time constraints in making decisions affect one's knowledge?

Again, the interpretation I have been proposing here can explain the results in the above cases. We can agree with Shin that a difference in *reliability* is not the best way to explain the epistemic difference between low-time Sally and more-time Sally; regardless, there does seem something right in the underlying reasons we might be tempted to appeal to in order to support this difference. After all, it seems like more-time Sally, given all her extra time, could have done *something* more to make a better decision, and that she *should not* have waited until the last two minutes to do so. This kind of behavior, though, is not necessarily characteristic of someone who lacks knowledge, but someone who is acting *irresponsibly*. This is, I think, a natural way that we would evaluate the actions of more-time Sally: we would, no doubt, ask why she had not done more research in the time she had available, why she would risk an inaccurate diagnoses by waiting until the last minute, etc. None of these evaluations, though, are necessarily indicative of more-time Sally's knowledge; instead, they are indicative of the fact that, given the constraints of her practical situation, more-time Sally was capable of seeking out more evidence, double-checking, etc., and failed to. More-time Sally is, then, an *epistemically irresponsible* agent,

30

but not necessarily someone who lacks the relevant knowledge. We can, then, explain what we think is wrong with the way that more-time Sally acted without having to appeal to anything about her knowledge. The results from Shin as well as those from Pinillos and Feltz and Zarpentine can, then, be explained in terms of the view that we tend to judge practical reasoning on the basis of how epistemically responsible someone is.

However, one might object that we cannot consistently judge that someone *knows* a proposition, but is acting on the basis of that proposition in an epistemically irresponsible way, while at the same time adhering to the RKP (or some similar principle). After all, the RKP states that if one knows a proposition then one has done enough to act on it in an appropriate way, and if one thinks that someone is acting in an epistemically irresponsible way, then to say that they can have knowledge while doing so just seems to be an implicit denial of RKP. One might also be suspect of the general conceptual plausibility of being able to act on one's knowledge in an epistemically irresponsible way. Let's look at these concerns in more detail.

5. Epistemically Irresponsible Action and the RKP

According to the view I have presented here, we hold people to two different normative epistemic standards when they act: one in terms of one's epistemic relationship with a proposition that they treat as a reason for acting, and another in terms of whether one is acting in an epistemically responsible way. Equivocating between these two kinds of judgments permits the indirect argument to go through on the basis of our evaluations of practical reasoning in situations that differ in terms of practical factors. Recall that premise (2) of the indirect argument stated:

2) Whether it's appropriate to treat *p* as a reason for acting depends on the

practical circumstances surrounding the truth of *p*.

Judgments that appear to support (2) can be indicative either as a judgment that pertains

to meeting the normative standard dictated by the RKP or one that pertains to the

normative standard of epistemically responsible action. We've seen that the empirical

evidence is ambivalent on the truth of (2); I have argued that this ambivalence is best

explained by the way that our epistemic evaluations are affected by our epistemic vigilance,

and the situations under which we are likely to evaluate someone as being sufficiently

epistemically vigilant.

Why not think, though, that these responsibility judgments in turn influence our

knowledge judgments? In other words, why not think that when we judge that someone is

acting in an epistemically irresponsible way we also thereby judge that someone does not

have knowledge? In order to argue against the indirect argument we would have to show

that we can make an evaluation that someone has knowledge that *p*, but is nevertheless

employing *p* in one's practical reasoning in an epistemically irresponsible way. How can

this be the case if RKP is true, and appropriate treatment of *p* as a reason for acting really is

dependent upon one's knowledge?

Here we need to distinguish two senses in which one can "appropriately" treat a

proposition as a reason for acting. One sense is that which is governed by a norm of

practical reasoning, like the RKP. In this sense, one inappropriately treats a proposition as a

reason for acting if one fails to know the relevant proposition, where this sense of propriety

pertains to one's epistemic relationship towards that proposition. Another sense, I have

suggested, is dictated by the expectations that one be epistemically vigilant: in order to

meet these demands one needs to have done enough to make sure that one is presenting oneself as a trustworthy source of information. In order to appropriately treat proposition *p* as a reason for acting, then, we might need to have a significant epistemic relationship towards *p*; if the RKP is correct, this relationship is knowledge. In order to act in an epistemically responsible way, on the other hand, we need to make sure that we are fulfilling our epistemic commitments towards others: we need to make sure that we are presenting ourselves as a potentially trustworthy source of information.

Importantly, while an evaluation that one does or does not know that *p* is an evaluation of one's epistemic relationship towards *p* itself, whether one is acting in an epistemically responsible or irresponsible way is not itself a function of one's epistemic relationship or cognitive attitude towards *p*[18]. To make sure that one is able to be relied upon as a trustworthy source of information in a given situation one might be required to double-check the source of one's evidence, in which case one would have to have some significant epistemic relationship with facts concerning one's belief-forming processes; or, one might have to have reason to believe that one has the knowledge one has, in which case one would have to have some significant epistemic relationship not towards *p*, but towards 'I know that *p*', etc.[19] There are potentially many different ways that one might be able to ensure that one is trustworthy given one's situation, a fact that is exemplified in what I argue below are natural ways that we chastise others for failing to act in epistemically

---

[18] If an evaluation of one's epistemic responsibility *were* an evaluation of one's relationship towards *p* itself then the view I am defending here would threaten to collapse into PE, since, if evaluations that one is acting epistemically irresponsibly on one's knowledge that *p* was an evaluation that one had failed in some way with regard to *p* itself, then it would seem that all instances of epistemically irresponsible action were then failures to know. Thanks to an anonymous referee for suggesting that I address this explicitly.

[19] I am assuming here the denial of the "KK principle", namely the principle that if one knows that *p* then one knows that one knows that *p*. I take it that it is not controversial to deny this principle.

responsible ways, e.g. cases in which we feel that someone should have double-checked the source of their evidence, or reflected on their knowledge, or calibrated their cognitive processes, etc.

We might apply Origgi's distinction between default and vigilant trust to make the distinction clearer: one might have to know a proposition in order to treat it as a reason for acting when there are no extenuating circumstances (i.e. in a default trust situation), whereas one would have to do more to make sure that one had the knowledge one portrayed oneself as having when the stakes are high (i.e. in a vigilant trust situation). Or, to continue Sperber et al.'s analogy, in order to successfully walk down the street by oneself one must possess a certain set of abilities, while in order to successfully walk down a crowded street one needs to both have these abilities, and make sure that one is employing them carefully. On the picture in which there are two different normative epistemic evaluations that pertain to treating a proposition as a reason for acting, there *is* an epistemic evaluation that is sensitive to the stakes of a situation. I have argued, however, that in order to make best sense of the empirical data, the kind of epistemic evaluation that is sensitive to stakes is one of acting in an epistemically responsible way, not one of acting in an "appropriate" way as is dictated by the RKP.

I have also suggested that cases in which one evaluates subjects as knowing but not acting in an epistemically responsible way are coherent, and, indeed, natural; more-time Sally, I suggested, was one of these cases. We might worry, though, that even this case is somewhat on the convoluted side, and is therefore not a great example of the "naturalness" of epistemic responsibility judgments. Once we have a notion of epistemically responsible action in place on our normative landscape, though, we can see that many other kinds of

everyday cases lend themselves to an interpretation in which a judged epistemic failing

pertains to responsibility and not knowledge.

Consider, for instance, an everyday case of *double-checking*. Say that it is very

important to you that I am on time to pick you up from the airport. I check the time of your

scheduled arrival online, through a reliable website, and commit the time to memory: 6

p.m. It is not difficult or tricky, and I easily remember it. When I meet you at the airport you

thank me for being on time and ask me how I knew when to pick you up. Upon giving my

answer, however, you chastise me: it was so important that I was there on time that I really

should have double-checked. I assure you that it was not difficult to remember, that the

website was reliable, that I've picked up many friends successfully in the past by using it,

etc. Still, you say, it would have been prudent to have checked at least once more, and

maybe even a few more times, especially given how important it was that you be picked up

on time.

Cases in which we feel that someone should have double-checked are common.

Furthermore, they seem to be cases in which, although it would be most natural to ascribe

knowledge – for instance, it seems natural to ascribe to me knowledge that your plane was

going to arrive at 6 p.m. – it seems equally natural to say that because of the importance of

being there on time I should have done something more epistemically. A judgment that I

violated some kind of normative requirement is not, in this case, most naturally explained

by a failure to have knowledge. The interpretation I have provided here supplies this

requirement: because the relevant truth was important to you, I should have been more

epistemically vigilant with regards to the source of my evidence, and thus the epistemic

demands for acting in an epistemically responsible way were higher[20,21].

Double-checking cases have, however, sometimes been interpreted differently.

Jessica Brown (2008), for example, presents a case of double-checking that she takes to

challenge the relationship between knowledge and epistemically appropriate practical

reasoning:

> *Surgeon*: A student is spending the day shadowing a surgeon. In the morning he
>
> observes her in clinic examining patient A who has a diseased left kidney. The
>
> decision is taken to remove it that afternoon. Later, the student observes the
>
> surgeon in theatre where patient A is lying anaesthetised on the operating table. The
>
> operation hasn't started as the surgeon is consulting the patient's notes. The student
>
> is puzzled and asks one of the nurses what's going on:
>
> Student: I don't understand. Why is she looking at the patient's records? She was in
>
> clinic with the patient this morning. Doesn't she even know which kidney it is?
>
> Nurse: Of course, she knows which kidney it is. But, imagine what it would be like if
>
> she removed the wrong kidney. She shouldn't operate before checking the patient's
>
> records. (176)

---

[20] That I did enough epistemically to have knowledge but not to act in an epistemically responsible way can make sense of why you and I might argue about whether I should, in fact, have double-checked. In response to your chastising, for instance, I might plead that I did, in fact, *know* that your plane would arrive at 6, and thus wonder why you would think I did anything wrong at all. This is unlikely to appease you: you might still insist that, despite knowing, I really should have double-checked. Making sense of this exchange does not require that we abandon RKP: the normative epistemic failing that you're accusing me of is one that pertains not to my knowledge, but to whether I acted on it in an epistemically responsible way.

[21] Again, here one is expected to exercise one's epistemic vigilance not towards another agent, but towards a non-agential source of evidence, namely the website one used to check the flight arrival time. And again, it seems natural to say that one can exercise one's epistemic vigilance in this way, since the website is a source of evidence and one trusts it.

Brown argues that it would not be appropriate for the surgeon to act on the basis of her knowledge, and that therefore knowledge is insufficient for appropriate practical reasoning. The PE proponent might instead interpret Brown's case as one in which the surgeon would not count as knowing which kidney was the correct one to remove without double-checking. As was the case with the double-checking case I presented above, however, without an antecedent commitment to PE it does not seem plausible to deny knowledge to the surgeon in this case. Brown presents the case as a counterexample to the sufficiency direction of the RKP, but her case only creates problems for a principle like RKP if we assume that the epistemic standard that the surgeon risks failing to meet should she not double-check is one that determined by what she knows. In other words, the "should" of "She shouldn't operate before checking the patient's records" is not necessarily the "should" of appropriateness, but rather the "should" of responsibility. We can maintain the RKP, then, by arguing that if the surgeon did not double-check, the epistemic standard she would have violated would be those required to act in an epistemically responsible way.

We can, of course, defuse the indirect argument by rejecting the idea that there is any important normative connection between knowledge and practical reasoning, or that the epistemic requirement for appropriate practical reasoning is not knowledge[22], or that there is no singular consistent epistemic requirement for practical reasoning. On the interpretation I am presenting here, however, we can reconcile the view that we can commit a normative epistemic failing when acting on a proposition we know, while maintaining the RKP (or something like it). If we accept this view, then we could also

---

[22] Defusing the indirect argument by rejecting the RKP is the strategy employed by Brown (2008).

account for some of the arguments that have been presented against the RKP that trade on intuitions that there are cases in which one has knowledge but is unable to act on it in an epistemically appropriate way.

Consider, for instance, the problem of what has been called *isolated knowledge*. Call a case of knowledge isolated if we know the relevant proposition but do not know any of the surrounding reasons that support its truth. Knowledge of trivia might be isolated in this way: I know that Lincoln was born in 1809, but I might not know anything else about either Lincoln or what happened in the year 1809. Jennifer Lackey (2007) argues that there is another kind of isolated knowledge, what she calls *isolated second-hand knowledge*. This is knowledge that we get through testimony from a reliable source, but, again, for which we do not know any of the reasons that make it true. Lackey provides an example in which a friend tells me that "The Diving Bell and the Butterfly was extremely moving." Since my friend is a reliable judge of movies, and since this proposition happens to be true, I come to know it. But I have not seen the movie myself, and know nothing about it. Lackey argues that it seems that there is something epistemically suspect about acting on isolated second-hand knowledge, and that since we can have isolated knowledge but we cannot act on it in an epistemically appropriate way, then a principle like RKP must be, strictly speaking, false.

However, we can also interpret cases in which one employs isolated knowledge in one's reasoning not as instances in which one violates a principle that links practical reasoning and knowledge, but in which one violates a norm of acting on that proposition in an epistemically responsible way. For instance, although I might have some isolated second-hand knowledge, because I do not know any of the reasons that make the relevant proposition true, and because I have no additional knowledge in the vicinity, I thereby lack

any reason to believe that I do, in fact, have the knowledge that I have. It might, then, be a characteristic of isolated knowledge that, because of its isolation, one is typically unable to act on it in an epistemically responsible way.

6. The Future of Pragmatic Encroachment

We started with something of a puzzle: the indirect argument for PE is valid, and appears sound. How, then, can we explain the fact that empirical results are ambivalent about whether people do, in fact, make knowledge judgments in ways that are consistent with PE? I have argued that the indirect argument rests of an equivocation between one normative epistemic evaluation that pertains to the requirements of RKP (or something like it), and one that pertains to the requirements of acting in an epistemically responsible way[23]. I have also argued that responsibility judgments are judgments that we naturally make because of the kinds of reasoners that we are, that responsibility judgments can explain the ambivalence of empirical results of experimental philosophy concerning cases that differ in practical circumstances, and that there is a coherent notion of irresponsible knowledge. The arguments I have presented here, however, are not intended as definitive refutations of pragmatic encroachment; indeed, one could accept everything that I have said here and still accept PE as a general thesis about knowledge. However, if we do accept

---

[23] One might wonder what kinds of experiments could be performed to test when one is evaluating an action as epistemically responsible or irresponsible. I take it that an experiment like that conducted by Pinillos (2012) is a good example, i.e. one in which participants are able to explain why they reported the judgment they did. If participant's' evaluations are consistently explained by judgments that do not pertain to knowledge, then this would be an indicated that participants are not making evaluations that they take to be reflective of a subject's knowledge in the first place. Thanks to an anonymous referee for suggesting I address this point.

everything that I have said here, there might not be a lot of reason left to think that PE is true.

Works Cited

Brown, Jessica. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm of Practical Reasoning." *Nous* 42.2: 167-189.

Buckwalter, Wesley and Jonathan Schaffer. 2013. "Knowledge, Stakes, and Mistakes." *Nous* 47.1: 1-34.

DeRose, Keith. 2009. *The Case for Contextualism Vol. 1 and 2.* Oxford: Oxford University Press.

Ditto, Peter H., Geoffrey D. Munro, Anne Marie Apanovitch, James A. Scepansky and Lisa K. Lockhart. 2003. "Spontaneous Skepticism: The Interplay of Motivation and Expectation in Responses to Favorable and Unfavorable Medical Diagnoses." *Personality and Social Psychology Bulletin* 29: 1120-1132.

Eagly, Alice H., Patrick Kulesa, Laura A. Brannon, Kelly Shaw and Sarah Hutson-Comeaux. 2000. "Why Counterattitudinal Messages Are as Memorable as Proattitudinal Messages: The Importance of Active Defense Against Attack." *Personality and Social Psychology Bulletin* 26: 1392-1408.

Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Feltz, Adam and Chris Zalpentine. 2010. "Do You Know More When It Matters Less?"

    Philosophical Psychology 23.5: 683-706.

Fricker, Miranda. 2007. *Epistemic Injustice*. New York: Oxford University Press.

Haidt, Jonathan. 2001. "The Emotional Dog and its Rational Tail: A Social Intuitionist

    Approach to Moral Judgment." *Psychological Review* 108(4): 814-834.

Ham, Jaap and Kees van den Bos. 2008. "Not fair for me! The influence of personal

    relevance on social justice inferences." *Journal of Experimental Social Psychology* 44:

    699-705.

Hawthorne, John. 2005. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Kahneman, D. and Frederick, S. 2002. Representativeness Revisited: Attribute Substitution

    in Intuitive Judgment. In T. Gilovich, D. Griffin, D. Kanheman (eds.). *Heuristics and*

    Biases: The Psychology of Intuitive Judgment. Cambridge: Cambridge University Press.

Lackey, Jennifer. 2007. "Norms of Assertion." *Nous* 41.4: 594-626.

Miller, Boaz. 2014a. "Catching the WAVE: The Weight-Adjusting Account of Values and

    Evidence." *Studies in History and Philosophy of Science* 47: 69-80.

---. 2014b. "Science, Values, and Pragmatic Encroachment on Knowledge." *European Journal*

    for the Philosophy of Science 4.2: 253-70.

Nagel, Jennifer. 2008. "Knowledge Ascriptions and the Psychological Consequences of

    Changing Stakes." *Australasian Journal of Philosophy* 86.2: 279-94.

Origgi, Gloria. 2012. "Epistemic Injustice and Epistemic Trust." *Social Epistemology* 26.2:

    221-235.

---. 2010. "Epistemic Vigilance and Epistemic Responsibility in the Liquid World of

    Scientific Publications." *Social Epistemology* 24.3: 149-59.

Pinillos, Angel. 2012. "Knowledge, Experiments, and Practical Interests." In Jessica Brown

    and Mikkel Gerken (eds.). *Knowledge Ascriptions.* Oxford: Oxford University Press, pp.

    192-219.

Pinillos, Angel and Shawn Simpson. 2014. "Experimental Evidence in Support of Anti-

    Intellectualism About Knowledge." In James R. Beebe (ed.). *Advances in Experimental*

    Epistemology.

Richter, Tobias, Sascha Schroeder and Britta Wohrmann. "You Don't Have to Believe

    Everything You Read: Background Knowledge Permits Fast and Efficient Validation of

    Information." *Journal of Personality and Social Psychology* 96.3: 538-558.

Schul, Yaacov, Ruth Mayo and Eugene Burnstein. 2004. "Encoding Under Trust and

    Distrust: The Spontaneous Activation of Incongruent Cognitions." *Journal of*

    Personality and Social Psychology 86(5): 668-679.

Shin, Joseph. 2014. "Time Constraints and Pragmatic Encroachment on Knowledge."

    Episteme 11.2: 157-180.

Sperber, Dan. 2013. "Speakers are Honest because Hearers are Vigilant: Reply to Kourken

    Michaelian". *Episteme* 10.1: 61-71.

Sperber, D., F. Clément, C. Heintz, O. Mascaro, H. Mercier, G. Origgi, and D. Wilson. 2010.

    "Epistemic vigilance". Mind & Language 25(4):359–393.

Sperber, Dan and Hugo Mercier. 2011. "Why Do Humans Reason? Arguments for an

    Argumentative Theory." *Behavioral and Brain Sciences* 34: 57-111.

Sripada, Chandra and Jason Stanley. 2012. "Empirical Tests of Interest-Relative

    Invariantism." *Episteme* 9.1: 3-26.

Stanley, Jason. 2004. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Stanley, Jason and John Hawthorne. 2008. "Knowledge and Action." *The Journal of*

    Philosophy 105(10): 571-90.

Weatherson, Brian. 2011. "Defending Interest-Relative Invariantism." *Logos and Episteme*

    2.4: 591-609.