

# Some Recent Work in Experimental Epistemology<sup>1</sup>

N. A. Pinillos (ASU)

## 1. Introduction

In this paper I survey some recent developments in experimental philosophy and discuss their bearing on two leading theories in epistemology: Contextualism and Interest Relative Invariantism. In the first part of the paper, I discuss some ways experimental philosophy may be relevant to assessing contextualism and IRI. In the second part, I discuss and critique some of the recent experimental work. I note that this paper is not a survey of all work on “experimental epistemology”. I will not discuss research on the instability of epistemic intuitions (including Gettier intuitions)<sup>2</sup>, the distinction between knowing-how and knowing-that<sup>3</sup>, or the (purported) moral dimension of knowledge<sup>4</sup>. Readers interested in a critical review of the first area are encouraged to read Jennifer Nagel (2007) and readers interested in the connection between morality and mental state attributions (including knowledge) are encouraged to read Joshua Knobe (2010).<sup>5</sup>

## 2. Contextualism and Experimental Philosophy

---

<sup>1</sup> I wish to thank Wesley Buckwalter, Josh May, Jennifer Nagel, Jonathan Schaffer, Jonathan Weinberg and an anonymous referee for their help in sorting out many of the issues raised here.

<sup>2</sup> Weinberg et al. (2001), Nichols et al. (2003), and Swain et al. (2008).

<sup>3</sup> Bengson et al. (2009)

<sup>4</sup> Beebe and Buckwalter (forthcoming)

<sup>5</sup> I also do not discuss empirical work on the role of confidence and stakes in decision making that might be brought to bear on folk attributions of knowledge. Readers interested on this topic are encouraged to read Nagel (2008, 2010).

Contextualism is the thesis that relative to different contexts of use, knowledge-attributing (or denying) sentences may express distinct contents (where the various contents will vary in epistemically interesting ways). On one common way of carrying out this proposal, an attribution of the form ‘A knows P’ can vary in content depending on the epistemic standards at play in the context of use. The sentence can be true relative to a low standards context but it can be false relative to a high standards context.

I start off by asking whether experimental philosophy can, in principle, shed light on contextualism. Given the methods characteristic of recent work in experimental philosophy, this will just be the question of whether the use of controlled experiments to investigate people’s use of ‘knows’ and related forms can shed light on contextualism. Clearly, the answer to this must be ‘yes’ if ordinary people’s use of ‘knows’ is relevant to assessing contextualism. For it is hardly credible to hold that although participants interested in the contextualism debate should pay close attention to the behavior of ordinary people, they should not pay attention to this information if it is collected using standard scientific methods.<sup>6</sup>

Now so long as ordinary people use ‘know’ competently across contexts, then the methods of experimental philosophy can be helpful in assessing contextualism. For we can use these methods to detect if there are differences in content expressed across the

---

<sup>6</sup> A worry about ecological validity: Situations in which subjects fill out surveys in classrooms or online (conditions under which social psychology experiments are often conducted) may be importantly different from everyday situations where ‘knows’ is used. I am not sure if this worry can be developed into a legitimate objection to particular experimental results. Thanks to an anonymous referee here.

contexts. If we can detect differences, then this would constitute some evidence for contextualism.

One can resist this conclusion by insisting that ordinary people are not competent users of ‘know’ or at least less competent than philosophers. As a consequence, investigating the behavior of ordinary people is not likely to yield significant insights.<sup>7 8</sup> This response, however, is not very promising. Arguably, context sensitivity is rampant in natural language. For other context sensitive construction like ‘I’, ‘you’, ‘tall’ and many others, we do not suppose that ordinary people’s behavior with respect to these words is a less reliable indicator of their context sensitivity than the behavior exhibited by philosophers. Barring some special argument, we should expect this to also hold for ‘knows’.<sup>9</sup>

It should not be surprising then that many philosophers think that folk use of ‘knows’ is relevant to assessing contextualism. Keith DeRose (2006), for example, is explicit about this when he says:

---

<sup>7</sup> See Antti Kauppinen (2007) for a criticism of experimental philosophy along these lines. Note, however, his claim is not about contextualism or context sensitivity in particular.

<sup>8</sup> A referee points out that even if ordinary people are in general less competent than philosophers, their judgments might still count as evidence in certain situations (perhaps because in some cases, philosopher’s judgments may be a product of bias associated with a quirk in philosophical training or because the task is one we don’t expect ordinary people to make a mistake).

<sup>9</sup> Please see Jessica Brown (forthcoming) for a more in-depth discussion of whether contextualism makes predictions about ordinary uses of ‘knows’ and how this relates to experimental philosophy. She concludes that contextualism does make predictions about ordinary usage and that experimental philosophy is relevant for assessing contextualism.

The best grounds for accepting contextualism concerning knowledge attributions come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk: What ordinary speakers will count as “knowledge” in some non-philosophical contexts they will deny is such in others (pg. 316).

The importance of looking at folk linguistic behavior is also emphasized by Peter Ludlow (2003). In that paper, he gathers evidence for contextualism by conducting a series of Google internet searches. This type of corpus research, like others in the literature, assumes that linguistic behavior of the folk is relevant to assessing contextualism.<sup>10</sup>

### **3. IRI and Experimental Philosophy**

In what follows, I assume that IRI (Interest Relative Invariantism) is just the conjunction of the following two claims: Invariantism about ‘knowledge’ is true and knowledge is sensitive to stakes.<sup>11</sup> The claim that knowledge is sensitive to stakes is just the idea that whether an agent who believes P also knows P may depend on the practical costs (for that agent) of being wrong about P: when the stakes are high, the epistemic standards for

---

<sup>10</sup> Pinillos (2011) argues that for certain expressions, whether or not they get a variantist semantics (contextualism/relativism) may be an external matter in the sense of Putnam and Burge. If so, appealing to ordinary behavior in the usual way may not be of much help.

<sup>11</sup> Some theorists may understand IRI to be a weaker thesis so that it is agnostic as to whether contextualism is true. To make the exposition simpler, I will assume the stronger interpretation. Nothing hangs on this, since the part of IRI that is relevant for us now is just that knowledge is sensitive to stakes.

attaining knowledge may be higher.<sup>12</sup> Now, since invariantism entails that contextualism is false and we saw that experimental philosophy can be relevant to assessing contextualism, it will also be relevant to assessing IRI. The more interesting question is whether experimental philosophy is relevant for establishing the second part: the claim that knowledge is sensitive to stakes. That is, we are now interested in the question of whether the scientific study of ordinary people's use of 'knowledge' is relevant to assessing the claim that knowledge is sensitive to stakes. (I point out that if it turns out that the scientific study of ordinary people's judgments are irrelevant to epistemology because the study--scientific or otherwise--of ordinary people's judgments is irrelevant, then not only will this be problematic for some areas of experimental epistemology, but it will also be problematic for standard armchair epistemology. As we will see below, a lot of practitioners appeal to ordinary uses of 'know' to support their theories.)

If we treat the claim that knowledge is sensitive to stakes as a substantive statement about knowledge similar to how we treat many claims in the physical sciences and mathematics, folk judgments would be largely irrelevant in assessing its merits.<sup>13</sup> Appealing to folk judgments to settle disputes in the physical sciences and mathematics is rare at best.

Some may worry, however, that the analogy between IRI and theses in the physical sciences and mathematics is strained. Goldman (2007) points out that the

---

<sup>12</sup> See Weatherson (MSab) for a variation which focuses on odds as opposed to stakes. Using this as background, he develops a criticism of the experimental work I discuss below. I was not aware of his criticisms until this manuscript was nearly finished.

<sup>13</sup> Kornblith (2007)

methods of the disciplines are very different. Epistemologists, but not scientists, make heavy appeal to intuitions concerning fantastical scenarios.

Furthermore, even if we grant that knowledge is a natural kind (and to be studied scientifically like we study water or helium), I think it may still be useful to appeal to folk judgments. It is commonly held that an investigation into a natural kind will begin by identifying the obvious instances of the kind.<sup>14</sup> It is less often noted, however, that what **is considered** ‘obvious’ can vary from person to person and across time. Many of us find it obvious that racist and sexist ideas are false, but as history reveals, many people in the grips of bias (even those that may have been considered “experts”) have thought (and many continue to think) otherwise. Relatedly, theorists often point out that philosophers’ judgments about purported cases of knowledge might be subject to bias.<sup>15</sup> A useful corrective against this danger may be to seek out judgments from people who are ignorant of the various theories of epistemology (non-philosophers) and so are less likely to exhibit those biases. Hence, even if we think that the study of knowledge is just like the study of other natural kinds, the study of folk judgments may still be valuable.

Now, instead of thinking of IRI as a thesis of the natural sciences, we can think of it as a metaphysical claim. Jessica Brown (forthcoming), for instance, suggests drawing an analogy between IRI and incompatibilism: the metaphysical thesis that free will is incompatible with determinism. Brown points out that the truth of this thesis is consistent

---

<sup>14</sup> Kornblith 2007.

<sup>15</sup> Goldman (2007) writes: ‘philosophers are leery about trusting the intuitions of other philosophical analysts who have promoted general accounts of the analysandum, e.g., knowledge or justification. Commitment to their own favored account can distort their intuitions, even with respect to their own (pre-theoretical) concept.’ (pg. 15)

with ordinary people having judgments that are not in accord with the thesis. She suggests something similar may be said about IRI and folk judgments about knowledge.

At this point, however, we should distinguish the strong claim that a metaphysical thesis entails some proposition about ordinary speakers from the weaker claim that the thesis can gain support from evidence concerning the behavior of ordinary people. For example, Dean Zimmerman (2007) thinks that presentism (the thesis that things in the present are all that exist) is supported by the observation that common sense beliefs accord with the thesis. But this is compatible with presentism not entailing or requiring anything about what common sense tells us. Understanding IRI on this model, we can still think of experimental philosophy as providing some evidence for or against IRI, all the while agreeing with Brown that the thesis does not entail or require any claim about how the folk behave.

The reasoning we just went through can be extended without appealing to the notion of common sense. Suppose that through the methods of experimental philosophy we discover that people are disposed to make knowledge ascriptions in some very simple and ordinary situations in such a way that if the ascriptions turn out to be true, they conform to a pattern of facts cleanly predicted by IRI. In this case, unless we have some special reason to think that the subjects are making some error, we should presume that the ascriptions are true, and conclude that we have found evidence in favor of IRI (In simple and ordinary cases, we don't often challenge people's attitude ascriptions---why should we challenge them here?). This is yet another way in which experimental

philosophy can be relevant to assessing IRI, even if we understand IRI as a thesis that is largely about the metaphysics of knowledge.<sup>16</sup>

It is worth pointing out, however, that appealing to common sense (or to what ordinary people believe) in order to assess a metaphysical thesis (or a thesis “about the world”) may be a sign that research is at an early and speculative stage. As the discipline matures, reliance on ordinary judgments and hence experimental philosophy might become less relevant.

Finally, instead of thinking of the idea that knowledge is sensitive to stakes on the model of science or metaphysics, we can think of it as a semantic or conceptual thesis. On this interpretation, competent users of the word (or concept) ‘know’ will use that word (and its related forms) in accordance with the idea that knowledge is sensitive to stakes.<sup>17</sup> On this understanding, we should expect that experimental work will be relevant to assessing IRI.

When theorists discuss IRI, they are seldom explicit about the exact role of ordinary judgments are supposed to play in the theory. However, defenders of IRI often do appeal to ordinary judgments and practices to support their claims. Jason Stanley

---

<sup>16</sup> We could also be justified in accepting the methodological assumption that folk judgments matter to epistemology without having a clear idea as to why this might be. If asked for justification, we may simply say that the assumption has been fruitful in research.

<sup>17</sup> Goldman (2007). The proposal on the table is compatible with the idea that the concept ‘knowledge’ is a primitive notion. If a concept is sensitive to some parameter, it does not also mean that the concept admits of a non-trivial analysis (involving necessary and sufficient conditions). Furthermore, and following a theme from Goldman, we should also consider the possibility that experimental work can give evidence for or against the notion that there is a single and unique knowledge concept shared by all humans.

(2005), for example, thinks that ‘our intuitive reactions’ to certain vignettes provide ‘a prima-facie case for the thesis that knowledge is not just a matter of non-practical facts, but is also a matter of how much is at stake’ (pg. 6). John Hawthorne, in defending IRI, claims that the inclination to claim that lottery propositions are not known (e.g. the inclination to say that one does not know she will lose the lottery tomorrow) is ‘not merely a datum about the inclinations of philosophers’. He also claims the inclinations are ‘widespread’ and this is a fact that is ‘uncontroversial’ (pg. 8).

Similarly, Hawthorne and Stanley (2008) remark that ‘ordinary folk appraisals of the behavior of others suggests that the concept knowledge is intimately intertwined with the rationality of action’ (pg. 571). And as is well appreciated, the presumed intimate connection between the rationality of action and knowledge is a core tenet of IRI. It looks then like many defenders of IRI accept that folk behavior is relevant to establishing their claims.<sup>18</sup> If they are right, then it is plausible to think that experimental philosophy can provide valuable data that is relevant to establishing their claims.

#### **4. Recent Experiments**

In the previous sections, we surveyed some theoretical grounds for thinking that experimental work is relevant to assessing both IRI and contextualism. In this section we begin to examine some of the experimental literature. The organization of the discussion is as follows. First I will simply describe the experimental work (4.1) and how it has been taken to be relevant to epistemology. I will then discuss how contextualists (4.2) and IRI

---

<sup>18</sup> It is important to not overstate this point. Many defenders appeal to theoretical principles (connecting knowledge and action, for example) to defend their theses. I am not at all claiming that the theorists believe that ordinary judgments make or break their theories.

defenders (4.3) may respond to these findings. Finally, I discuss some recent work aimed to support contrastivism, a particular brand of contextualism (4.4). Two points are worth mentioning before I continue. First, I have been selective in choosing which experiments to discuss. My choice reflects an attempt to find a unifying theme in the type of criticisms we can raise against tempting interpretations of these results, and an attempt to incorporate the criticisms that have recently surfaced in blogs and papers in progress (this is all that is available at this time). Second, the criticisms I raise against the experiments should not be seen as aiming to establish that the experiments are not externally valid. Instead, the criticisms can be seen as fruitful avenues to pursue, which favor the idea that the experimental results are not trouble for IRI or contextualism. To fully vindicate these points, further experiments would have to be conducted (and all parties should agree to this since the experimental philosophers also agree that their work is preliminary).

#### **4.1 Experimental Work**

Consider the following pair of now famous vignettes constructed by Keith DeRose (1992, pg. 193). The relevant difference between them are the stakes, or practical interests, at issue.

*Low* My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this

case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” I reply, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.”

*High* My wife and I drive past the bank on a Friday afternoon, as in [Low], and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, “Banks do change their hours. Do you know the bank will be open tomorrow?” Remaining as confident as I was before that the bank will be open then, still, I reply, “Well, no. I’d better go in and make sure.”

The intuitive responses to these cases, acknowledged by many philosophers, is that in Low, the agent's utterance of 'I know it'll be open' is true although her denial 'No [I don't know]' in High is also true.<sup>19</sup> Now, different theorists draw different conclusions from this. IRI theorists take these intuitions to directly reflect the idea that knowledge is sensitive to stakes. Others take these intuitions to support contextualism. They might claim that when stakes are high the standards for knowledge operative in the conversation are raised and this explains the difference in judgments. Jennifer Nagel, on the other hand, takes the intuitions to not support IRI over classical invariantism (where knowledge is not sensitive to stakes).<sup>20</sup> And Patrick Rysiew thinks these reactions reflect pragmatic features of the conversational context and consistent with invariantism.<sup>21</sup> Despite the differences, many are in agreement about the data. The data is that the intuitive reactions to these cases are as stated.

It turns out, however, that there is reason to reject this assumption. Three separate studies failed to confirm that ordinary people have the predicted responses to the bank vignettes (and similar thought experiments).<sup>22</sup> It might be thought then that these new

---

<sup>19</sup> Jonathan Schaffer (2006), however, holds that the vignettes properly modified to only reflect differences in stakes does not elicit the predicted intuitions (in him).

<sup>20</sup> Nagel (2008)

<sup>21</sup> Rysiew (2007)

<sup>22</sup> These are the three studies: Feltz and Zarpentine (forthcoming), May et al. (forthcoming), and Buckwalter (forthcoming). To be clear, these authors suppose that they are giving evidence that undercuts some motivations for IRI or contextualism. They don't take themselves to be refuting those theories and do not think their studies are the last word on the matter. They all recommend further studies. I should also mention that Neta and Phelan (ms) have done experiments that concern evidence and practical interests. Their studies, however, do not directly concern the concept 'knowledge'.

results are very relevant to concerns in epistemology. Buckwalter (forthcoming), for example, takes his study to call into question ‘the conventional wisdom that contextualism and SSI [IRI] really have folk attributions of knowledge on their side’ (pg 12). Similarly, Jonathan Schaffer and Joshua Knobe (forthcoming) remark that ‘it seems as if the current debate over contextualism...is based on a myth’ (pg 10). Furthermore, they go on to accept the ‘interpretation of the data on which stakes have little or no impact on knowledge ascriptions’ (pg. 11). As a result of this, they think the results give evidence against IRI. Jessica Brown (forthcoming), on the other hand, comes to a very different conclusion. She thinks that the result impugns only traditional contextualism about knowledge but not IRI. Independent on who is right, it is clear then that many theorists are taking these new studies to be highly relevant to issues in epistemology.

Before we assess these studies, it is worth pointing out that IRI makes an “existential” claim.<sup>23</sup> The theory doesn’t say that stakes will always play a role in knowledge. It just says that sometimes it does. So if for a pair of vignettes some studies fail to detect the effect of stakes and if in a totally different pair some studies detect the effect of stakes, we should not conclude that the evidence is inconclusive. Rather, we should conclude (all things being equal) that the totality of the evidence favors IRI. This should be kept in mind as we discuss studies that fail to detect the effect of stakes and different studies that were able to detect it.

Now, in assessing the merits of these studies, we must be sure to separate three different types of criticisms. First, there is the idea that experimental philosophy is, in

---

<sup>23</sup> Weatherson (MSa)

principle, irrelevant to epistemology. Second, we can grant that experimental philosophy is, in principle, relevant to epistemology but argue that direct surveys of the type characteristic of experimental philosophy do not get at the right sorts of intuitions that are at issue in philosophical discussions.<sup>24</sup> Third, we might have no problem with survey methods in general but instead may wish to criticize specific details of experimental design. In many cases, these criticisms can lead to improved studies. The sorts of problems we will raise below fall in the third camp.

The design of these studies vary slightly but many of them consist of presenting subjects with one of the two bank vignettes (“Low” or “High” stakes) modified somewhat and asking subjects about the extent to which they agree (or disagree) with either the ascription (or the denial) of knowledge made by a protagonist in the vignette.<sup>25</sup> It was found that (in general) there were no statistically significant differences found between responses given to the two vignettes.<sup>26</sup> This then looks like some remarkable evidence that folk use of knowledge does not accord with what many philosophers were

---

<sup>24</sup> Cullen (2009), Bealer (1999).

<sup>25</sup> In May et al. (forthcoming) the subjects were asked not whether an assertion or denial made by a protagonist is true but rather whether a sentence about knowledge (as displayed outside the text of the vignette) is true. This sort of experiment is less relevant for contextualism since it is not clear whether the sentence should be interpreted within the context of the vignette. The experiment is more relevant for an invariantist thesis. And here, May et al. did report a statistically significant difference in answers given to the High and Low stakes vignettes. May, however, reports to me that he doesn’t take this evidence to favor IRI because subjects still tended to agree with the positive knowledge ascription in both cases. I still take this to be evidence for IRI (since if the experiment has no confounds, stakes is playing some significant role in knowledge ascriptions) but perhaps not very strong evidence.

<sup>26</sup> These were between-subjects studies. Some within-subjects studies were run. I will not discuss these here.

predicting. The results are particularly noteworthy because the researchers used nearly the same vignettes traditional philosophers use to motivate their theories.

I would like to focus on a feature that is common to almost all vignettes in the studies. With the exception of one experiment (which we will come to soon), all of them contain positive ascriptions of knowledge made by protagonists in both the Low and High vignettes. This is a key difference between the experiments and the original bank thought experiment utilized by philosophers. It is plausible to think that this feature of the experimental design may be responsible for why the probes are not detecting sensitivity to stakes. However, the details of this critique will depend on whether contextualism or IRI are true. Consequently, I discuss each of these perspectives in turn.

Before continuing, I wish to issue a caveat. Insofar as IRI and contextualism are concerned, what matters is whether the new results are problematic for those theories. In the next sections, I describe how those theories have the tools to explain the experimental results. If this is right, then the new results may not be that problematic for those theories. But to be fair, the authors of the experiments do not claim that those theories cannot explain the results. So even if I am right, it does not mean those authors have made a mistake. Again, my goal is mainly to assess the extent to which the studies create trouble for IRI and contextualism. Presumably, this is the question that matters the most for contemporary epistemology.

#### **4.2 A Response on Behalf of Contextualists**

To start, I point out that contextualists need not accept stakes matter at all. They may hold that other features of the context are responsible for fixing the content of ‘knowledge’

constructions. Nonetheless, contextualism seems ready made to explain how stakes can affect knowledge.<sup>27</sup> So in what follows I provide a contextualist diagnosis of the experimental philosophy results. The discussion here closely resembles some ideas in DeRose (forthcoming).

Knowledge ascriptions are subject to rules of accommodation.<sup>28</sup> Here's an example of this type of rule at work (I assume here that different contexts can fix different "epistemic standards", but this is not required). Suppose that at a particular time, there are certain epistemic standards operative at a context of utterance so that some knowledge ascribing sentence S expresses a false content relative to that context although nobody utters S. Now imagine a situation that is just like the previous one except that somebody actually utters S. Accommodation allows for the standards to be lower in the second situation so that S expresses a true proposition relative to the "updated" context. Uttering S can change the context resulting in the expression of a true proposition. As I explain below, the contextualist can appeal to this rule to show that the experiments cited do not give great evidence against contextualism.

The experimental results would create trouble for contextualism if the thesis predicted that (a) the vignettes invoke different epistemic standards, and (b) these standards are so different that the sentence uttered (positive knowledge ascription) 'I know the bank will be open in Saturday' would be true relative to one context (Low stakes) but false relative to the other one (High stakes). Although, a contextualist may agree that the differences in stakes give rise to distinct epistemic standards (enough to

---

<sup>27</sup> See DeRose (1992).

<sup>28</sup> Lewis (1979)

warrant differences in content for the sentence in question), they can also hold that through a rule of accommodation, the positive ascription of knowledge in the high stakes vignette has the effect of lowering the epistemic standards. The result of this is that the overall epistemic standards across the vignettes are not different enough to elicit the expected judgments. If this is right, then the results of these experiments are perfectly consistent with contextualism.<sup>29</sup>

Let us turn now to the one experiment where the high stakes vignette contains a *denial* of knowledge.<sup>30</sup> The result of this experiment is that people tended to agree with the denial of knowledge in the High condition and also with the ascription of knowledge in the low stakes case. This looks like it bodes well for contextualism except that the statistics concerning the differences in attitudes were not statistically significant. Some may interpret this to be evidence against contextualism.

By way of response, I emphasize that we should be cautious before drawing conclusions from a single study reporting a null result (a failure to find statistically significant difference in responses). In fact, since the p-value for this study is .23, the only thing we can conclude is that if the null hypothesis is true (stakes have no effects), then we should expect the sort of result we got in this experiment 23% of the time. But from this fact we can hardly draw the conclusion that there is good reason to accept the null hypothesis (I am not suggesting that the authors of the study draw this conclusion).

---

<sup>29</sup> This discussion assumed a version of contextualism where it makes sense to speak of varying epistemic standards. This is not necessary, however. A rule of accommodation can work on other contextual parameters that are relevant for fixing the content of knowledge sentences.

<sup>30</sup> Feltz and Zarpentine (forthcoming).

There is one more important point the contextualist can make in response to the studies. Contextualists can think that stakes affect folk use of ‘knowledge’ in virtue of the fact that high stakes make salient the possibility of error. And although some experiments (from the three papers we are discussing) tended to show that salience of error play no role in knowledge ascriptions, a recent study by Schaffer and Knobe (2010) gives some evidence to the contrary. S&K claim that a problem with previous experiments was that the possibility of error was not made salient enough to the experimental subjects. The contextualist can then take this result not only to directly support contextualism, but now also has some reason to question the assumption that the difference in stakes (in the bank vignettes) were made salient enough to the experimental subjects.<sup>31</sup>

#### **4.3 A Response on Behalf of IRI.**

I now want to discuss some ways in which an IRI defender can try to account for the experimental results. Here, I will also be concerned with the two variations in the experiments cited (one with the denial of knowledge in the high stakes condition and the other with the positive attribution).

Concerning the studies containing the denial of knowledge in the High conditions, the IRI defender could give the same response as the contextualist. The failure to find an effect there (the obtaining of a null result) should not be taken by itself to be very good evidence against their theory (I am not suggesting that any of the authors have claimed otherwise). Further work would need to be done to establish this claim. However, the contextualist responses to the other studies (containing the positive attributions of

---

<sup>31</sup> Curiously, S&K’s vignettes do not contain any ascriptions of knowledge.

knowledge in the High conditions) is not available to IRI defenders, since rules of accommodation assume contextualist machinery.

In the studies in question, the protagonist in the high stakes vignette asserts that he knows the bank will be open on Saturday. In low stakes, the protagonist asserts the same thing. The main finding as we recall (which might be thought to create trouble for IRI) was that there was no statistically significant difference in the level of agreement that subjects displayed across the two vignettes.<sup>32</sup>

Now, whether these results spell trouble for IRI will crucially depend on the assumption that the evidence the protagonist in the vignette has available to him (for the proposition that the bank will be open on Saturday) across the vignettes is the same. After all, the relevant probe here should concern how knowledge ascriptions vary with stakes—not how they vary across different amounts of evidence available. Keeping this in mind, let us see how the experimentalist might ensure that the perceived evidence available is kept constant across the vignettes.<sup>33</sup> The strategy here is to construct pairs of vignettes that have identical texts except for the description of what is at stake. This seems like a reasonable way to ensure the perceived evidence is the same across the vignettes.

However, it will be relevant in what follows that the amount of evidence available to the protagonist is not made wholly explicit—or at least it is open ended.<sup>34</sup> The only part of the vignette that speaks directly to the evidence is when the protagonist claims “I was

---

<sup>32</sup> As mentioned before, one of the results due to May et al (forthcoming) did turn up a statistically significant difference between the High and Low conditions.

<sup>33</sup> This is an especially difficult thing to do in a between-subjects study where subjects do not see both vignettes.

<sup>34</sup> See also DeRose (forthcoming)

there last week (and I know the bank will be open on Saturday)”. It is plausible that this sentence could be associated with a variety of distinct impressions concerning the strength of evidence at issue. For instance, readers could think that the protagonist thinks the bank will be open on Saturday because she saw a sign that said the bank will be open the following Saturday, or because she saw a sign which says the bank is (in general) open on Saturdays, or because she asked an employee (or perhaps a customer) etc. These interpretations correspond to distinct impressions about the quality of evidence the protagonist possesses for the proposition that the bank will be open on Saturday.

Keeping these ideas in mind, the IRI defender can begin to make the case that the perceived evidence available to the protagonist is not the same across the vignettes. If this is right, then we should not trust the study to have revealed that stakes play no role in knowledge attributions. The idea comes from how stakes and evidence interact to yield knowledge. Consider a simple model of how this might work (assuming that knowledge is sensitive to stakes). Suppose that due to the open-endedness of the vignettes, the evidence available to the protagonist could be interpreted as either strong or weak (depending on the other clues available to the subject). Now in low stakes, having weak evidence is enough for having knowledge that the bank will be open on Saturday. But in high stakes, knowledge requires strong evidence. Here’s then what might very well happen when people see in the High condition that the protagonist claims he knows. If the protagonist claims he knows, subjects will have some reason to think that what he said is true. In light of this, subjects will then have some extra reason to think that the evidence available is in fact strong. On the other hand, in the Low condition, when the

protagonist claims he knows, there is no similar reason to think that the evidence is strong (since low strength evidence is enough for knowledge). Hence it is plausible to hold that subjects have different impressions concerning the strength of evidence available to the protagonist across the vignettes. If this is right, the results do not impugn IRI. For it is not the case that the only relevant variable that changes across the vignettes is perceived stakes. By following through on this line of reasoning, the IRI defender may very well be able to explain the results.<sup>35</sup>

Finally, I should mention some experimental work carried out which may support IRI. Pinillos (forthcoming) reports on experiments he conducted that have less of an issue about keeping the evidence constant across the probes. In these studies, Pinillos presents subjects with a pair of stories about an agent that has to come to a certain conclusion by gathering some evidence. The stories vary with what is at stake for the protagonist (“high stakes” and “low stakes”). Pinillos then asked subjects to say how much evidence the agent in question needs to gather before he knows the conclusion. It was discovered that the responses to the high stakes case were much higher than in the lower stakes case. Along with other results presented in that paper, Pinillos takes this data to provide experimental support for IRI.

#### **4.4. Schaffer and Knobe**

I want to now focus on some recent work by Schaffer and Knobe (2010), which seems to support a particular (non-traditional) type of contextualism: contrastivism. According to contrastivism, claims that may look on the surface to be simply about an agent knowing

---

<sup>35</sup> **Related objections** can be found in the Certain Doubts Blog: [http://el-prod.baylor.edu/certain\\_doubts/?p=1999](http://el-prod.baylor.edu/certain_doubts/?p=1999). See especially, entries by Jennifer Nagel.

some proposition at a time are in reality claims about that agent knowing that proposition (at the time) *rather than some other proposition*. In the words of S&K:

Knowledge so conceived is not a two-place relation between a subject and a proposition, but is rather a three-place relation between a subject, a proposition (the fact), and a contrast proposition (the foil). All knowledge takes the form: *s* knows that *p* rather than *q*. (pg 14)

Contrastivism is still a contextualist theory in that utterances of the form ‘A knows that P’ can vary in content and truth value in epistemically interesting ways. But this view differs from some traditional contextualist theses in that, for example, it is denied here that ‘knowledge’ denotes different relations depending on the context of utterance.

Now, S&K provide three experiments that not only support contextualism in general but support contrastivism in particular (as opposed to traditional contextualism). I will focus on one of these experiments (perhaps the strongest one).

In this probe, subjects received one of two stimuli. They both contained the following “jewel-thief” vignette:

Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter

forgot to wear gloves. He also forgot about the security camera. Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and indentified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information.

Now on the first condition, this passage was followed by this (thief-contrast):

Everyone is now asking the big question: Who stole the rubies? The news reporter is about to write a story about Mary. He is wondering if Mary now knows who stoles the rubies. He writes: "Mary now knows that Peter stole the rubies"

Please tell us whether you agree or disagree with the news reporter's claim "Mary now knows that Peter stole the rubies."

But subjects in the second condition read the following passage instead (jewel-contrast):

Everyone is now asking the big question: What did Peter steal? The news reporter is about to write a story about Mary. He is wondering if

Mary now knows what Peter stole. He writes: “Mary now knows that Peter stole the rubies”

Please tell us whether you agree or disagree with the news reporter’s claim, “Mary now knows that Peter stole the rubies.”

The results here are interesting. Subjects tend to agree with the knowledge assertion in the first condition but disagree with it in the second condition. Since the assertions are made up of the exact same words and the facts of the case including Mary’s state of mind and evidence appear the same across the two scenarios, then the best way of explaining the result is to say the assertions express different contents. And this suggests contextualism. Moreover, the difference between contexts does not seem to be one that corresponds to different epistemic standards at play understood in the traditional sense. Rather, the difference seems to correspond to the different contrast cases at play. Hence, this sort of experiment tends to support contrastivism over traditional contextualism.

I now give a critique of this study which is similar to one given by De Rose (forthcoming). Let us look carefully at the vignettes. Like before, it is crucial that the vignettes are set up so that (the perceived) amount of evidence Mary has for the proposition that Peter stole the rubies is the same across the vignettes. This is presumably established by providing subjects nearly identical vignettes where the parts that discuss the evidence available to Mary are

identical. These parts imply that Mary had plenty of evidence that Peter was the robber. But it was not said that Mary had no evidence that it was rubies that were stolen. Instead, we are told that Mary “has no further information”. Now, it is fairly well accepted that quantified claims such as this one have variable interpretations depending on the context in question. For example, we could imagine a context in which the truth of an utterance of “Mary has no further information” is compatible with Mary having plenty of further information which is irrelevant to purposes of the conversation. So suppose, for example, that there was a discussion about who the robber was. Mary could divulge everything she knew about the robber while leaving out information that it was rubies that were stolen. Yet, in that situation one could still truthfully utter “Mary has no further information”. So my point here is simple: the quantified sentence “Mary has no further information” could express a truth in a context in which Mary has further information to the effect that the rubies were stolen. The crucial issue for S&K is that the context of the vignette must not be such a context. The context of the vignette must be one in which “Mary has no further information” rules out the possibility that she has information to the effect that rubies were stolen. I believe, however, that in the first condition (thief-contrast), the sentence is naturally interpreted as being such that its truth (in the context) is consistent with Mary having further information about what was stolen. But this is not the case in the second condition. Let us turn to this.

In the “thief-contrast” scenario, right after we are told that Mary “has no further information” (which is supposed to imply she has no idea what was stolen) we are also told this: “everyone is now asking the big question: who stole the rubies”. But what should a reader make of this? If everyone is asking “who stole the rubies” then it must be widely assumed that it was rubies that were stolen.<sup>36</sup> We should assume that Mary has information that it was rubies that were stolen. So then when we were just told that Mary has “no further information” we should interpret this quantified statement to be compatible with Mary having further information to the effect that the rubies were stolen. In sharp contrast, focusing on the second vignette (jewel-contrast) there is no similar motivation to think that Mary has this information. What follows from this then is that the evidence available to Mary may be different across the vignettes. In the first vignette, Mary has information that rubies were stolen while in the second she does not. No wonder then that subjects think that Mary knows that Peter stole the rubies in the first vignette but not the second vignette. This interpretation of the responses, moreover, is compatible with invariance and traditional contextualism. There is no need to appeal to contrastivism.<sup>37</sup>

---

<sup>36</sup> Arguably, the construction “who stole the rubies” presupposes that rubies were stolen. If we are told that everyone is asking that question, then it is plausible that everyone, including Mary, is aware that (or at least assuming that) it was rubies that were stolen. So it is natural to conclude that Mary has further information (in this condition) about the rubies.

<sup>37</sup> After this manuscript was nearly completed, Jonathan Schaffer reported to me that he conducted some new experiments which aimed to get around the possible confound mentioned here. His results disfavor the idea I am pressing here.

## 5. Conclusion

I have surveyed some recent work in experimental philosophy. This work counts as among the very first in the growing field of ‘experimental epistemology’. The perspective taken here, though friendly to the general approach, has been somewhat critical of particular studies insofar as one may wish to say they are very problematic for IRI or contextualism. One advantage of experimental work, however, is that criticisms can often be tested empirically. In this sense, armchair criticisms will often be incomplete yet still useful in guiding further research.

## WORKS CITED

Bealer, G. 1999: Intuition and the autonomy of philosophy. In M. DePaul and W. Ramsey (eds) *Rethinking Intuition: The Psychology of Intuition and Its Role In Philosophical Inquiry*. Lanham, Maryland: Rowman and Littlefield Publishers Inc.

Beebe, James and Wesley Buckwalter, (forthcoming) *The Epistemic Side-Effect Effect*. *Mind & Language*.

Bengson, J., Moffett, M., & Wright, J.C. (2009). "The Folk on Knowing How." *Philosophical Studies*, 142(3): 387-4

Brown, J (forthcoming) “Experimental Philosophy, Contextualism and SSI”. *Philosophy and Phenomenological Research*.

Buckwalter, W. (forthcoming) “Knowledge isn’t Closed on Saturdays: A Study in Ordinary Language’ *Review of Philosophy and Psychology*.

Cullen, Simon (2010). *Survey-Driven Romanticism*. *Review of Philosophy and Psychology* 1(2)

DeRose. K. 1992. “Contextualism and Knowledge Attributions” *Philosophy and Phenomenological Research* 52(4): 913-929

DeRose, K. 2006. “ ‘Bamboozled by Our Own Words’: Semantic Blindness and Some Arguments Against Contextualism”. *Philosophy and Phenomenological Research*, 73(2): 316-338.

DeRose, K. (ms) Contextualism, Contrastivism and X-phi Surveys.

Fantl, J. & M. McGrath 2007. “On Pragmatic Encroachment in Epistemology” *Philosophy and Phenomenological Research*. 75: 558-89

Feltz, A. & Zarpentine C. (forthcoming) “Do You Know More When it Matters Less”. *Philosophical Psychology*.

Goldman, A. 2007. ‘Philosophical Intuitions: Their Target, Their Source and Their Epistemic Status’ *Grazer Philosophische Studien* 74, 1–26.

Hawthorne, J. 2004. *Knowledge and Lotteries*. OUP

Hawthorne, J. and J. Stanley. 2008 ‘Knowledge and Action’. *Journal of Philosophy*.

Kauppinen, A. 2007 The rise and fall of experimental philosophy. *Philosophical Explorations*, 10, 95-118.

Knobe, Joshua (forthcoming). Person as Scientist, Person as Moralist. *Behavior and Brain Sciences*.

Kornblith, H. 2007. Naturalism and Intuitions. *Grazer Philosophische Studien* 74, 27-49.

Lewis, David. Scorekeeping in a Language Game. *Journal of Philosophical Logic* (8) 339-359.

Ludlow, P. 2005 “Contextualism and The New Linguistic Turn in Epistemology” in Preyer and Peter (eds) *Contextualism in Philosophy: Knowledge, Meaning and Truth*. 11- 50

May, J, W. Sinnott-Armstrong, Jay. G. Hull, Aaron Zimmerman (forthcoming) *Review of Philosophy and Psychology*.

Nagel, Jennifer (2007). Epistemic Intuitions. *Philosophy Compass* 2:6

Nagel, J. 2008. “Knowledge Ascriptions and The Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy*. 86: 279-94.

Nagel, J. 2010. "Knowledge Ascriptions and The Psychological Consequences of Thinking About Error". *Philosophical Quarterly* 60

Neta, Ram and Mark Phelan (ms) "Evidence that Stakes Don't Matter for Evidence."

Nichols, Shaun, Stephen Stich and Jonathan M. Weinberg, (2003). *Metaskepticism: Meditations in Ethno-Epistemology* In S. Luper (ed.), *The Sceptics* (Ashgate), pp. 227-2

Pinillos, Ángel 2011. Time Dilation, Context and Relative Truth. *Philosophy and Phenomenological Research*

Pinillos, Ángel (forthcoming) 'Knowledge, Experiments and Practical Interests. New Essays on Knowledge Ascriptions. Eds Jessica Brown and Mikkel Gerken. Oxford University Press.

Rysiew, P. (2007). "Speaking of Knowing" *Noûs*. 41:4

Schaffer (2006). "The Irrelevance of the Subject: Against Subject Sensitive Invariantism" *Philosophical Studies* 127, 87-107

Schaffer, J and Joshua Knobe (2010). Contrastivism Surveyed. *Noûs*

Swain, S., Alexander, J., and Weinberg J. 2008: The instability of philosophical intuitions: running hot and cold on truetemp". *Philosophy and Phenomenological Research*, 76, 138 – 155.

Stanley, J. 2005. *Knowledge and Practical Interests*. OUP

Weathersom, B. (MSa) 'Defending Interest, Relative Invariantism'

Weatherson, B. (MSb) 'Knowledge Bets and Interests'

Weinberg, J., Nichols, S., and Stich, S. 2001: Normativity and epistemic intuitions. *Philosophical Topics*, 29, 429–460.

Zimmerman, Dean. 2007. *The Privileged Present: Defending an 'A-Theory' of Time*. In *Contemporary Debates in Metaphysics* : Blackwell. Eds Ted Sider, John Hawthorne and Dean Zimmerman.

