Episteme, 9, 1 (2012) 3–26 © Cambridge University Press
doi:10.1017/epi.2011.2

EMPIRICAL TESTS OF INTEREST-RELATIVE INVARIANTISM

CHANDRA SEKHAR SRIPADA AND JASON STANLEY
sripada@umich.edu

ABSTRACT

According to Interest-Relative Invariantism, whether an agent knows that *p*, or possesses other sorts of epistemic properties or relations, is in part determined by the practical costs of being wrong about *p*. Recent studies in experimental philosophy have tested the claims of IRI. After critically discussing prior studies, we present the results of our own experiments that provide strong support for IRI. We discuss our results in light of complementary findings by other theorists, and address the challenge posed by a leading intellectualist alternative to our view.

Suppose Hugo lives in rural Montana, and he has to decide whether or not to go to the market – the only market anywhere near where he lives – in order to get supplies for his baby. He could go today, or he could put it off until Saturday. In one hypothetical situation, it's not a very weighty decision. His wife is home for the weekend and is breastfeeding. They have enough diapers and formula to last until Monday. In another situation, it is a very weighty decision. His wife is away for the weekend, and they have no food for the baby – the milk is spoiled, and they have no formula. Does whether Hugo *knows* that the store will be open this weekend depend upon whether he finds himself in one situation or the other? What about what Hugo has *evidence* to believe, is *justified* in believing, or just *ought* to believe?

According to the thesis of *Interest Relative Invariantism about Knowledge*, how much is at stake in a situation does potentially have a direct impact on whether a participant in the situation knows something at the time at which the situation occurs. One could adopt Interest Relative Invariantism about epistemic notions other than knowledge as well. One could hold that how much is at stake in a situation potentially has a direct impact on whether something counts as evidence at all to participants in a situation. Alternatively, one could hold that how much is at stake in a situation potentially has a direct impact on the *quality of evidence* some information provides to an agent in that situation. Either of these is a version of Interest Relative Invariantism about *Evidence*. Finally, one could be an Interest Relative Invariantist about what one ought to believe, or more generally, how *confident* one ought to be. An Interest Relative Invariantist about the norms of confidence would hold that how confident one ought to be (including whether one ought to adopt a full belief) about a certain proposition depends in part on how much is at stake.

Some of these versions of Interest Relative Invariantism (henceforth IRI) might sound like non-starters. One of the morals of Pascal's Wager is, after all, that one should not confuse what one epistemically ought to believe with what one ought to believe practically.

In perhaps other terms, what makes something a good practical reason is not what makes something a good epistemic reason, and vice versa. IRI about what one ought to believe sounds like it violates this moral of Pascal's Wager. However, this is far too quick. The Interest Relative Invariantist about norms of belief or confidence holds that there is a practical element even in epistemic reason. This does not mean that there is no distinction between practical and epistemic reason, and between practical and epistemic 'ought'.[1]

Fantl and McGrath (2002, 2009), Hawthorne (2004), and Stanley (2005) argue for IRI about various epistemic properties and relations. They exploit several different kinds of arguments for Interest Relative Invariantist theses. Most of the arguments for IRI about knowledge (at least) are fairly theoretical, and enter in as defensive maneuvers against skeptical threats, or threats to various theoretical principles about knowledge. For example, Fantl and McGrath (2009) hold that IRI about Knowledge rescues fallibilism about knowledge from the objection from Concessive Knowledge Attributions, statements of the form 'I know that p, though it is possible that q', where q is some proposition obviously inconsistent with p. Hawthorne (2004) argues that it allows for a non-skeptical response to various versions of the problem of skepticism formulated with Lottery proposition type scenarios. Stanley (2004), Hawthorne (2004), and Hawthorne and Stanley (2008: 588ff.) argue that the Interest Relativity of Knowledge is a corollary of principles connecting knowledge and action. More generally, the thought is that there is some connection between properties from the domain of practical rationality and epistemic properties and relations, which entails that the obtaining of the latter has a practical dimension.

Defenders of IRI about an epistemic property or relation concede that the thesis is surprising. But they nevertheless seem to think IRI about Knowledge, or the norms for belief or evidence is supported by, or at the very least *consistent* with, ordinary judgments speakers have about hypothetical cases. For example, defenders of IRI seem to think that speakers would be more reluctant to grant that Hugo knows that a particular store has diapers in the hypothetical high-stakes situation described above than in the hypothetical low-stakes situation (Stanley, 2005).

However, a number of recent experimental philosophy studies challenge the view that IRI about any epistemic property or relation is even consistent with ordinary judgments speakers have. Summarizing these studies in a forthcoming paper, Jonathan Schaffer and Joshua Knobe write:

> a recent series of empirical studies threatens to undermine this whole debate. … Strikingly, the results suggest that people simply do not have the intuitions they were purported to have [by advocates of IRI].

In fact, Schaffer and Knobe go further than even many of the authors of the experimental philosophy studies in drawing a conclusion about IRI. As they write about IRI models of knowledge:

> These models predict that the higher the stakes for the subject, the lower the propensity for the ascriber to ascribe knowledge (all else equal). Such models (unsupplemented by pragmatic or performance based factors) seem to predict *the exact opposite* of the data.

1 These points are ably made in detail in Schroeder (forthcoming).

The hypothetical cases discussed by advocates of IRI are not novel with them. In fact, the main cases discussed by Stanley (2005) are due originally to *contextualists* in epistemology, such as Stewart Cohen and Keith DeRose. Contextualists about knowledge ascriptions hold that predicates of the form 'knows that *p*' are context-sensitive in a distinctively epistemological way. The novel contribution of advocates of IRI is to point out an alternative explanation of these sorts of intuitions. If Schaffer and Knobe are right that the folk lack the intuitions about hypothetical cases that many epistemologists have thought they had, they will have made an important discovery.

Our purpose in this paper is to provide our own assessment of the 'intuitive' case for IRI about a range of epistemic properties and relations. We review previous studies and highlight certain confounding variables that may have been present in these studies. We then present our own studies that strongly support the position that folk have stakes-sensitive intuitions in the way predicted by IRI. We also use our data to shed novel light on the nature of the connections between notions in the domain of practical rationality and epistemic properties and relations. Thus, we seek not merely to show that the folk do have stakes-sensitive intuitions, but to use the methodology of experimental philosophy to illuminate the theoretical principles that underlie these intuitions.

We begin with a brief review of the studies that form the basis for Schaffer and Knobe's conclusion. Our review is organized around four methodological issues that arise in these studies. This review sets the stage for the presentation of our own studies in section 2, where we attempt to address these methodological concerns. In section 3, we turn to recent results by Angel Pinillos that nicely complement our own. In section 4, we raise both theoretical and empirical concerns for Jennifer Nagel's alternative 'strict invariantist' account of the data. Finally, in section 5, we show how our results provide evidence for novel connections between practical and theoretical rationality that support IRI.

1. CONCERNS ABOUT PRIOR EXPERIMENTAL STUDIES

1.1. Assessing knowledge claims in embedded contexts

Previous experimental studies frequently tested the stakes-sensitivity hypothesis by assessing subjects' levels of agreement with statements of the following general form (see e.g. Buckwalter 2010; Feltz and Zarpentine forthcoming):

> When S says, 'Hannah knows that the bank will be open on Saturday,' what she says is true.

Here the knowledge claim is embedded into a relative clause uttered by a third party. Subjects are asked to assess whether or not someone's sincere assertion is true, where the assertion is a knowledge-attribution. The choice to utilize embedded knowledge claims reflects the way so-called 'bank cases' were initially presented in the work of advocates of contextualism about knowledge ascriptions, such as Keith DeRose, and represented in Stanley (2005). However, embedding raises certain interpretational difficulties.

It is likely that subjects in experimental studies will tend not to directly contradict others' assertions. This may be due to non-epistemic factors, such as social graces and wariness to offend or anger. It may also reflect epistemic factors, such as a principle of

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

charity that holds that others' sincere assertion that *p* provides strong evidence that *p*. These factors would tend to dampen the effect of stakes on knowledge attributions. This is because subjects' hypothesized tendency to deny knowledge claims in high stakes conditions would be countervailed by their tendency to avoid contradicting the knowledge claims asserted by the characters in the vignette.[2]

There is a sociological reason why these vignettes have the extra complexity involved in embeddings. This is that the vignettes were originally devised to provide support for contextualism about knowledge ascriptions, the view that predicates such as 'knows that snow is white' are context-sensitive in a distinctively epistemological way. Contextualism is a linguistic thesis about the context-sensitivity of epistemic vocabulary. Embedding is necessary for assessing the hypothesis of contextualism because we need to know whether speakers have the context-sensitive intuitions in question (see DeRose, forthcoming: section 1.2). In assessing contextualism, we are assessing whether the utterance of a certain sentence can have one interpretation in one conversational context, and a different one in another conversational context.

In contrast, embedding is not at all necessary for testing IRI. In fact, embedding is a distracting complication in evaluating IRI. Interest Relative Invariantism is a metaphysical thesis about epistemic properties and relations. That is, IRI is a thesis about epistemic *facts*. In assessing whether IRI has intuitive or counterintuitive consequences, we should only be asking subjects whether their judgments about the epistemic facts vary between circumstances differing only in what is at stake. It only confuses matters to ask subjects about their judgments about assertive utterances of sentences that express those facts.

It is only a sociological accident that embedded knowledge claims occur in these cases. Given the methodological and interpretational difficulties associated with eliciting subjects' assent to embedded knowledge claims, we believe that studies aiming to test interest-relative invariantism (as opposed to contextualism) should instead ask subjects to *directly* assess claims that the relevant character in the vignette has the epistemic property in question.[3]

1.2. *Knowledge and mere truth, or mere true belief*

A number of philosophers have recognized that there are standard contexts in which people will agree that someone knows something, even when it is utterly clear that they only possess a true belief. As John Hawthorne (2000) writes:

2 Psychologists call subjects' tendency towards non-contradiction *acquiescence bias*. We thank Jennifer Nagel for calling our attention to the extensive psychological literature (e.g. Podsakoff et al. 2003) that documents this phenomenon. DeRose (forthcoming: section 1.3) makes a distinct but related point. If contextualism is true, then the principle David Lewis calls 'accommodation' predicts that we will shift the context to one that makes the utterance true. So one should avoid having embedded occurrences of knowledge attributions that the experimental subject is supposed to contradict. We do not of course assume the truth of contextualism, so our point is not about accommodation, but about charity.

3 Jennifer Nagel raises the worry that participants in these prior experimental philosophy studies weren't evaluating the matrix predicate at all, but rather were looking at the gist and judging the truth value of the embedded proposition. That is, they were taking the phrase 'what Hannah said' to refer to the embedded proposition, rather than to the mental state attribution, and evaluating the truth of the former.

> Suppose I ask in an ordinary setting whether someone knows whether Boston is the capital of Massachusetts. Suppose it turns out that he does truly believe this though the epistemic credentials of his path to that belief are decidedly shaky: Perhaps he got it from a book that misprinted most of the state capitals though not this one. Perhaps he got the information from someone that he had good reason to distrust (who happened to be sincere on this occasion or else who tried to lie and accidentally told the truth on this occasion). Would your acceptance of the statement 'Boston is the capital of Massachusetts' fail to be knowledge in such cases? Not so, or not clearly so.

With other epistemologists, we do not take these facts about folk judgments to show that knowledge is *merely* truth, or that knowledge is merely true belief. Though the folk sometimes focus on truth or true belief, knowledge is epistemically richer. One hypothesis as to why the folk are so easily led to mistake truth or true belief for knowledge is that the true belief aspects of knowledge are often more salient and overshadow the justificatory or basing dimension. In many contexts, all we care about is that an agent had a true belief, or that the content of the knowledge is true. In such contexts, we don't care about the justificatory or basing requirements of knowledge – we care only about its *factivity*, i.e. that knowledge entails truth. So we are led to grant that someone knows something even though they clearly do not satisfy anything beyond the truth or true belief requirements of knowledge.

An early experimental study of people's knowledge attributions by Jonathan Weinberg, Shaun Nichols, and Stephen Stich – perhaps because they were cognizant of this very issue – presented subjects with hypothetical cases and asked them whether the character in the vignette 'Really knows' that p or whether he or she 'only believes' that p (Weinberg et al., 2001). Explicitly contrasting knowledge and mere belief may help subjects focus on more than just the belief aspect of knowledge. But there is still a worry that people will focus only on the factivity of knowledge in assessing whether or not someone knows. IRI about Knowledge, it bears emphasis, is a thesis about just the epistemic aspects of knowledge. We need a way to filter out the proclivity to focus on the non-epistemic aspects of knowledge. One way to do so is to ask about the epistemic aspects of knowledge, such as quality of evidence and warranted confidence, e.g. how confident one ought epistemically to be.[4] One can probe judgments about a person's quality of evidence for p without having to stipulate that p is true, and the same applies to judgments about warranted confidence. Probing these other judgments thus provides a way to circumvent the people's tendency to focus on the factivity of knowledge.

In the previous studies of IRI that concerned knowledge, no specific efforts appear to have been made to ensure that subjects were focused on the epistemic aspects of knowledge, rather than just the true belief requirement. This is noteworthy because if subjects are just focused on the true belief requirement of knowledge, this would tend to produce a pattern of results in which one observes high knowledge ratings for both high and low stakes vignettes, and therefore attenuation of the difference in ratings between high and low stakes vignettes, since stakes are not predicted to have much effect on truth or even mere true belief. This pattern of results was exactly what was found, raising the worry that it was only true belief, or even mere truth, that was being probed. Because of this

4 There are interesting questions arising for the advocate of knowledge-first epistemology from potential divergences between folk intuitions about knowledge and folk intuitions about evidence.

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

complication, a thorough study even of IRI about knowledge should elicit subjects' judgments about other epistemic properties that are constitutively connected to knowledge.

1.3. Conveying that the stakes are low

Philosophical readers have two important advantages over experimental subjects in understanding the key features of IRI-relevant thought experiments (such as DeRose's bank cases as well as other cases). First, philosophical readers are privy, prior to reading the cases, to the specific hypotheses that are being tested. This greatly aids their abilities to pick out the salient features of the cases and avoid getting hung up on irrelevant details. Second, philosophical readers are presented with both the target case as well as the matched contrast case. Having access to both cases greatly facilitates one's ability to mentally move back and forth between the cases and 'home in on' the intuitive judgments that are supposed to differ across the two cases.

Because experimental subjects are deprived of these important methods for understanding the cases, it raises a key question. Even when the cases explicitly state the stakes at issue, do subjects apprehend these stakes in a sufficiently clear way at the time they make their epistemic judgments? For example, consider the various versions of DeRose's bank cases tested in a number of studies including May et al. (2010), Buckwalter (2010), and Feltz and Zarpentine (forthcoming).

> *Standard Low Stakes Bank Case*
> 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. As they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Hannah notes that she was at the bank 2 weeks before on a Saturday morning, and it was open. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, 'I know the bank will be open tomorrow. So we can deposit our paychecks tomorrow morning'.

While the case says 'it isn't very important that their paychecks are deposited right away', does this mean subjects will infer that the stakes for the Hannah are low? Perhaps not. A person who wants to deposit a check and falsely believes that the bank is open tomorrow will make a trip all the way to the bank only to find it closed, and will thus have to return a second time. This would seem to be a fairly bad outcome. Moreover, the subjects' impression that the additional trip to the bank is costly is heightened by the fact that Hannah and Sarah are actively deliberating about it. Finally, subjects do not get to see the matched contrast case in which the stakes are obviously much higher. Thus they cannot use the High Stakes Bank case to help them infer that the stakes in Low Stakes Bank are supposed to be interpreted as being quite low.

In short, while the Low Stakes Bank case conveys the lowness of the stakes to philosophical readers, it is likely that experimental subjects will often fail to interpret the case in this way. Another example of this same problem occurs in Felz and Zarpentine (forthcoming) who present the following Low Stakes Ravine case:

> Bill, Jim, and Sarah are hiking and they come to a ravine. There is a bridge five feet over the ravine. Bill sees Jim and Sarah cross the bridge and Jim says to Sarah 'Bill knows that the bridge is stable enough to hold his weight.'

Subjects are then asked whether what Jim says is true. In the paired contrast case, the ravine is 1,000 feet deep. We believe that an experimental subject who is unaware of the hypothesis being tested and who reads just the preceding Low Stakes case (without reading the matched High Stakes case) will likely perceive the stakes involved are quite high. First, a fall of five feet will surely be painful, if not seriously damaging, or even fatal.[5] Secondly, unless the stakes are indeed high in the situation, it's very odd that Jim bothers not only to reflect upon whether Bill knows that the bridge is stable, but also to communicate that information to Sarah. Once again, these cases are perhaps effective for illustrating the stakes effect to philosophers who know the hypothesis being tested and get to see both cases. But it is unlikely that they effectively convey the stakes to the folk.

When redesigning Low Stakes vignettes to correct for the preceding problem, it is easy to make another kind of mistake. Consider the following vignette:

> When Hannah drinks cow's milk, she gets a very slight headache. The headache is really very mild and it does not bother Hannah very much. The headache does not interfere with Hannah's other activities. Because the headache she gets from cow's milk is so mild, it is not very important whether what is in the glass turns out to be cow's milk or soymilk.

An experimentalist might use a vignette like this to try to convey that the stakes for Hannah are low. The problem, however, is that the narrator spends considerable time describing the putatively low stakes consequence in careful and exacting detail. Given that subjects plausibly expect narrators to spend time describing things that are important, and ignore things that are genuinely unimportant, the vignette above ends up sending strong cues to subjects that the consequences of drinking cow's milk are indeed in some way important. Why else would the narrator spend so much time describing this consequence if the consequence is *not* important or significant to the character in the vignette in some way? Call this 'the problem of narrator cues'. It is highly non-trivial to design a Low Stakes vignette that avoids this problem.

1.4. *High stakes subjects think and behave differently*

Suppose Hi and Lo are given a day to ascertain the truth of some proposition. For Hi, everything is riding on being right, while for Lo, being right matters hardly at all. And let's further suppose that both Hi and Lo are perfectly aware of the stakes that are involved. Would you predict that Hi and Lo would differ in the ways they gather data, the reasoning processes they employ, the exhaustiveness of their search for evidence, and so on? It seems perfectly obvious that they would differ in all these respects.

5 For example, in the summer of 2011, Sammy Wanjiru, the great Kenyan marathoner, suffered a fatal blow to the head after falling from a first floor balcony.

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

Moreover, these differences are directly relevant to the truth conduciveness of their respective inquiries. Thus the following seems to be plausible:

(1) Other things being equal, the magnitude of the stakes for Ss being right positively correlates with the assumed truth conduciveness of S's inquiry.

This raises a problem. IRI itself predicts that, other things being equal, higher stakes tend to *diminish* attributions of knowledge. But (1) predicts an effect in the opposite direction. That is, since a person in a high stakes situation is assumed to have engaged in a more truth conducive pattern of inquiry, subjects will be *more* likely to attribute knowledge to this person.

The overall situation is shown in Figure 1. Stakes are predicted to produce two divergent effects on the epistemic facts. If our judgments support the predictions of IRI, then the studies should show that higher stakes diminish our proclivity to think that subjects have the relevant epistemic properties (path a). But higher stakes are also predicted to increase our confidence that the characters in the vignette have gathered more evidence. So stakes are relevant via an *indirect* pathway (paths b and c). On this pathway, higher stakes lead to the perception that the person has engaged in a more exhaustive and thus more truth conducive inquiry (path b), which in turn leads us to think that they are more likely to have the desirable epistemic properties (path c).

Variable 2 in Figure 1 is a *suppressor variable* (Maasen and Bakker 2001) because it attenuates the effect of interest, i.e. the direct relationship between stakes and judgments about the epistemic facts. To illustrate suppression, consider the relationship between the size of a fire and the time it takes to put it out. One might plausibly hypothesize that, other things being equal, larger fires take more time to put out than smaller fires. But if larger fires also lead to more fire trucks being dispatched, then the difference in putting-out-time between larger and smaller fires will be attenuated, or may even be reversed. The presence of a suppressor variable does not imply any problem with the *truth* of the original hypothesis. It might very well be true (and presumably is true!) that, other things being equal, larger fires take longer to put out than smaller fires. The presence of suppression, rather,

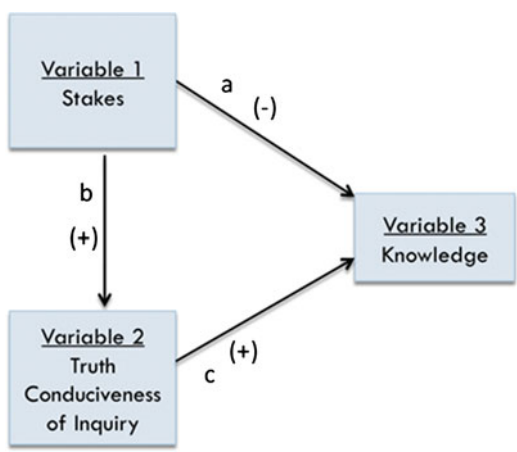

Fig. 1. (Colour Online) Stakes exert two ***opposing*** effects on knowledge attributions.

introduces problems for *detecting* the hypothesized relationship via experimental studies. Such studies must remove the suppressor or else measure it and control for its effect.

Previous experimental studies of IRI appear not to have taken note of the potentially important suppressor effect shown in Figure 1, or sought to remove or control for it (the exception is Angel Pinillos's studies, which we discuss in section 4). We hypothesized that this may in part explain why many of these studies found only small differences in judgments between the high and low stakes conditions.

2. EXPERIMENTS

Subjects and procedures

All extant studies that tested IRI faced several of the preceding concerns. Therefore, we undertook studies of our own. We investigated a total of three pairs of vignettes, each of which had a low stakes and high stakes version. Thus there were six conditions in total. We used a between-subject design (subjects were assigned to only one condition of the study). Fifty subjects participated in each condition, for a total of 300 subjects for the whole experiment.

We call the first pair of vignettes the 'Basic' vignette, because the vignettes that follow modify features of this vignette in certain ways. The low and high stakes versions of the Basic vignette went as follows.

Basic Low Stakes

Hannah has a gene that causes her to experience a slightly dry mouth when she eats pine nuts. Hannah is very much aware of this, and has known this for a very long time.

One evening, Hannah and her sister Sarah are at a new restaurant that has just opened. Hannah orders a plate of noodles. When her food is brought to the table, Hannah notices something that looks like nuts sprinkled on her noodles and wonders what it is. Sarah says, 'The noodles may be topped with pine nuts.' Hannah notes that the menu says her dish does not contain pine nuts. Based on this, Hannah forms the belief that the noodles are not topped with pine nuts.

If it turns out that the noodles are topped with pine nuts, then when Hannah eats the dish, her mouth will get a little dry. Since Hannah has plenty to drink with her meal, it does not matter very much whether or not the noodles are topped with pine nuts.

Basic High Stakes

Hannah has a gene that makes her seriously allergic to pine nuts. Eating only a single pine nut will cause her to go into shock and die. Hannah is very much aware of this, and has known this for a very long time.

One evening, Hannah and her sister Sarah are at a new restaurant that has just opened. Hannah orders a plate of noodles. When her food is brought to the table, Hannah notices something that looks like pine nuts sprinkled on her noodles and wonders what it is. Sarah says, 'The noodles may be topped with pine nuts.'

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

Hannah notes that the menu says her dish does not contain pine nuts. Based on this, Hannah forms the belief that the noodles are not topped with pine nuts.

If it turns out that the noodles are topped with pine nuts, then when Hannah eats the dish, she will go into shock and die. Since eating even a single pine nuts will cause her to die, it matters a lot whether or not the noodles are topped with pine nuts.

After reading *one* of the preceding cases, subjects next answered *both* of the questions in Table 1.

In designing the Basic vignette, we were attempting to address the concerns raised in the previous section about prior experimental philosophy studies. For example, given complications that arise from probing assent to knowledge claims in embedded contexts, we instead asked subjects to directly assess whether the character in the vignette knows that *p* (Question 2). In response to concerns that subjects are focusing too much on the factivity of knowledge, we asked subjects to assess the character's quality of evidence for *p* (Question 1), in addition to asking subjects whether the character knows that *p* (Question 2).

The Basic Low Stakes case above describes in detail the consequences for Hannah of being wrong about whether her noodles are topped with pine nuts. It thus raises the 'problem of narrator cues' raised in 1.3, i.e. the problem that by spending so much time in describing the consequences, the case implicitly conveys that the consequences are in some way important or significant. To address this concern, we conducted an additional study. In the following 'Implicit/Explicit' vignette, the fact that the stakes are low is implicitly conveyed, rather than explicitly described.

Implicit Low Stakes

Hannah likes the taste of most foods and is not a very picky eater.

One evening, Hannah and her sister Sarah are at a brand new restaurant that has just opened up. Hannah orders a plate of noodles. When her food is brought to the table, Hannah notices something that looks like nuts sprinkled on her noodles and wonders what it is. Sarah says, 'The noodles may be topped with pine nuts.' Hannah notes that the menu says her dish does not contain pine nuts. Based on this, Hannah forms the belief that her noodles are not topped with pine nuts.

The matched high stakes case explicitly describes the high stakes involved.

Table 1. Questions used for the Basic vignette and Implicit/Explicit vignette (vignettes 1 and 2)

Question	Question wording	Anchors for 7-point scale
1	What is the strength of Hannah's evidence that her noodles are not topped with pine nuts?	Very weak evidence, Very strong evidence
2	Suppose it turns out that her noodles are not topped with pine nuts. Please rate how strongly you agree or disagree with the following sentence: 'Hannah *knows* her noodles are not topped with pine nuts.'	Strongly agree, Strongly disagree

Explicit High Stakes

Hannah is seriously allergic to pine nuts. Eating only a single pine nut will cause her to go into shock and die. Hannah is very much aware of this and has known this for a very long time.

One evening, Hannah and her sister Sarah are at a brand new restaurant that has just opened up. Hannah orders a plate of noodles. When her food is brought to the table, Hannah notices something that looks like nuts sprinkled on her noodles and wonders what it is. Sarah says, 'The noodles may be topped with pine nuts.' Hannah notes that the menu says her dish does not contain pine nuts. Based on this, Hannah forms the belief that her noodles are not topped with pine nuts.

After reading *one* of the preceding vignettes, subjects next answered *both* of the questions in Table 1.

Finally, in order to address the concern that high stakes lead subjects to infer that the characters in the vignette engaged in more truth conducive inquiry (i.e. the suppressor hypothesis raised in 1.4), we also tested a pair of cases in which Hannah is ignorant about the stakes involved. In this vignette, Hannah is completely ignorant that she has a gene that makes her allergic to the relevant food item. We thought it is implausible that an adult could really be unaware that she has a life-threating allergy to pine nuts, since pine nuts are so commonly encountered in day-to-day life. Thus we created a new case in which Hannah is allergic to *Mongolian pine nuts*, a fictional kind of nut that we thought subjects would perceive as fairly exotic and rarely encountered.[6]

Ignorant Low Stakes

Hannah has a gene that causes her to experience a slightly dry mouth when she eats Mongolian pine nuts. Hannah has absolutely no idea that she has this gene, nor is there any way she could know that she has this gene.

One evening, Hannah and her sister Sarah are at a brand new Mongolian restaurant that has just opened up. Hannah orders a plate of noodles. When her food is brought to the table, Hannah notices something that looks like nuts sprinkled on her noodles and wonders what it is. Sarah says, 'I heard that Mongolian dishes are often served topped with Mongolian pine nuts.' Hannah notes that the menu says her dish does not contain Mongolian pine nuts. Based on this, Hannah forms the belief that the noodles are not topped with Mongolian pine nuts.

If it turns out that the noodles are topped with Mongolian pine nuts, then when Hannah eats the dish, her mouth will get a little dry. Since Hannah has plenty to drink with her meal, it does not matter very much whether or not the noodles are topped with Mongolian pine nuts.

6 Somewhat annoyingly, in spring 2011, Chinese pine nuts caused a episode in France in which consumers of the nuts suffered a bitter aftertaste for 6–8 weeks. This news item, thankfully, appears to have largely been relegated to European news sources, while the current study was conducted exclusively with subjects who reside in the USA.

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

Ignorant High Stakes

Hannah has a gene that makes her seriously allergic to Mongolian pine nuts. Eating only a single Mongolian pine nut will cause her to go into shock and die. Hannah has absolutely no idea that she has this gene, nor is there any way she could know that she has this gene.

One evening, Hannah and her sister Sarah are at a brand new Mongolian restaurant that has just opened up. Hannah orders a plate of noodles. When her food is brought to the table, Hannah notices something that looks like nuts sprinkled on her noodles and wonders what it is. Sarah says, 'I heard that Mongolian dishes are often served topped with Mongolian pine nuts.' Hannah notes that the menu says her dish does not contain Mongolian pine nuts. Based on this, Hannah forms the belief that the noodles are not topped with Mongolian pine nuts.

If it turns out that the noodles are topped with Mongolian pine nuts, then when Hannah eats the dish, she will go into shock and die. Since eating even a single Mongolian pine nut will cause her to die, it matters a lot whether or not the noodles are topped with Mongolian pine nuts.

After reading *one* of the preceding two vignettes, subjects next answered *both* of the questions in Table 2.

Results

Results for all three pairs of cases are presented in Figure 2. For question 1 which concerns Hannah's quality of evidence, there was a significant effect of stakes in all three vignettes, with subjects saying Hannah's evidence was weaker in the high stakes case compared to the low stakes condition (Basic Vignette: $t(98) = 1.98$, $p = 0.05$; Implicit/Explicit: $t(98) = 2.29$, $p = 0.02$; Ignorant: $t(98) = 4.15$, $p < 0.001$).

For question 2 which concerns judgments about knowledge, there was a highly statistically significant effect of stakes in vignettes 2 and 3 (Implicit/Explicit: $t(98) = 3.43$, $p = 0.001$; Ignorant: $t(98) = 3.61$, $p < 0.001$). However, in vignette 1 (i.e. the Basic vignette), a statistically significant stakes effect was not observed ($t(98) = 0.25$, $p =$ n.s.).

Discussion

Overall, we found a statistically significant effect of stakes on epistemic properties in all three vignettes. In the Basic vignette, we observed stakes effects on judgments of quality

Table 2. Questions used for the Ignorant vignette (vignette 3)

Question	Question wording	Anchors for 7-point scale
1	What is the strength of Hannah's evidence that her noodles are not topped with Mongolian pine nuts?	Very weak evidence, Very strong evidence
2	Suppose it turns out that her noodles are not topped with Mongolian pine nuts. Please rate how strongly you agree or disagree with the following sentence: 'Hannah *knows* her noodles are not topped with Mongolian pine nuts.'	Strongly agree, Strongly disagree

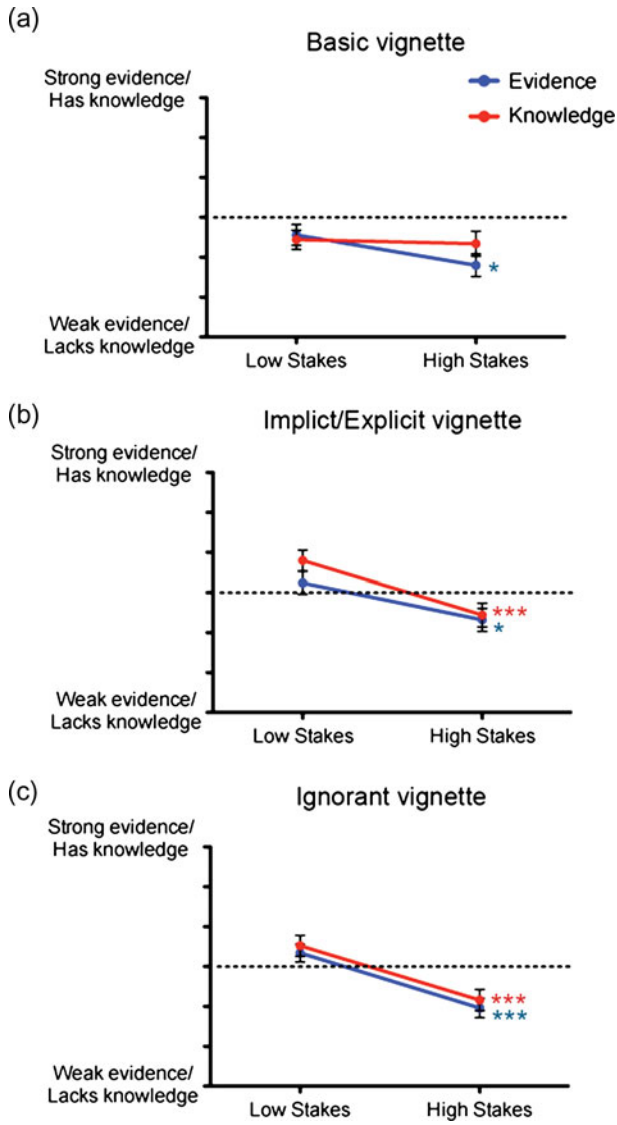

Fig. 2. (Colour Online) Results for questions about quality of evidence (blue) and knowledge (red) for the low stakes and high stakes versions of all three vignettes. * = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$.

of evidence, but not on judgments of knowledge. One explanation for this difference is the hypothesized tendency for the folk to focus on the factivity of knowledge (see section 1.2). The fact that we saw effects of stakes on quality of evidence in all three vignettes provides tentative support for the view that questions about quality of evidence may be more sensitive as probes of stakes effects compared to questions about knowledge.

As predicted, we observed a more robust effect of stakes on both quality of evidence and knowledge in the Implicit/Explicit vignette than the Basic vignette. Moreover, the larger stakes effect was driven by higher ratings for quality of evidence and knowledge

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

in the Implicit Low Stakes case. That is, while ratings in the Explicit High Stakes case did not differ from ratings in the Basic High Stakes case ($t(98) = 0.23$, p = n.s.), ratings in the Implicit Low Stakes case were significantly higher than in the Basic Low Stakes case ($t(98) = 3.87$, $p < 0.001$). This result supports the view that the problem of narrator cues (see section 1.3) is indeed a problem. When low stakes are made explicit and described in detail (i.e. as in the Basic Low Stakes case), people's ratings for quality of evidence and knowledge are very similar to their ratings in a high stakes case. But when low stakes are left implicit and not explicitly described in detail, ratings for quality of evidence and knowledge are significantly higher.

We also predicted that we would observe larger stakes effects in the Ignorant vignette relative to the Basic vignette. Results confirmed this prediction, though not in the way we expected. We predicted that the larger stakes effect in the Ignorant vignette would be driven by *lower* ratings for quality of evidence and knowledge in Ignorant High Stakes versus Basic High Stakes. This prediction is rooted in the suppressor hypothesis discussed in section 1.4. That is, we hypothesized that subjects would assume that a non-ignorant high stakes character has performed more truth conducive forms of inquiry. Such an assumption would be blocked for a character ignorant of the high stakes involved, thus leading to diminished quality of evidence and knowledge ratings. We did *not* find the predicted difference in the Ignorant High Stakes case versus Basic High Stakes case ($t(98) = 0.43$, p = n.s.), providing evidence against the suppressor hypothesis.

Instead, we found that in Ignorant Low Stakes, subjects rated quality of evidence and knowledge to be significantly *higher* than in the Basic Low Stakes ($t(98) = 3.04$, $p < 0.01$). One explanation for this result borrows from the logic of the problem of narrator cues discussed earlier. Recall that explicit and detailed discussion by the narrator of the consequences involved in a low stakes situation serves to send the cue that the stakes are actually important or significant in some way. It may be that there is also a second source of confounding cues. In the Basic Low Stakes vignette, the case reads 'Hannah has a gene that causes her to experience a slightly dry mouth when she eats pine nuts. Hannah is very much aware of this, and has known this for a very long time.' It may be that the fact that a character is 'very much aware' of some fact, and 'has been aware of this for a long time', sends the cue that the fact at issue is in some way important or significant. After all, why would Hannah be so very aware of something of little significance? If so, then IRI predicts that Ignorant Low Stakes, which removes this unwanted cue, would yield significantly higher ratings for quality of evidence and knowledge, which is exactly what we found. In sum, the Basic Low Stakes case may give rise to two kinds of confounding cues. Narrator cues are removed in Implicit Low Stakes, and character cues are removed in Ignorant Low Stakes. This would explain why quality of evidence and knowledge ratings are significantly higher in the latter two cases than in the Basic Low Stakes case.

Overall, the sizes of the stakes effects observed in this study were modest – typically around one point on a seven-point scale. But this should not be surprising. IRI claims that that knowledge and other epistemic properties are *sensitive* to stakes, not that they are entirely dependent on stakes and stakes alone. IRI theorists agree with traditional epistemologists that knowledge and other epistemic properties are sensitive to factors other than stakes (such as reliability). It is not as if advocates of IRI hold that manipulating the stakes can magically transform epistemic dirt into gold, and vice versa. Seen in this light, the modest-sized effects observed in these experiments are not a problem for IRI, and indeed are what should be expected given sensible versions of IRI.

3. TESTING IRI WITH ALTERNATIVE EXPERIMENTAL DESIGNS

Nearly all studies of IRI used a similar experimental design. This design holds fixed (or at least tries to hold fixed) the evidence available to the characters, and measures subjects' judgments about whether the characters know some proposition, or possess some other epistemic properties or relations. This allows one to test, given that a character performs some fixed amount of evidence gathering, whether knowledge shifts depending on stakes. An alternative way that one might test IRI essentially flips the standard design on its head. One could allow characters in the vignette to collect less versus more evidence, and ask how much evidence they need to collect in order to know some proposition. This allows one to test, given that a character knows some proposition, whether the amount of evidence required to qualify the person as having knowledge shifts depending on stakes.

Call these two kinds of designs 'evidence-fixed' and 'evidence-seeking' designs respectively. Angel Pinillos has developed a series of elegant experiments that use evidence-seeking designs to test IRI (Pinillos, manuscript). In one experiment, subjects are presented with one of the following two cases.

> *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, 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 the vignette, all subjects were asked 'How many times do you think Peter [John] has to proofread his paper before he knows that there are no typos? ____ times.' Subjects were told to insert an appropriate number in the blank space. Results showed a highly significant difference between the two cases. In particular, subjects gave higher answers in the high stakes condition (median = 5) than in the low stakes condition (median = 2).

Pinillos ran a number of other studies using evidence-seeking designs to systematically address alternative explanations of the results. For example, one might raise an objection that an additional difference between Low Typo and High Typo is that, in High Typo,

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

John might be more reluctant to form a *belief* that there are no typos in his paper. Since knowledge requires belief, then this might explain why high-stakes John must perform more proofreads than low-stakes Peter before he is credited with knowing that there are no typos. To address this 'belief objection', Pinillos ran another study with the following additional paragraph at the end of each version of the vignette:

> *Typo-Explicit Belief*
> It turns out that right after Peter [John] finished writing his paper, he formed the belief that there are no typos in his paper—and in fact there are no typos. But does he know this? How many times do you think that Peter [John] has to proofread his paper before he knows there are no typos? ___ times.

Results showed that stipulating belief in this way did not erase the asymmetry – answers were still significantly higher in the high stakes condition compared to the low stakes condition. As an additional check against the belief objection, Pinillos ran another pair of vignettes in which the character is completely ignorant of the stakes involved in the situation. Since the character is ignorant, it cannot be that perception of higher stakes lead to reluctance on the part of the character to form the relevant belief about the paper's not containing any typos. Once again, results showed a robust asymmetry between high and low stakes, providing additional evidence that the belief objection is off the mark.

Notice that Pinillos's results mirror and complement our own. We studied a number of vignettes with evidence-fixed designs, and found diminished ratings for knowledge and quality of evidence in vignettes involving high stakes compared to otherwise matched vignettes involving low stakes. Pinillos studied a number of vignettes with evidence-seeking designs, and found that in order for a person to know some proposition, the person has to collect more evidence in high stakes situations compared to otherwise matched low stakes situations. These convergent results provide strong support for the IRI thesis that knowledge (and other epistemic properties such as quality of evidence) are indeed sensitive to stakes.

4. JENNIFER NAGEL'S VIEW

The challenge from Schaffer and Knobe (forthcoming) was to show that we do have stakes-sensitive intuitions about epistemic notions. We believe our experimental data, coupled with complementary findings from Angel Pinillos, address this challenge. But there is another challenge to IRI, one that is consistent with the presence of the stakes-sensitive intuitions we have shown to be present in ordinary judgments. This approach grants that the folk have stakes-sensitive intuitions but denies that these intuitions are properly interpreted as evidence for IRI. The challenge is best articulated in the work of Jennifer Nagel (2011, 2008). As Nagel writes:

> Using complex cases which readily produce variations in expected epistemic behavior, contextualists and relativists have diagnosed these variations in our expectations as arising from shifting epistemic standards … I argue that our varied expectations for epistemic behavior might be better seen as arising from an invariant expectation that subjects will think adaptively; in fact, I'll aim to

> show that we naturally see anxious high-stakes and casual low-stakes subjects as needing to engage in different behavior in order to meet precisely the same standard. (2011: 3)

Nagel proposes a psychological force called *epistemic anxiety* that serves to regulate the cognitive effort directed to ascertaining the truth of practically relevant propositions in order to bring the amount of effort expended into line with the practical costs of being wrong about the relevant propositions. High stakes situations generate elevated levels of epistemic anxiety. As a result, people in high stakes situations tend to utilize evidence-gathering strategies that are more thorough and accurate than the strategies used in low stakes situations. This claim is perhaps unsurprising, but for added measure Nagel provides additional evidence derived from a substantial body of psychological research. She sums up the evidence as follows:

> In general, high-stakes subjects think more systematically and less heuristically, relying more on deliberate and controlled cognition and less on first impressions and automatic responses [Kunda 1990; Lerner and Tetlock 1999]. Many cognitive biases—a recent survey article on accountability counts sixteen—are known to be attenuated when subjects take themselves to be shifted into a higher-stakes condition [Lerner and Tetlock 1999]. (2008: 282)

In addition to deploying more systematic, thorough and accurate strategies, high stakes subjects also exhibit a lowered *need-for-closure*. Here 'closure' is technical term introduced by the psychologist Arie Krugalanski to refer to the state immediately after the transition from the state of entertaining a hypothesis or conjecture to the state of having a genuine belief. Nagel writes:

> Achieving closure or judgmental commitment on a question puts an end to the experience of ambiguity and delivers the sense of having a firm answer. The opposite of closure is openness or judgmental non-commitment, in which we are able to continue juggling alternative possibilities, perhaps lingering in ambiguity or confusion. (2008: 286)

Thus elevated levels of epistemic anxiety tend to produce two important differences in the thought and behavior of high versus low stakes subjects – the former rely on *more systematic and thorough* evidence-gathering strategies and the they also exhibit a *diminished need-for-closure*. Nagel uses these descriptive generalizations to produce two quite distinct kinds of arguments against IRI.

The first kind of argument, the origins of which Nagel attributes to Kent Bach, says that stakes affect knowledge only by affecting our perceptions of whether the character *believes* the relevant proposition. In particular, a high stakes subject experiences higher epistemic anxiety and lower need-for-closure, and thus she will be more hesitant to form a settled belief. Because belief is required for knowledge, a high stakes subject will thus lack knowledge as well:

> High-Stakes Hannah is a low need-for-closure subject: she is strongly averse to making a mistake about the banking hours, and willing to entertain hypotheses that would make her initial evidence inconclusive. Her Low Stakes counterpart is in a neutral need-for-closure condition: Low-Stakes Hannah is not described as being under anything like the pressure for immediate decision characteristic of high-need-for-closure subjects, nor is she strongly averse to making the wrong call about Saturday banking—the scenario suggests that getting it wrong would be a mild inconvenience at

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

> worst. If our intuitions about the shift in Stanley's cases are driven by the contrast between low and neutral need-for-closure, then something like the Bach objection will be right: we ascribe knowledge in Low Stakes and deny it in High because we naturally attribute higher and lower confidence belief to the contrasted subjects, or confident belief and a state of evidence assessment that precedes fixed belief. (2008: 289)

Nagel anticipates an important objection to this Stakes Affect Belief argument. What if we simply stipulate that the contrasted characters both form a settled belief in which they have equal levels of confidence? For example, in DeRose and Stanley's bank cases, we could stipulate that Hannah forms a confident belief that the bank will be open on Saturday. The Stakes Affect Belief argument predicts that, given this stipulation, the asymmetric intuitions between high and low stakes should disappear. But a quick check of intuitions suggests that this does not in fact happen. Evidence from experimental studies confirms this impression. In the three vignettes we presented (Basic, Explicit/Implicit and Ignorant), as well as in Pinillos's Typo-Explicit Belief vignette, the fact that the subject formed a belief in the relevant proposition is clearly stated in the vignette. Yet an asymmetry in knowledge (or quality of evidence) judgments between high and low stakes is observed in all of the vignettes.

Perhaps recognizing that this objection has considerable bite, Nagel offers a second kind of argument against IRI. This argument is subtler than the Stakes Affect Belief Argument. In particular, this second argument admits to two interpretations, a weaker and a stronger version. We will argue that the weaker version of the argument simply *is* IRI masquerading under a different name. The stronger version of the argument is distinct from IRI, but this version is implausible.

Suppose that Paul is in a high stakes situation, and he uses an evidence-gathering strategy that one would expect to be deployed in a low stakes situation. For example, suppose that his life depends on the bank being open Saturday, and he forms the belief that the bank will open on Saturday by recalling a vague memory of being at the bank a Saturday several months ago. If we conclude that Paul does not know the bank will be open on Saturday, what are the mediating inferences that might have gotten us to that conclusion?

The weaker version of Nagel's account addresses this question as follows. Paul is in a high stakes situation, which generates considerable epistemic anxiety. However, the evidence-gathering strategy he uses falls well short of what the epistemic anxiety aroused by his current situation demands. Paul's evidence-gathering strategy seems only good enough to satisfy the epistemic anxiety associated with a low stakes situation. The gap between the quality of the evidence-gathering strategy he deploys and what the situation demands leads to the perception that Paul is engaged in 'wishful thinking' or has 'unfounded confidence'. Indeed, on the present view, *the very same evidence-gathering strategy*, and in particular one that sufficed for knowledge in a low stakes situation, might be deemed inadequate in a high stakes situation in virtue of failing to satisfy the elevated epistemic anxiety aroused by the higher stakes. This is how Nagel puts the point:

> On this view, the anxious high stakes may need to visit the branch to double-check the banking hours in order to seem as accurate as the casual low stakes subject who just remembers a recent Saturday visit to the bank; if the high-stakes subject doesn't work harder than his low stakes counterpart, then he will come across as having a problematic basis for belief. Although we expect different behavior of these two subjects, we are not holding them to different standards: as far as

> reliability is concerned, both subjects are benchmarked against the performance that would be anticipated if nothing opposed the level of epistemic anxiety that subjects automatically feel on their behalf. (2011: 419–20)

Notice that this way of understanding Nagel's account (i.e. the weaker version) is *exactly* in line with IRI. IRI claims the practical costs of being wrong enter into the possession conditions for knowledge and other epistemic properties and relations. Thus an evidence-gathering strategy that would suffice for knowledge in a low stakes situation might be perceived as inadequate (e.g. as an instance of wishful or careless thinking) when the stakes are high. Nagel seems to be making exactly this same point, though she does not acknowledge the role that IRI-licensed inferences play in her account. As we see it, the key problem of this weaker version of Nagel's account is that it too freely helps itself to terms such as 'wishful thinking' and 'unfounded confidence' when describing characters in high stakes situations. What makes it the case that these epistemic pejoratives are deserved in the first place is precisely the fact that when the stakes are high, the possession conditions for knowledge and other epistemic properties become more stringent.

The challenge for Nagel, then, is to clarify how stakes-sensitive intuitions arise in a way that *differs* from IRI. To do this, Nagel needs to add an additional argument to the weaker version of her view that we've discussed so far. We call the extra element that Nagel needs the Stakes Affect Strategies Argument. Suppose there are two situations. In low-stakes situation **L**, most people use the relatively easy and only moderately reliable evidence gathering strategy **1**, and they form a settled belief rather quickly. In high-stakes situation **H**, most people use the highly demanding evidence-gathering strategy **2**, and they form a settled belief rather slowly. Suppose Paul forms a belief rather quickly in **H**. And let us suppose that we intuitively judge that Paul lacks knowledge. Here is a *non-IRI* explanation for why Paul lacks knowledge: Paul forms a belief quickly in **H**. Since the strategy he uses (let us call it **3**) cannot be the highly demanding one that most people use (i.e. **2**), there must be some *epistemically damaging factors* **F** (such as haste, distraction, etc. ...) that explain why Paul does not use **2**. Whatever these factors turn out to be, they will often lead us to believe that **3** is some highly unreliable strategy. In particular, we tend to infer that **3** is epistemically deficient not only compared to **2**, *but also compared to* **1** *as well*. Call an evidence-gathering strategy *defective* if it exhibits insufficient reliability to qualify a person for knowledge in *both high and low stakes* situations. If **3** is defective, then Paul lacks knowledge, and the reason has nothing to do with knowledge itself being stakes-sensitive. The explanation for why Paul lacks knowledge is because the strategy he uses (i.e. **3**) lacks adequate reliability, irrespective of stakes. Moreover, one does not need IRI to deliver the conclusion that using a defective evidence-gathering strategy disqualifies one from knowledge. This is something that the traditional intellectualist epistemologist would have said anyways.

We've now characterized two versions of Nagel's view, a weaker and a stronger version. It is somewhat unclear which version Nagel intends, though in some cases, it does appear that she is indeed advancing the stronger version:

> For example, if someone is in a high-stakes situation and declines to pursue readily available evidence on a question that should be provoking high epistemic anxiety, it would be natural for us to attribute to him some desire or condition overshadowing his natural desire for increased cognitive effort. If we see this condition as the basis of his belief, then his judgment may naturally seem *less reliable than the judgment of his low-stakes counterpart*. (2011: 15–16; emphasis ours)

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

Elsewhere she says,

> Because High Stakes subjects ordinarily experience low need-for-closure, it's only psychologically plausible that High-Stakes Hannah could have high confidence of the sort stipulated on her slim evidence if she is compromised in her accuracy, for example by thinking hastily or in a way biased by wishful thinking … Ordinarily, high- and low-stakes subjects think differently about problems complex enough to generate contrasting epistemic intuitions: we can make high-stakes subjects think like low-stakes subjects on such problems, *but only by putting them under conditions where they think less accurately than their low-stakes counterparts*. To the extent that our knowledge ascriptions are sensitive to stakes through being sensitive to lowered accuracy, they cannot be used by the advocate of IRI to establish the most provocative part of his thesis, namely that knowledge itself involves certain factors that are irreducibly practical rather than truth-conducive. (2008: 291–2; emphasis ours)

The emphasized text in the passages above suggests that when a high-stakes subject does not exhibit the care and consideration we expect of subjects facing high stakes, we tend to infer that the strategy used is inadequate not just relative to the stakes (which is something an IRI-theorist also thinks), but rather that the strategy used is inadequate *simpliciter* (i.e. that it is epistemically defective).

We've already explained our main objection to the weaker version of Nagel's view, i.e. that it simply *is* a version of IRI. Now we'll discuss some objections to the stronger version of Nagel's view. A key problem with the Stakes Affect Strategies argument is that it is only applicable to hypothetical cases where we are uncertain about the strategy actually used by the characters. But a number of the cases used to illustrate IRI *stipulate* the evidence gathering strategy used by the characters. For example, in all three vignettes we presented (Basic, Explicit/Implicit, and Ignorant), the vignette explicitly states the evidence-gathering strategy the character used: 'Hannah notes that the menu says her dish does not contain pine nuts. Based on this, Hannah forms the belief that her noodles are not topped with pine nuts', and this description does not differ between low and high stakes versions of the cases. Thus it is not plausible that people perceive differences in the actual evidence-gathering strategies the characters deployed, and in particular that the strategy used in the high stakes case is perceived to be epistemically defective, while the strategy used in the low-stakes case is not perceived this way. Rather, a more plausible hypothesis is that, given that the two characters use the *same* evidence-gathering strategies, stakes are affecting whether the use of this strategy is sufficient for the character's having knowledge.

An additional challenge to the Stakes Affect Strategies argument consists in cases in which the character is ignorant of the fact that she is in a high stakes situation, and instead thinks that she is in a low stakes situation. Our experimental studies, as well as the studies of Pinillos discussed above, found that stakes continue to exert a significant effect on knowledge even when the character is completely ignorant of the stakes. This challenges the Stakes Affect Strategies argument because it blocks the key inference on which the argument depends. That is, it would seem that, if a character does not know the stakes involved, from the fact that she uses an evidence-gathering strategy that is less accurate than a typical high-stakes strategy, we cannot infer that that epistemically damaging factors (such as haste and distraction) are at play. But if this is so, then why do we refuse to credit her with knowledge? The most plausible explanation is the one put forward by IRI:

the strategy she uses is epistemically inadequate given the more stringent possession conditions for knowledge associated with high stakes situations.[7]

5. WHAT PRACTICAL FACTORS ARE EPISTEMICALLY RELEVANT?

One traditional way to argue for IRI is via links between knowledge and action of the sort defended by Fantl and McGrath (2002, 2009), Hawthorne (2004), Stanley (2005), and Hawthorne and Stanley (2008). Here is a version of this argument for IRI that depends on one direction of a principle that Hawthorne and Stanley (2008) call the 'Reason-Knowledge Principle':

> Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff one knows that p.

The relevant direction is from right to left:

> If one knows that p, then it is appropriate to treat the proposition that p as a reason for acting.

One can use this direction of the Reason-Knowledge Principle to argue for skepticism. If knowledge gives one an appropriate reason for acting, then it is appropriate for someone to bet her life on some mundane proposition that she ordinarily takes herself to know, such as the proposition that her refrigerator is in good working order. But it is not appropriate to take one's mundane ordinary knowledge as a reason for acting when stakes have been elevated. So it appears that the Reason-Knowledge principle threatens an untenable skepticism.

The argument for IRI is that it allows one to preserve the Reason-Knowledge Principle without embracing skepticism. As Hawthorne and Stanley (2008) write:

> If knowledge is constitutively related to one's practical environment, then it is open to us to claim that while one may know that p in a situation where not much is at stake as concerns the proposition that p, one loses knowledge once one enters an environment where a good deal is at stake as regards the truth or falsity of the proposition that p. On such a view knowledge … vanishes when the stakes go up.

When I face a bet upon which my life depends, certain propositions are no longer appropriate to take as reasons for acting. But given the stakes-sensitivity of knowledge, they are also no longer known. This method of arguing for IRI suggests that the practical factors that affect knowledge are precisely those that are relevant for when it is appropriate to take something as a reason for acting. Jeremy Fantl and

7 Nagel does discuss Ignorant High Stakes cases in Nagel (2008, 2010). In short, she provides an error theory for these cases. We believe it is a cost for Nagel's view that it has one account for non-ignorant high stakes cases and an entirely different error-theoretic account for ignorant high stakes cases. IRI provides a much more uniform account of how stakes effect epistemic properties and relations that applies to both non-ignorant as well as ignorant high stakes cases.

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

Matthew McGrath (2009: 84ff.) develop this argument from the Reason-Knowledge principle to IRI in admirable detail.

One might think that the conditions under which it is appropriate to take a proposition as a reason for acting are affected only by what is in an agent's purview. If so, then if an agent is unaware, and couldn't possibly be aware, of some factors that impinge on the utility of the choice she faces, then these factors cannot affect the conditions under which it is appropriate to take a proposition as a reason for acting. However, as we have seen from Ignorant High Stakes, we *do* take factors outside an agent's purview to affect whether or not an agent has knowledge or good evidence. Insofar as the knowledge–action links are an argument for IRI, the notion of 'appropriate for action' that is relevant in them must be sensitive to factors of which the agent is potentially not aware.

The thought that the practical factors that are epistemically relevant are not confined to what the agent recognizes about her situation has been recognized by some advocates of IRI. Stanley (2005: ch. 5) argues that the notion relevant to epistemology is what he calls a *serious practical question*. When a proposition is a serious practical question in a situation, then higher standards of evidence are in play. The intuitive idea is that a proposition is a serious practical question for a person at a time if and only if that person ought to take that proposition into account in decision-making.

In trying to explicate the notion of a serious practical question, Stanley (2005: 94–5) appeals to warranted expected utility:

> A proposition p is practically irrelevant if and only if, where $a_1 \ldots a_n$ are the actions at my disposal, the differences between the warranted expected utilities of $a_1 \ldots a_n$ relative to the nearest states of the world in which p are not meaningfully different from the warranted expected utilities of $a_1 \ldots a_n$ relative to the nearest states of the world in which $\sim$p.

Immediately after characterizing the notion of a serious practical question in terms of *warranted* expected utility, Stanley notes that a more objective notion is ultimately needed – that a serious practical question may need to be defined in terms of *objective* utility, rather than warranted expected utility:

> I have characterized this account in terms of the *warranted* expected utilities of the actions at an agent's disposal, because the agent might not be aware of what is in her own best interest, though, if she had some additional information, she would be. So the expected utility calculation, on this account, should be thought of as pertaining to the agent's warranted expected utilities, and not her subjective credences. *But warranted expected utility is probably not sufficiently impersonal of a notion to do the required epistemic work. There may be facts relevant to the utility calculation that the agent is not epistemically responsible for knowing. So a more impersonal notion of utility may be required to capture the notion of a serious practical question.* (Stanley 2005; emphasis ours)

Our findings corroborate Stanley's suspicion. In Ignorant High Stakes, it is made vivid that the relevant agent is not only unaware of the stakes, but is not responsible for knowing the stakes ('Hannah has absolutely no idea that she has this gene, nor is there any way she could know that she has this gene.'). Nevertheless, the stakes are intuitively epistemically relevant. Therefore, intuition supports the view that it is the presence of the stakes themselves, rather than the agent's awareness of the stakes, which is epistemically relevant.

Thus, our findings show that the intuitive case for IRI is one that supports the thesis that it is the impact of a decision on the objective utility of an agent that has epistemic relevance. This is a step forward in our understanding of the basis for IRI. For example, insofar as one is drawn to IRI by the kinds of knowledge-action links discussed by Hawthorne, Stanley, Fantl, and McGrath, we must understand these normative principles as concerning reasons for acting that are not merely *subjective normative reasons* (which, very roughly, are reasons that obtain in the nearest world in which the agent's (conscious) beliefs are true, and are thus, by definition, within the agent's purview). Rather, the relevant normative principals appear to be sensitive to *objective normative reasons*, which hold irrespective of whether the agent believes that such reasons obtain.

ACKNOWLEDGEMENTS

Thanks to Jim Joyce and Alexandra Plakias for discussion of the issues in this paper. We also presented this paper in a conference sponsored by the University of Geneva. We are grateful to the audience there for discussion, in particular Igor Douven and Pascal Engel. We are particularly indebted to Jeremy Fantl for helping us think through some of the points in the last section. Finally, Jennifer Nagel provided very helpful written comments on the penultimate draft.

REFERENCES

Buckwalter, Wesley. 2010. 'Knowledge isn't Closed on Saturday: A Study in Ordinary Language.' *Review of Philosophy and Psychology*, 1(3): 395–406.

DeRose, Keith. Forthcoming. 'Contextualism, Contrastivism, and X-Phi Surveys.' *Philosophical Studies*.

Fantl, Jeremy, and McGrath, Matthew. 2002. 'Evidence, Pragmatics, and Justification.' *Philosophical Review*, 111: 67–94.

2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Feltz, Adam, and Zarpentine, Chris. Forthcoming. 'Do you Know More When it Matters Less?' *Philosophical Psychology*.

Hawthorne, John. 2000. 'Implicit Belief and a Priori Knowledge.' *Southern Journal of Philosophy*, 38 (S1): 191–210.

2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hawthorne, John, and Stanley, Jason. 2008. 'Knowledge and Action.' *Journal of Philosophy*, 105: 571–90.

Maasen, Gerard H., and Bakker, Arnold B. 2001. 'Suppressor Variables in Path Models.' *Sociological Methods and Research*, 30(2): 241–70.

May, Joshua, Sinnott-Armstrong, Walter, Hull, Jay, and Zimmerman, Aaron. 2010. 'Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study.' *Review of Philosophy and Psychology*, 1(2): 265–73.

Nagel, Jennifer. 2008. 'Knowledge Ascriptions and the Psychological Consequences of Changing Stakes.' *Australasian Journal of Philosophy*, 86(2): 279–94.

2010. 'Epistemic Anxiety and Adaptive Invariantism.' *Philosophical Perspectives*, 24(1): 407–35.

Pinillos, Angel. Forthcoming. 'Knowledge, Experiments and Practical Interests.' In J. Brown and M. Gerken (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press.

Podsakoff, Philip M., MacKenzie, Scott B., Lee, Jeong-Yeon, and Podsakoff, Nathan P. 2003. 'Common Method Biases in Behavioral Research: A Critical Review of the Literature and Recommended Remedies.' *Journal of Applied Psychology*, 88(5): 897–903.

https://doi.org/10.1017/epi.2011.2 Published online by Cambridge University Press

Schaffer, Jonathon, and Knobe, Joshua. Forthcoming. 'Contrastive Knowledge Surveyed.' *Nous*.
Schroeder, Mark. Forthcoming. 'Stakes, Withholding, and Pragmatic Encroachment on Knowledge.' *Philosophical Studies*.
Stanley, Jason. 2005. *Knowledge and Practical Interests*. New York: Oxford University Press.
Weinberg, Jonathan M., Nichols, Shaun, and Stich, Stephen. 2001. 'Normativity and Epistemic Intuitions.' *Philosophical Topics*, 29: 429–60.

CHANDRA SEKHAR SRIPADA is Assistant Professor of Philosophy and Psychiatry at the University of Michigan, Ann Arbor. His research lies at the intersection of philosophy (especially moral psychology), neuroscience, and psychiatry. His papers have appeared in journals in philosophy (Philosophy and Phenomenological Research, Philosophical Studies), and the sciences (Proceedings of the National Academy of Sciences, Biological Psychiatry). University of Michigan, 2215 Angell Hall, 435 South State Street, Ann Arbor, MI 48109-1003, sripada@umich.edu, (734) 358-3297.

JASON STANLEY is Professor of Philosophy at the Rutgers University, New Brunswick, member of the Rutgers University Center for Cognitive Science, and associate faculty member in the Department of Linguistics. He previously taught at the University of Michigan at Ann Arbor and Cornell University. He has published three books with Oxford University Press, Knowledge and Practical Interests (2005), Language in Context (2007), and Know How (2011). Department of Philosophy, College Avenue Campus, 1 Seminary Place, New Brunswick, NJ 08901.