

Interest-Relative Invariantism

Brian Weatherson

1 Introduction

One of the initial motivations for epistemological contextualism was that the appropriateness of self-ascriptions of knowledge seemed to depend, in some circumstances, on factors that were traditionally thought to be epistemologically irrelevant. So whether our hero *S* was prepared to say *I know that p* would depend not just on how strong *S*'s evidence for *p* was, or how strongly they believed it, but on factors such as how much it mattered whether *p* was true, or what alternatives to *p* were salient in their thought or talk.

It was immediately noted that this data point, even if accepted, is consistent with a number of theories of the truth of knowledge ascriptions. It might be that things like stakes and salient alternatives affect the assertability conditions of knowledge ascriptions, but not their truth conditions.¹ But let's assume that we've convinced ourselves that this isn't right, and that whether *S* can truly (and not just appropriately) say *I know that p* depends on things like the stakes or salient alternatives.

It still doesn't follow that contextualism is true. It might be that in all contexts, whether an utterance of *S knows that p* is true depends on the stakes for *S*, or on the salient alternatives for *S*. That would be true, the idea is, whether *S* is talking about herself, or someone else is talking about her. The stakes, or salient alternatives, would affect the truth conditions of *S*'s utterance not because she is the one doing the talking, but the one being talked about. The practical and theoretical situation of the ascriber of the knowledge ascription may be relevant, even if the practical and theoretical situation of the ascriber need not be.

This line of thought leads to the idea that knowledge itself is interest-relative. Whether *S knows that p* is true, i.e., whether *S* knows that *p*, depends on how much it matters to *S* that *p* is true, or on which alternative are salient to *S*. The thesis that knowledge is interest-relative is consistent with contextualism. It could be that whether a knowledge ascription is true depends on the interests of both the ascriber, and the ascribee. In this entry, however, I'm going to largely focus on the view that knowledge is interest-relative, but contextualism is false. On this view, the interests of the ascribee do matter to the truth of a knowledge ascription, but the interests of the ascriber do not.

This view is naturally called **interest-relative invariantism**, since it makes knowledge interest-relative, but it is a form of anti-contextualism, i.e., invariantism.

¹Note to editors: If there is an article on warranted assertability manoeuvres in the book, it would be good to refer to it here. Otherwise I'll add a reference to a journal article or two.

The view is sometimes called **subject-sensitive invariantism**, since it makes knowledge relevant to the stakes and salient alternatives to the subject. But this is a bad name; of course whether a knowledge ascription is true is sensitive to who the subject of the ascription is. I know what I had for breakfast and you (probably) don't. What is distinctive is which features of the subject's situation that interest-relative invariantism says are relevant, and the name interest-relative invariantism makes it clear that it is the subject's interests. (I'll shorten 'interest-relative invariantism' to IRI in what follows.)

There are a number of ways to motivate and precisify IRI. I'll spend most of this entry going over the choice points, starting with the points where I think there is a clearly preferable option, and ending with the choices where I think it's unclear which way to go. Then I'll discuss some general objections to IRI, and say how they might be answered.

2 Motivations

There are two primary motivations for IRI. One comes from intuitions about cases, the other from a pair of principles. It turns out the two are connected, but it helps to start seeing them separately.

Jason Stanley (2005) starts with some versions of the 'bank cases' due originally to Keith DeRose (1992). These turn on idiosyncratic, archaic details of the US payments system, and I find it hard to have clear intuitions about them. A cleaner pair of examples is provided by Angel Pinillos (2012); here are slightly modified versions of his examples.

Ankita and Bojan each have an essay due. They have, surprisingly, written word for word identical papers, and are now checking the paper for typos. The paper has no typos, and each has checked their paper twice, with the same dictionary, and not found any typos. They are, in general, equally good at finding typos, and have true beliefs about their proficiency at typo-spotting.

The only difference between them concerns the consequence of a typo remaining. If the paper is a borderline A/A- paper, a typo might mean Ankita gets an A- rather than an A. But the grade doesn't matter to her; she's already been accepted into a good graduate program next year so long as she gets above a C. But Bojan's instructor is a stickler for spelling. Any typo and he gets a C on the paper. And he has a very lucrative scholarship that he loses if he doesn't get at least a B on this paper. (Compare the Typo-Low and Typo-High examples in Pinillos (2012, 199).)

The intuition that helps IRI is that Ankita knows she has no typos in her paper, and should turn it in, while Bojan does not know this, and should do a third (and perhaps fourth or fifth) check. Contextualists have a hard time explaining this; in this very context I can say “Ankita knows her paper has no typos, but Bojan does not know his paper has no typos”. If the intuition is right, it seems to support interest-relativity, since the difference in practical situation between Ankita and Bojan seems best placed to explain their epistemic difference. That’s the argument from cases.

The argument from principles takes off from the idea that knowledge plays an important role in good deliberation, and that knowledge does not require maximal confidence. It is easiest to introduce with an example, though note that we aren’t going to rely on epistemic intuitions about the example. Chika looked at the baseball scores last night before going to bed and saw that the Red Sox won. She remembers this when she wakes up, though she knows that she does sometimes misremember baseball scores. She is then faced with the following choice: take the red ticket, which she knows pays \$1 if the Red Sox won last night, and nothing otherwise or the blue ticket, which she knows pays \$1 iff $2+2=4$. Now consider the following principle, named by Jessica Brown (forthcoming):

K-Suff If S knows that p , then S can rationally take p as given in practical deliberation.

The following trio seems to be inconsistent:

1. Chika knows the Red Sox won last night.
2. Chika is rationally required to take the blue ticket.
3. K-Suff is true.

By 1 and 3, Chika can take for granted that the Red Sox won last night. So the value of the red ticket, for her, is equal to its value conditional on the Red Sox winning. And that is \$1. So it is at least as valuable as the blue ticket. So she can’t be rationally required to take the blue ticket. Hence the three propositions are inconsistent.

This is worrying for two reasons. For one thing, it is intuitive that Chika knows that the Red Sox won. For another thing, it seems this form of argument generalises. For almost any proposition at all, if Chika knows the red ticket pays out iff that proposition is true, she should prefer the blue ticket. So she knows very little.

How could this argument be resisted? One move, which we’ll return to frequently, is to deny K-Suff. Maybe Chika’s knowledge that the Red Sox won is insufficient; she needs to be certain, or to have some higher order knowledge. But denying K-Suff alone will not explain why Chika should take the blue ticket. After

all, if **K-Suff** is false, the fact that Chika knows the payout terms of the tickets is not in itself a reason for her to choose the blue ticket.

So perhaps we could deny that she is rationally required to choose the blue ticket. This does seem extremely unintuitive to me. Intuitions around here do not seem maximally reliable, but this is a strong enough intuition to make it worthwhile to explore other options.

And IRI provides a clever way out of the dilemma. Chika does not know the Red Sox won last night. But she did know that, before the choice was offered. Once she has that choice, her knowledge changes, and now she does not know. The intuition that she knows is explained by the fact that relative to a more normal choice set, she can take the Red Sox win as a given. And scepticism is averted because Chika does normally know a lot; it's just in the context of strange choices that she loses knowledge.

The plotline here, that principles connecting knowledge and action run up against anti-sceptical principles in contrived choice situations, and that IRI provides a way out of the tangle, is familiar. It is, simplifying greatly, the argumentative structure put forward by Hawthorne (2004), and by Fantl and McGrath (2002, 2009), and by Weatherson (2012). It does rely on intuitions, but they are intuitions about choices (such as that Chika should choose the blue ticket), not about knowledge directly.

Some discussions of IRI, especially that in Hawthorne and Stanley (2008) use a converse principle. Again following the naming convention suggested by Jessica Brown (forthcoming), we'll call this **K-Nec**.

K-Nec An agent can properly use p as a reason for action only if she knows that p .

I'll mostly set the discussion of **K-Nec** aside here, since my preferred argument for IRI, the argument from Chika's case, merely relies on **K-Suff**. But it is interesting to work through how **K-Nec** helps plug a gap in the argument by cases for IRI.

Buckwalter and Schaffer (2015) argue that the intuitions behind Pinillos's examples are not as solid as we might like. It's true that experimental subjects do say that Bojan has to check the paper more times than Ankita does before he knows that the paper contains no typos. But those subjects also say he has to check more times before he believes that the paper has no typos. And, surprisingly, they say that he has to check more time before he guesses the paper has no typos. They suggest that there might be interest-relativity in the modal 'has' as much as in the verb 'knows'. To say someone 'has' to X before they Y, typically means that it is improper, in some way, to Y without doing X first. That won't be a problem for the proponent of IRI as long as at least in some of the cases Pinillos studies, the relevant senses of propriety are connected to knowledge. And that's plausible for belief; Bojan has to

know the paper is typo free before he (properly) believes it. At least, that's a plausible move given K-Nec.

There is one other problem for argument from cases for IRI. Imagine that after two checks of the paper, we tell Bojan that Ankita has checked her paper in just the same way he has checked his. And we tell him that Ankita does not overly care whether her paper is typo-free, but is confident that it is. We then ask him, does Ankita know her paper is typo free? Many philosophers think Bojan should answer "No" here. And that isn't something IRI can explain. According to IRI, he should say, "I don't know." He can't say Ankita does know, since he doesn't know their common paper has no typos. But it's hard to see why he should deny knowledge. Keith DeRose (2009, 185) thinks this case is particularly hard for IRI to explain, while Brian Kim (forthcoming) offers some possible explanations. This objection doesn't tell against the claim that knowledge is interest-relative, but it does threaten the invariantism. An interest-relative contextualist should say that everyone should deny Bojan knows his paper is typo free, and Bojan should deny Ankita knows her paper is typo-free.

3 Odds and Stakes

Interest-relative invariantism says that the interests of the subject matter to what she knows. This is a fairly vague statement though; there are a number of ways to make it precise. Right now I have interests in practical questions (such as whether I should keep writing or go to lunch) and in theoretical questions (such as whether IRI is true). Do both kinds of interests matter? We'll return to that question in the next section. For now we want to ask a prior question: when do practical interests matter for whether a subject knows that p ? There are two main answers to this question in the literature.

Stakes When the agent has a possible bet on p that involves large potential losses, it is harder to know that p .

Odds When the agent has a possible bet on p that involves long odds, it is harder to know that p .

The difference between these two options becomes clear in a simple class of cases. Assume the agent is faced with a choice with the following structure:

- There is a safe option, with payout S .
- And there is a risky option, with good payout G if p is true, and bad payout B if p is false.

These choices need not involve anything like a ‘bet’, in the ordinary folk sense. But they are situations where the agent has to make a choice between a path where the payouts are p -dependent, and one where they are independent of p . And those are quite common situations.

The **Stakes** option says that the relevant number here is the magnitude $S - B$. If that is large, then the agent is in a high-stakes situation, and knowledge is hard. If it is low, then the agent is in a low stakes situation, and knowledge is relatively easy. (Perhaps the magnitude of $G - S$ is relevant as well, though the focus in the literature has been on examples where $S - B$ is high.)

The **Odds** option says that the relevant number is the ratio:

$$\frac{S - B}{G - S}$$

If that number is high, the agent faces a long odds bet, and knowledge is hard. If that number is low, the agent faces a short odds bet, and knowledge is relatively easy.

If our motivation for IRI came from cases, then it is natural to believe **Stakes**. Both Bojan and Chika face bets on p at long odds, but intuition raises greater worries about whether Bojan knows that p to whether Chika does.

But if our motivation for IRI came from principles, then it is natural to believe **Odds**. One way to think of the argument from principles for IRI is that it is a way to make all four of the following intuitive claims true:

1. Agents should maximise evidential expected utility; i.e., they should choose the option whose expected utility is highest if the utilities are the agent’s own, and the probabilities are the evidential probabilities given the agent’s evidence.
2. If an agent knows that p , they can ignore possibilities where p is false; i.e., they can make whatever choice is the rational choice given p .
3. Chika cannot ignore possibilities where the Red Sox lost; she should consider those possibilities because it is in virtue of them that the evidential expected utility of taking the red ticket is higher.
4. Agents with Chika’s evidence, background and dispositions typically know that the Red Sox won.

The first three principles imply that Chika does not know the Red Sox won. The only way to square that with the anti-sceptical fourth principle is to say that Chika is in some way atypical. And the only way she has been said to be atypical is in the practical choices she faces. But note it is not because she faces a high-stakes choice: precisely one dollar is at stake. It is because she faces a long (indeed infinitely long) odds bet.

In the general case we discussed above, agents maximise expected utility by taking the risky choice iff:

$$\frac{S-B}{G-S} < \frac{Pr(p)}{1-Pr(p)}$$

where $Pr(p)$ is the probability of p given the agent's evidence. The actual magnitudes at play don't matter to what choice maximises expected utility, just the odds the agent faces. So if one's motivation to keep IRI is to square expected utility maximisation with natural principles about knowledge and action, it seems the relevant feature of practical situations should be the stakes agents face.

Why could it seem stakes matter then? I think it is because in high stakes situations, the odds an agent faces are typically long ones. It is much easier to lose large amounts of utility than to gain large amounts of utility. Bojan stands to lose a lot from a typo in his paper; he doesn't stand to lose much by taking the time to check it over. So a high stakes situation will, at least typically, be a long odds situation. So if we say the odds the agent faces are relevant to what they know, we can explain any intuition that the stakes at play are relevant.

Jessica Brown (2008, 176) also notes that cases where the agent faces long odds but low stakes raise problems for the stakes-based version of IRI.

4 What Kind of Interests?

Let's return to the question of whether theoretical interests are relevant to knowledge, or only practical interests. There is some precedent for the more restrictive answer. Stanley's book on IRI is called *Knowledge and Practical Interests*. And he defends a theory on which what an agent knows depends on the practical questions they face. But there are strong reasons to think that theoretical reasons matter as well.

In the previous section, I suggested that agents know that p only if they would maximise expected utility by choosing the choice that would be rational given p . That is, agents know that p only if the answer to the question *What choice maximises expected utility?* is the same unconditionally as it is conditional on p . My preferred version of p generalises this approach. An agent knows that p only if the rational answer to a question she faces is the same unconditionally as it is conditional on p . What it is for an agent to face a question is dependent on the agent's interests. If that's how one thinks of IRI, the question of this section becomes, should we restrict questions the agent faces to just being questions about what choice to make? Or should they include questions that turn on her theoretical interests, but which are irrelevant to choices before her. There are two primary motivations for allowing theoretical interests as well as practical interests to matter.

The first comes from the arguments for what Jeremy Fantl and Matthew McGrath call the Unity Thesis (Fantl and McGrath, 2009, 73–6). They are interested in the thesis that whether or not p is a reason for an agent is independent of whether the agent is engaged in practical or theoretical deliberation. But we don't have to be so invested in the ideology of reasons to appreciate their argument. Note that if only practical interests matter, then the agent should come up with different answers to the question *What to do in situation S* depending on whether the agent is actually in S , or they are merely musing about how one would deal with that situation. And it is unintuitive that this should matter.

Let's make that a little less abstract. Imagine Chika is not actually faced with the choice between the red and blue tickets. In fact, she has no practical decision to make that turns on whether the Red Sox won. But she is idly musing over what she would do if she were offered the red ticket and the blue ticket. If she knows the Red Sox won, then she should be indifferent between the tickets. After all, she knows they will both return \$1. But intuitively she should think the red ticket is preferable, even in the abstract setting. And this seems to be the totally general case.

The general lesson is that if whether one can take p for granted is relevant to the choice between A and B, it is similarly relevant to the theoretical question of whether one would choose A or B, given a choice. And since those questions should receive the same answer, if p can't be known while making the practical deliberation between A and B, it can't be known while musing on whether A or B is more choiceworthy.

Weatherson (2012) suggests another reasons for including theoretical interests in what's relevant to knowledge. There is something odd about the following reasoning: *The probability of p is precisely x , therefore p , in any case where $x < 1$.* It is a little hard to say, though, why this is problematic, since we often take ourselves to know things on what we would admit, if pushed, are purely probabilistic grounds. The version of IRI that includes theoretical interests allows for this. If we are consciously thinking about whether the probability of p is x , then that's a relevant question to us. Conditional on p , the answer to that question is clearly no, since conditional on p , the probability of p is 1. So anyone who is thinking about the precise probability of p , and not thinking it is 1, is not in a position to know p . And that's why it is wrong, when thinking about p 's probability, to infer p from its high probability.

Putting the ideas so far together, we get the following picture of how interests matter. An agent knows that p only if the evidential probability of p is close enough to certainty for all the purposes that are relevant, given the agent's theoretical and practical interests. Assuming the background theory of knowledge is non-sceptical, this will entail that interests matter.

5 Global or Partial

So far I've described three ways to refine the defence of IRI.

1. The motivation could come from cases or principles.
2. The relevant feature that makes it hard to have knowledge could be that the agent faces a high-stakes choice, or a long-odds choice.
3. Only practical interests may be relevant to knowledge, or theoretical interests may matter as well.

For better or worse, the version of IRI I've defended has fairly clear commitments on all three; in each case, I prefer the latter option, though there are also good arguments for the former. From here on, I'm much less sure of the right way to refine IRI.

IRI, like contextualism, was introduced as a thesis about knowledge. But it need not be restricted that way. It could be generalised to a number of other epistemically interesting notion. At the extreme, we could argue that every epistemologically interesting notion is interest-relative. Doing so gives us a global version of IRI.

Jason Stanley (2005) comes close to defending a global version. He notes that if one has both IRI, and a 'knowledge first' epistemology (Williamson, 2000), then one is a long way to towards globalism. Even if one doesn't accept the whole knowledge first package, but just accepts the thesis that evidence is all and only what one knows, then one is a long way towards globalism. After all, if evidence is interest-relative, then probability justification, rationality and evidential support are interest-relative too.

Katherine Rubin (2015) objects to globalist versions of IRI. But the objections she gives turn, as she notes, on taking stakes not odds to be relevant.

If a non-global version of IRI could be made to work, it would have some theoretical advantages. It's nice to be able to say that Chika should take the blue ticket because the evidential probability of the Red Sox winning is lower than the evidential probability of two plus two being four. But that won't be a non-circular explanation if we also say that something is part of Chika's evidence in virtue of being known.

On the other hand, the motivations for interest-relativity of knowledge seem to generalise to all other non-gradable states. In ordinary cases, Chika could use the fact that the Red Sox won as a given in practical or theoretical reasoning. That is, she could properly treat it as evidence. But she can't treat it as evidence when deciding which ticket to take. So at least what she can properly treat as evidence seems to be interest-relative, and from there it isn't obvious how to deny that evidence itself is interest-relative too.

There remains a question of whether gradable notions, like epistemic probabilities, are also interest-relative. One of the aims of my first paper on IRI (Weatherson,

2005) was to argue that probabilistic notions are interest-invariant while binary notions are interest-relative. But if propositions that are part of one's evidence have maximal probability (in the relevant sense of probability), and evidence is interest-relative, that combination won't be sustainable.

In short, while the non-global version of IRI allows for some nice reductive explanations of why interests matter, the global version is supported by the very intuitions that motivated IRI. There is a danger here that whatever way the IRI theorist goes, they will run into insuperable difficulties. Ichikawa et al. (2012) argue strongly that this danger is real; there is no plausible way to fill out IRI. I'm not convinced that the prospects are quite so grim, but I think this is one of the more pressing worries for IRI.

6 Belief, Justification and Interest

If we decide that not everything in epistemology is interest-relative, then we face a series of questions about which things are, and are not, interest relative. One of these concerns belief. Should we say that what an agent believes is sensitive to what her interests are?

Note that the question here concerns whether belief is constitutively related to interests. It is extremely plausible that belief is causally related to interests. As Jennifer Nagel (2008) has shown, many agents will react to being in a high-stakes situation by lowering their confidence in relevant propositions. In this way, being in a high-stakes situation may cause an agent to lose beliefs. This is not the kind of constitutive interest-relativity that's at issue here, though the fact this happens makes it harder to tell whether there is such a thing as constitutive interest-relativity of belief.

I find it useful to distinguish three classes of views about beliefs and interests.

1. Beliefs are not interest-relative. If knowledge is interest-relative, the interest-relativity is in the conditions a belief must satisfy in order to count as knowledge.
2. Beliefs are interest-relative, and the interest-relativity of belief fully explains why knowledge is interest-relative.
3. Beliefs are interest-relative, but the interest-relativity of belief does not fully explain why knowledge is interest-relative.

In Weatherson (2005), I suggested an argument for option 2. I now think that argument fails, for reasons given by Jason Stanley (2005). I originally thought option 2 provided the best explanation of cases like Chika's. Assume Chika does the rational thing, and takes the blue ticket. She believes it is better to take the blue ticket. But that would be incoherent if she believed the Red Sox won. So she doesn't believe

the Red Sox won. But she did believe the Red Sox won before she was offered the bet, and she hasn't received any new evidence that they did not. So, assuming we can understand an interest-invariant notion of confidence, she is no less confident that the Red Sox won, but she no longer believes it. That's because belief is interest-invariant. And if all cases of interest-relativity are like Chika's, then they will all be cases where the interest-relativity of belief is what is ultimately explanatory.

The problem, as Stanley had in effect already pointed out, is that not all cases are like Chika's. If agents are mistaken, and especially if they are irrational, about the choice they face, this explanation won't go through. Assume Dian faces the same choice as Chika, and this is clear, but he irrationally believes that the red ticket pays out \$2. So he prefers the red ticket to the blue ticket, and there is no reason to deny he believes the Red Sox won. Yet taking the red ticket is irrational; he wouldn't do it were he rational. Yet it would be rational if he knew the Red Sox won. So Dian doesn't know the Red Sox won, in virtue of his interests, while believing they did.

Note this isn't an argument for option 1. Everything I said about Dian is consistent with the Chika-based argument for thinking that belief is interest-relative. It's just that there are cases where the interest-relativity of knowledge can't be explained by the interest-relativity of belief. So I now think option 3 is correct.

We can ask similar questions about whether justified belief is interest-relative, and whether if so this explains the interest-relativity of knowledge. I won't go into as much detail here, save to note that on my preferred version of IRI, Dian's belief that the Red Sox won is both justified and rational. This is, to put it mildly, not the most intuitive of verdicts, but I suspect the alternative verdicts lead to worse problems elsewhere. Rather than delving deeper into the details of IRI though, let's turn to some objections to the view.

7 Debunking Objections

Many arguments against IRI are, in effect, debunking arguments. The objector's immediate conclusion is not that IRI is false, but that it is unsupported by the arguments given for it.

Arguments that people do not have the intuition that, for example, Bojan lacks knowledge that his paper is typo-free, do not immediately show that IRI is false. That's because the truth of IRI can be made compatible with that intuition in two ways. For one thing, it is possible that people think Bojan knows because they think Bojan betting that his paper is typo free is, in the circumstances, a good bet.² For another thing, intuitions around here might be unreliable. Remember that one of the original motivations for IRI was that it was the lowest cost solution to the

²Compare the response to Feltz and Zarpentine (2010) that I make in Weatherson (2011, §1), or the response to Lackey (2010) by Masashi Kasaki (2014, §5).

preface paradox and lottery paradox. We shouldn't expect intuitions to be reliable in the presence of serious paradox. That consideration cuts both ways; it makes debunking objections to arguments for IRI from intuitions about cases look very promising. And I think those objections are promising; but they don't show IRI is false.

Similarly, objections to the premises of the argument from principles don't strictly entail that IRI is false. After all, IRI is an existential thesis; it says sometimes interests matter. The principles used to defend it are universal claims; they say (for example) it is always permissible to act on knowledge. Weaker versions of these principles might still be consistent with, or even supporting of, IRI. But this feels a little desperate. If the premises of these arguments fail, then IRI looks implausible.

But there are still two methodological points worth remembering. Sometimes it seems that critics of principles like K-Suff reason that K-Suff entails IRI, and IRI is antecedently implausible, so we should start out suspicious of K-Suff. Now why might IRI be antecedently implausible?

I think to some extent it is because it is thought to be so revolutionary. The denial of interest-relativity is often taken to be a "traditional" view. This phrasing appears, for example, in Boyd (forthcoming), and in Ichikawa et al. (2012), and even in the title of Buckwalter (2014). And if this were correct, that would be a mark against interest-relativity. The "inherited experience and acumen of many generations of men" (Austin, 1956-7, 11) should not be lightly forsaken. The problem is that it isn't true; there is nothing particularly novel about IRI. As Stephen R. Grimm (2015) points out, you can see a version of the view in Locke, and in Clifford. What's really radical, as Descartes acknowledged, is to think the perspective of the Cartesian meditator is the right one for epistemology.

Perhaps what is unintuitive about IRI is that it makes knowledge depend on factors that are not 'truth-directed', or 'truth-conducive'. There is a stronger and weaker version of the principle that might be being appealed to here. The stronger version is that IRI makes practical matters into one of the factors on which knowledge depends, and this is implausible. But IRI doesn't do this. It is consistent with IRI to say that only truth-conducive features of beliefs are relevant to whether they amount to knowledge, but how much of each feature one needs depends on practical matters. The weaker principle is that IRI makes knowledge counterfactually sensitive to features irrelevant to the truth, justification or reliability of the belief. This is true, but it isn't an objection to IRI. Any theory that allows defeaters to knowledge, and defeaters to those defeaters, will make knowledge counterfactually sensitive to non-truth-conducive features in just the same way. And it is independently plausible that there are defeaters to knowledge, and they can be defeated.³

³The argument of the last two sentences is expanded on greatly in Weatherson (2014, §3). The

These are all reasons to think that IRI is not antecedently implausible. There is one reason to think it is antecedently plausible. On a functionalist theory of mind, belief is a practical notion. And it is plausible that knowledge is a kind of success condition for belief. Now it's possible to have non-practical success conditions for a state our concept of which is practical. But I don't find that a natural starting assumption. It's much more intuitive, to me at least, that the norms of belief and the metaphysics of belief would be tightly integrated. And that suggests that IRI is, if anything, a natural default.

That's not an argument for IRI, or of course for K-Suff. And there are important direct objections to K-Suff. Jessica Brown (2008) and Jennifer Lackey (2010) have examples of people in high stakes situations who they say are intuitively described as knowing something, but not being in a position to act on it. I'm sympathetic to the two-part reply that Masashi Kasaki (2014) makes to these examples. The first thing to note is that these are hard cases, in areas where several paradoxes (e.g., lottery, preface, sceptical) are lurking. Intuitions are less reliable than usual around here. But another thing to notice is that it is very hard to say what actions are justified by taking p for granted in various settings. Brown and Lackey both describe cases where doctors have lots of evidence for p , and given p a certain action would maximise patient-welfare, but where intuitively it would be wrong for the doctor to act that way. As it stands, that's a problem for IRI only if doctors should maximise epistemic expected patient-welfare, and that principle isn't true. Kasaki argues that there isn't a way to fill out Lackey's example to get around this problem, and I suspect the same is true for Brown's example.

Finally, note that K-Suff is an extensional claim. Kenneth Boyd (forthcoming) and Baron Reed (2014) object to a principle much stronger than K-Suff: the principle that what an agent knows should explain why some choices are rational for them. Both of them say that if IRI is inconsistent with the stronger principle, that is a serious problem for IRI. (In Boyd's case this is part of an argument that IRI is unmotivated; in Reed's case he takes it to be a direct objection to IRI.) Now I think IRI is inconsistent with this principle. Chika doesn't know the Red Sox won because she can't rationally choose the red ticket, not the other way around. But I don't see why the principle is so plausible. It seems plausible to me that something else (e.g., evidence) explains both rational choice and knowledge, and the way it explains both things makes IRI true.

idea that knowledge allows for defeaters is criticised by Maria Lasonen-Aarnio (2014a). Eaton and Pickvance (2015) make an objection to IRI that misses the significance of this point.

8 Direct Objections

Let's close with direct arguments against IRI. There are two kinds of arguments that I won't address here. One of these is the argument, developed in Ichikawa et al. (2012) that there isn't a good way to say how far interest-relativity should extend. As I noted above, I agree this is a deep problem, and don't think there is a good answer to it in the existing literature. The other kind are objections that only apply to the Stakes version of IRI, not the Odds version. One instance of this kind is the Dutch Book argument deployed by Baron Reed (2014). I think several instances of that kind of argument are successful. But the theory they succeed against is not IRI, but a sub-optimal version of IRI. So I'll stick to objections that apply to the Odds version.

IRI does allow knowledge to depend on some unexpected factors. But so do most contemporary theories of knowledge. Most contemporary theories allow for knowledge to be defeated in certain ways, such as by available but unaccessed evidence (Harman, 1973, 75), or by nearby possibilities of error (Goldman, 1976), or by mistakes in the background reasoning. The last category of cases aren't really contemporary; they trace back at least to Dharmottara Nagel (2014, 58). And the theories allow for defeaters to be defeated. And once we work through the details of what can defeat a defeater, it turns out many surprising things can affect knowledge.

Indeed, for just about any kind of defeater, it is possible to imagine something that in some ways makes the agent's epistemic position worse, while simultaneously defeating the defeater.⁴ If interests matter to knowledge because they matter to defeaters, as is true on my version of IRI, we should expect strange events to correlate with gaining knowledge. For example, it isn't surprising that one can gain knowledge that p at exactly the moment one's evidential support for p falls. This consequence of IRI is taken to be obviously unacceptable by Eaton and Pickvance (2015), but it's just a consequence of how defeaters generally work.

IRI has been criticised for making knowledge depend on agents not allowing agents to get knowledge by not caring, as in these vivid quotes:

Not giving a damn, however enviable in other respects, should not be knowledge-making. (Russell and Doris, 2009, 433)

If you don't know whether penguins eat fish, but want to know, you might think ... you have to gather evidence. [But if IRI] were correct, though, you have another option: You could take a drink or shoot heroin. (Cappelen and Lepore, 2006, 1044–5)

⁴The argument of the last two sentences is expanded on greatly in Weatherson (2014, §3), where it is credited to Martin Smith. The idea that knowledge allows for defeaters is criticised by Maria Lasonen-Aarnio (2014b).

Let's walk through Cappelen and Lepore's case. IRI says that there are people who both have high confidence that penguins eat fish, and they have this confidence for reasons that are appropriately connected to the fact that penguins eat fish. But one of them really worries about sceptical doubts, and so won't regard the question of what penguins eat as settled. The other brushes off excessive sceptical doubts, and rightly so; they are, after all, excessive. IRI says that the latter knows and the former does not. If the former were to care a little less, in particular if they cared a little less about evil demons and the like, they'd know. Perhaps they could get themselves to care a little less by having a drink. That doesn't sound like a bad plan; if a sceptical doubt is destroying knowledge, and there is no gain from holding on to it, then just let it go. From this perspective, Cappelen and Lepore's conclusion does not seem like a *reductio*. Excessive doubt can destroy knowledge, so people with strong, non-misleading evidence can gain knowledge by setting aside doubts. And drink can set aside doubt. So drink can lead to knowledge.⁵

But note that the drink doesn't generate the knowledge. It blocks, or defeats, something that threatens to block knowledge. We should say the same thing to Russell and Doris's objection. Not giving a damn, about scepticism for example, is not knowledge-making, but it is knowledge-causing. In general, things that cause by double prevention do not make things happen, although later things are counterfactually dependent on them (Lewis, 2004). And the same is true of not caring.

Finally, it has been argued that IRI makes knowledge unstable in a certain kind of way (Lutz, 2014; Anderson, 2015). Practical circumstances can change quickly; something can become a live choice and cease being one at a moment's notice. If knowledge is sensitive to what choices are live, then knowledge can change this quickly too. But, say the objectors, it is counterintuitive that knowledge changes this quickly.

Now I'm not sure this is counterintuitive. I think that part of what it takes to know p is to treat the question of whether p as closed. It sounds incoherent to say, *I know a is the F , but the question of who is the F is still open*. And whether a question is treated as open or closed does, I think, change quite rapidly. One can treat a question as closed, get some new reason to open it (perhaps new evidence, perhaps an interlocutor who treats it as open), and then quickly dismiss that reason. So I'm not sure this is even a problem.

But to the extent that it is, it is only a problem for a somewhat half-hearted version of IRI. The puzzles the objectors raise turn on cases where the relevant practical options change quickly. But even once a practical option has ceased to

⁵ Wright (2004) notes that there often is not value in holding on to sceptical doubts, and the considerations of this paragraph are somewhat inspired by his views. My sense is that sceptical doubts are less debilitating than heroin addiction, so don't go down that route.

be available, it can be hard in practice to dismiss it from one's mind. One may often still think about what to do if it becomes available again, or about exactly how unfortunate it is that the option went away. As long as theoretical as well as practical interests matter to knowledge, it will be unlikely that knowledge will be unstable in just this way. Practical interests may change quickly; theoretical ones typically do not.

References

- Anderson, Charity. 2015. "On the Intimate Relationship of Knowledge and Action." *Episteme* 12:343–353, doi:10.1017/epi.2015.16.
- Austin, J. L. 1956-7. "A Plea for Excuses." *Proceedings of the Aristotelian Society* 57:1–30.
- Boyd, Kenneth. forthcoming. "Pragmatic Encroachment and Epistemically Responsible Action." *Synthese*, doi:10.1007/s11229-015-0878-y.
- Brown, Jessica. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning." *Noûs* 42:167–189, doi:10.1111/j.1468-0068.2008.00677.x.
- . forthcoming. "Impurism, Practical Reasoning and the Threshold Problem." *Noûs*, doi:10.1111/nous.12008.
- Brown, Jessica and Gerken, Mikkel (eds.). 2012. *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Buckwalter, Wesley. 2014. "Non-Traditional Factors in Judgments about Knowledge." *Philosophy Compass* 7:278–289, doi:10.1111/j.1747-9991.2011.00466.x.
- Buckwalter, Wesley and Schaffer, Jonathan. 2015. "Knowledge, Stakes and Mistakes." *Noûs* 49:201–234, doi:10.1111/nous.12017.
- Cappelen, Herman and Lepore, Ernest. 2006. "Shared Content." In Ernest Lepore and Barry C. Smith (eds.), *The Oxford Handbook of Philosophy of Language*, 1020–1055. Oxford: Oxford University Press.
- DeRose, Keith. 1992. "Contextualism and Knowledge Attributions." *Philosophy and Phenomenological Research* 52:513–529s.
- . 2009. *The Case for Contextualism: Knowledge, Skepticism and Context*. Oxford: Oxford.

- Eaton, Daniel and Pickvance, Timothy. 2015. "Evidence against pragmatic encroachment." *Philosophical Studies* 172:3135–3143, doi:10.1007/s11098-015-0461-x.
- 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. 2010. "Do You Know More When It Matters Less?" *Philosophical Psychology* 23:683–706, doi:10.1080/09515089.2010.514572.
- Goldman, Alvin I. 1976. "Discrimination and Perceptual Knowledge." *The Journal of Philosophy* 73:771–791.
- Grimm, Stephen R. 2015. "Knowledge, Practical Interests and Rising Tides." In David K. Henderson and John Greco (eds.), *Epistemic Evaluation: Purposeful Epistemology*, 117–137. Oxford: Oxford University Press.
- Harman, Gilbert. 1973. *Thought*. Princeton: Princeton University Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John and Stanley, Jason. 2008. "Knowledge and Action." *Journal of Philosophy* 105:571–90.
- Ichikawa, Jonathan Jenkins, Jarvis, Benjamin, and Rubin, Katherine. 2012. "Pragmatic Encroachment and Belief-Desire Psychology." *Analytic Philosophy* 53:327–343, doi:10.1111/j.2153-960X.2012.00564.x.
- Kasaki, Masashi. 2014. "Subject-Sensitive Invariantism and Isolated Secondhand Knowledge." *Acta Analytica* 29:83–98, doi:10.1007/s12136-013-0215-3.
- Kim, Brian. forthcoming. "In Defense of Subject-Sensitive Invariantism." *Episteme*, doi:10.1017/epi.2015.40.
- Lackey, Jennifer. 2010. "Acting on Knowledge." *Philosophical Perspectives* 24:361–382.
- Lasonen-Aarnio, Maria. 2014a. "The Dogmatism Puzzle." *Australasian Journal of Philosophy* 92:417–432, doi:10.1080/00048402.2013.834949.
- . 2014b. "Higher-Order Evidence and the Limits of Defeat." *Philosophy and Phenomenological Research* 88:314–345.

- Lewis, David. 2004. "Causation as Influence." In John Collins, Ned Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*, 75–106. Cambridge: MIT Press.
- Lutz, Matt. 2014. "The Pragmatics of Pragmatic Encroachment." *Synthese* 191:1717–1740, doi:10.1007/s11229-013-0361-6.
- Nagel, Jennifer. 2008. "Knowledge ascriptions and the psychological consequences of changing stakes." *Australasian Journal of Philosophy* 86:279–294, doi:10.1080/00048400801886397.
- . 2014. *Knowledge: A Very Short Introduction*. Oxford: Oxford University Press.
- Pinillos, Ángel. 2012. "Knowledge, Experiments and Practical Interests." In Brown and Gerken (2012), 192–219.
- Reed, Baron. 2014. "Practical Matters Do Not Affect Whether You Know." In Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, 95–106. Chichester: Wiley-Blackwell, 2nd edition.
- Rubin, Katherine. 2015. "Total Pragmatic Encroachment and Epistemic Permissiveness." *Pacific Philosophical Quarterly* 96:12–38, doi:10.1111/papq.12060.
- Russell, Gillian and Doris, John M. 2009. "Knowledge by Indifference." *Australasian Journal of Philosophy* 86:429–437, doi:10.1080/00048400802001996.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.
- Weatherson, Brian. 2005. "Can We Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19:417–443, doi:10.1111/j.1520-8583.2005.00068.x.
- . 2011. "Defending Interest-Relative Invariantism." *Logos & Episteme* 2:591–609.
- . 2012. "Knowledge, Bets and Interests." In Brown and Gerken (2012), 75–103.
- . 2014. "Probability and Scepticism." In Dylan Dodd and Elia Zardini (eds.), *Scepticism and Perceptual Justification*, 71–86. Oxford: Oxford University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- Wright, Crispin. 2004. "Warrant for Nothing (And Foundations for Free)?" *Proceedings of the Aristotelian Society, Supplementary Volume* 78:167–212.