

Pragmatic Encroachment

Jeremy Fantl and Matthew McGrath

In his classic (1953) article, Richard Rudner claims that

in accepting a hypothesis the scientist must make the decision that the evidence is *sufficiently* strong or that the probability is *sufficiently* high to warrant the acceptance of the hypothesis. Obviously, our decision regarding the evidence and respecting how strong is ‘strong enough’, is going to be a function of the *importance*, in the typically ethical sense, of making a mistake in accepting or rejecting the hypothesis... *How sure we need to be before we accept a hypothesis will depend on how serious a mistake would be.* (2, emphasis Rudner’s)

According to Rudner, an adequate account of the conditions of warranted hypothesis acceptance must include reference to an ethical or more broadly a pragmatic factor.

Rudner explicitly confines his discussion to the evidence or probability needed to be warranted in accepting a hypothesis, where acceptance for him seems to be subject to voluntary control, at least in certain cases: we *decide* the evidence is sufficiently strong. But in the past decade a number of philosophers have offered views similar to Rudner’s about a broader range of epistemic concepts. For example, regarding knowledge, Jeremy Fantl and Matthew McGrath ((2002) and (2007)), John Hawthorne (2004), and Jason Stanley (2005)¹ have recommended views according to which, whether a subject knows something to be the case depends on their practical situation.

This sort of conclusion shouldn’t strike us as immediately implausible. After all, there are widely acknowledged links between the practical and the epistemic. What you should do in a

¹ See also Hawthorne and Stanley’s co-authored (forthcoming).

certain choice situation is widely acknowledged to be determined not merely by how good or preferred certain possible outcomes of your acts would be but how probable these outcomes are given that you act in certain ways. Or consider again the importance of not being wrong. This seems obviously relevant to whether you should inquire further into whether p and to whether your evidence is strong enough to justify you in ignoring the chance that p is false. If your practical situation can be relevant to these matters, it might not seem to be a huge step to think it can be relevant to whether you know.

Some of our intuitions about specific cases seem to support the claim that knowledge can depend on practical factors. Consider DeRose's famous (1992) "Bank Cases":

Bank Case A (Low Stakes). My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." I reply, "No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon."

Bank Case B (High Stakes). My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a *very* bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I'd better go in and make sure." (913)

It looks like Keith speaks truly in Case A in attributing knowledge to himself that the bank will be open tomorrow, while he also speaks truly in Case B in denying himself knowledge. The only thing that changes in the two cases is how important it is for Keith to be right about whether the bank will be open tomorrow. Therefore, it looks like how important it is for Keith to

be right about whether the bank will be open tomorrow is relevant to whether Keith knows that the bank will be open tomorrow. And relevant in a clear way: holding fixed Keith's evidence concerning whether the bank will be open tomorrow, whether he knows it will be open varies with variations in how important it is for him to be right about this.

But here we find some odd consequences. If this is the proper lesson to draw from the Bank Cases, it would appear to follow that two subjects can have the same evidence concerning whether the bank will be open tomorrow, even though one of them knows it'll open tomorrow and the other doesn't. And this will be just because one subject believes the bank will be open tomorrow and the other doesn't.² Nor because the bank will be open tomorrow and false in the other, nor because one subject is Gettiered and the other not. What makes the difference in knowledge has nothing to do with these traditional factors. In fact, one subject might have *more* evidence than another that the bank will be open tomorrow – be better informed, have done more checking, etc. – but because much more is at stake for the more well-informed subject, the more well-informed subject can fail to know that the bank will be open tomorrow while the less-informed subject knows that the bank will be open tomorrow. All this is hard to swallow.

DeRose himself draws a very different moral from the cases. He grants that Keith speaks truly in Case A when he says "I know it'll be open" and he also speaks truly in Case B when he responds "no" to the question, "Do you know that the bank will be open tomorrow?" But he

² Kent Bach questions whether Keith is allowed to stipulate that he remain as confident even once the stakes change, and also questions whether, even if he retains the same degree of belief – the same credence – between the two cases, his credence remains sufficient for outright belief. His worries are shared by Jennifer Nagel, though she ends up granting the possibility that Keith retains belief between the cases. If one thinks Keith can't sincerely say "I better go in and make sure" if he has a belief that the bank will be open tomorrow, then we can alter the case by stipulating that Keith says, "I know it's open tomorrow." We can further emend the case by having Keith's wife say, "No you don't. This is important, Keith. You need to go in and make sure they're open tomorrow." The question would then be whether what Keith said is false. We see no great gulf between the original case B and this altered one.

insists that it does not follow that the bank cases show that pragmatic factors can make a difference to knowledge. Why not? DeRose's answer, to put it roughly, is that Keith's pragmatic situation affects what he means by 'know' but makes no difference to whether he knows. Put more precisely, his answer has two parts. First, there is a *contextualist* claim: which epistemic relation Keith's use of 'know' picks out varies with his pragmatic situation, and so varies across the bank cases. If the epistemic relations picked out vary with the pragmatic situation – the higher the stakes, roughly, the more demanding the epistemic relation picked out – what Keith says in case A might be true while what he said in case B might be false. That is the first part – the claim that pragmatic factors matter to the content of knowledge-attributions. The second part is the denial that they matter to knowledge. Or to put it more precisely: for each epistemic relation K that can be picked out by 'know' in some context, pragmatic factors make no difference to whether you stand in K to a proposition. If this is true of all such relations, then the sentence 'Pragmatic factors make no difference to whether one knows' comes out invariantly true.

What is important for our purposes is that the contextualist finds a way to salvage the intuitions in the cases while avoiding some of the problematic consequences of letting knowledge involve the subject's practical situation. Thus, an important motivation for contextualism, it might be thought, is the avoidance of those problematic consequences.³

But contextualism comes with problems of its own; for one thing, the problematic consequences remain, but bumped up a level. The contextualist, it seems, must grant that two subjects might have the same evidence that the bank will be open tomorrow, but one could truly self-attribute knowledge while the other truly self-denies knowledge. One subject could have

³ This point is emphasized by DeRose in his (forthcoming).

more evidence than the other, be better informed about whether the bank will be open tomorrow, have done more checking into whether it will be open tomorrow, but the less well-informed subject might be able to truly self-attribute knowledge while the better-informed subject cannot. In addition, there are independent difficulties. For example, Stanley (2005) takes the contextualist about knowledge attributions to be positing a sort of context-sensitivity which departs significantly from the familiar sorts.

So, perhaps we should deny contextualism while maintaining our commitment to the intuitions in the Bank Cases – our commitment to the view that whether a subject knows something depends on the subject’s practical situation. This is the *subject-sensitive invariantism* of Hawthorne and Stanley.⁴ It is an *invariantism* because it holds that the proposition expressed by a knowledge-attribution does not vary from attributor context to attributor context. But it is *subject-sensitive* because it holds that whether a subject knows something is sensitive to the practical situation of the subject.

Note, though, that we are not forced to a denial of contextualism simply by an acceptance of subject-sensitivity. If the primary motivation for contextualism is the preservation of the view that a subject’s practical situation must be irrelevant to whether that subject knows, then there will be some pressure, upon accepting subject-sensitivity, to reject contextualism. But there might be other pressures in favor of contextualism. For example, as DeRose (2004) and Stanley both point out, subject-sensitive invariantists have difficulty accommodating third-person cases, in which a High-Stakes subject denies knowledge to a Low-Stakes subject, even though each subject has the same evidence – evidence, when we are thinking about the Low-Stakes subject alone, we are tempted to think is good enough to give the Low-Stakes subject knowledge.

⁴ Hawthorne calls his view ‘sensitive moderate invariantism’ and Stanley calls his ‘interest-relative invariantism’. The label ‘subject sensitive invariantism’ is due to DeRose (2004).

(Cohen's (1999) airport case is a good example.) We can adopt contextualism to handle these cases, while still accepting subject-sensitivity if we think it is the most natural way to handle cases like the Bank Cases, or if there are independent arguments for subject-sensitivity – for the claim that your practical situation is relevant to whether you know.⁵

One might well ask under what conditions the *relevance* of the practical or pragmatic becomes the fearsome *encroachment* in Jonathan Kvanvig's label, "pragmatic encroachment". So far, we have treated 'pragmatic' as equivalent to 'practical'. But there is a broader use of the former term, to cover not only factors having to do with action and preference but also features of a speech context that have been standardly taken not to be relevant to the content of what is said but only to the appropriateness of saying it, e.g., the salience of error-possibilities. Since John Hawthorne tentatively proposes that salience of possible error is relevant to knowledge, let us use 'pragmatic' in the broader sense.

We can distinguish at least two grades of pragmatic relevance. But it is clear that only one of them deserves the label "pragmatic encroachment." We focus on knowledge, but a similar account could be given for any epistemic feature:

First grade: Knowledge has a pragmatic condition. That is, there is some interesting⁶ true principle of one or both of the following forms, where PC is pragmatic:

in order to know that *p*, you must satisfy condition PC.

⁵ Contextualism might also provide a way of dealing with some of the counterintuitive consequences we discussed above. If the attributor A has a certain high stakes practical situation in mind, perhaps that tends to affect the content of 'knows' in A's context, so that an attribution of knowledge can be true in her context only if the subject meets suitably high standards. Thus, it is hard to truly utter the likes of 'S1 and S2 have the same evidence concerning whether *p* but only S1 knows, because S2 is in a high stakes situation'.

⁶ We add the qualification that the principle should be 'interesting' to rule out the truth of principles such as 'In order to use your knowledge that *p* in practical reasoning, you must know that *p*' as entailing any grade of pragmatic relevance.

in order to satisfy condition PC, you must know that p .

Thus, examples of such a condition would include: *being rational to act as if p* (Fantl & McGrath 2002), *being rational to ignore in action the chance that p is false* (Fantl & McGrath forthcoming), *being appropriate to use p as a premise in practical reasoning* (Hawthorne 2004, Stanley 2005), *being such that no counterpossibilities of error are salient* (Hawthorne 2004).

Notice that Rudner wants to go further than this, with respect to hypothesis acceptability. For him, it's not merely that in order to be warranted in accepting a hypothesis, you must satisfy some pragmatic condition, though that is true; it's that whether a hypothesis is acceptable for you can vary with mere variations in a pragmatic factor, in particular with mere variations in how important it is to be right about the hypothesis. This suggests a second, deeper, grade of pragmatic relevance:

Second grade: Knowledge can vary with mere variation of pragmatic factors. That is, there are cases of knowledge such that if we merely vary a pragmatic factor present in that case, and leave everything else the same (as much as is possible), we can arrive at a case of ignorance.

The bank cases support this thesis: we have merely varied the importance for Keith of being right about whether the bank will be open tomorrow, leaving everything else, as much as possible, the same, and with this pragmatic variation comes a variation in knowledge.

The second grade is certainly more controversial than the first. Suppose there is pragmatic relevance only of the first grade, and not the second. Then we would not have to cope with the peculiarity of statements such as 'I know that p but I wouldn't know that p if more were riding on whether p '. One way to accept the first grade without the second is to accept the claim that knowledge that p requires epistemic certainty, which in turn requires the satisfaction of

pragmatic conditions, such as *being rational to ignore the chance that not- p in one's decision-making*. The obvious worry about such a view is that it makes knowledge too hard to come by.

It is clear that, of these two grades of pragmatic relevance, only the second grade should count as pragmatic encroachment on knowledge. In previous work, we have endorsed the second grade – and so endorsed pragmatic encroachment – because we think it is a simple consequence of the denial of a supervenience thesis concerning knowledge, which we have called “epistemological purism”:

Epistemological purism: two subjects alike with respect to their strength of epistemic position with respect to p are alike with respect to whether they know that p (or at least with respect to whether they are in a position to know that p).

It is a tricky matter to specify precisely what strength of epistemic position involves. We take it that your strength of epistemic position with respect to p is determined by your standing on truth-relevant dimensions with respect to p , including reliability, strength of evidence, epistemic probability, and the like. Truth-relevant dimensions, intuitively, can be thought of as dimensions a higher standing on which with respect to p places you in a better position with respect to the truth of p . So, how important it is for your life that you are right about whether p doesn't count as a truth-relevant dimension, nor does your standing on it seem to be fixed by your standing on truth-relevant dimensions with respect to p .

We will say more about the grounds on which we reject this thesis later on. For the moment, notice that one could accept the second grade without denying epistemological purism, as Stanley (2005) does. Stanley argues that variations in knowledge that p , due to pragmatic factors, give rise to variations in your standing along the truth-relevant dimensions. So, in Bank Case A, Keith has better evidence for the proposition that the bank will be open tomorrow than

he does in Case B: in Case A, part of his evidence is *the bank will be open tomorrow*, which is superbly good evidence for *the bank will be open tomorrow*; in case B, Keith lacks this evidence. Why the difference in evidence? Because of a difference in knowledge. Stanley accepts Williamson's (2000) E=K thesis – the thesis that your total evidence is just the same as the totality of what you know. Stanley suggests the same goes for other epistemic concepts we would like to think of as picking out standings on truth-relevant dimensions. When you lose knowledge that *p*, due to changing stakes, you thereby lose probability 1 for *p*. If Stanley is right, then, on the one hand, epistemological purism is not threatened by pragmatic encroachment, but on the other, epistemological purism doesn't deserve its name, because there are no purely truth-relevant dimensions of the relevant sort that come even close to providing a supervenience base for knowledge.

One might think that Stanley's position, although it saves one from the denial of epistemological purism, is a more radical form of pragmatic encroachment than ours, which denies purism. For Stanley thinks there is no genuinely epistemic notion that is unsullied by pragmatic factors. Change the stakes and you change not only knowledge but evidence, justification, reasons, probability, etc. There is pragmatic encroachment, in Juan Comesaña's (2008) words, *all the way down*.

One nice feature Stanley's view might appear to have is that it avoids having to explain away the peculiarities such as 'A and B have the same evidence for *p*, but only A is in a position to know that *p*'.⁷ Stanley does have to explain away peculiar temporal and modal claims about

⁷ We say his view 'appears' to have this nice feature. But we doubt it does. Notice that even if E=K, and even if there is pragmatic encroachment all the way down, we can still ask whether some epistemic properties supervenes others. For example, consider the following, where evidence is non-factive with respect to *p* is evidence which does not entail *p*:

pragmatic factors making a difference to knowledge, but perhaps these aren't as counterintuitive. Nonetheless, the cost of following Stanley in explaining away these peculiarities seems greater than the cost of the peculiarities explained away. That's because your standing on many epistemic dimensions isn't plausibly affected by mere changes in practical environment. Consider probability: you are offered a high-stakes bet on the proposition this die will come up 6, that doesn't seem to lower its probability for you and it certainly does not raise the probability for you that it will come up 1-5, or any of 1-5 individually. We see no reason to think that matters are different when the probabilities approach 1, and so when it can seem plausible that, before being offered the bet, you know.

Putting Stanley's deep pragmatic encroachment aside, here is what we take to be the simplest and most convincing case for accepting pragmatic encroachment of the second grade.

Think about Keith in Bank Case A. Unless we are willing to accept the claim that knowledge requires epistemic certainty, we should say that he knows the bank will be open tomorrow (or if one doesn't like this claim because it is about the future, make any adjustments to the example you like, but preserve lack of certainty). At any rate, if knowledge doesn't require certainty, there will be some case relevantly like Bank Case A in which the subject has knowledge. So, Keith knows in Bank Case A that the bank will be open tomorrow. Keith also knows that if it will be open tomorrow, the option with the best outcome is to wait until tomorrow. Given how very well Keith knows this conditional, Keith also knows that waiting until tomorrow will have the best outcome of his available options. And this is what he is justified in doing.

Two subjects alike with respect to their non-factive evidence regarding p are alike with respect to whether they are justified in believing that p). It is not clear to us how Stanley's acceptance of "pragmatic encroachment all the way down" enables him to avoid denying this antecedently plausible claim.

Now consider Keith in Bank Case B. Assume, for reductio, that the second grade of pragmatic encroachment fails: knowledge can't vary with mere changes of pragmatic factors. Then, since Keith knows the bank will be open tomorrow in Case A, he knows it in Case B. But he also seems to know that if it will be open tomorrow the available option with the best outcome is to wait until tomorrow to cash the check. Knowing this so very well – just as well as in Case A – Keith will also know that the available option in Case B with the best outcome is to wait until tomorrow. Now ask yourself: if you know that of all your available options, option O will have the best outcome, what should you do? It is hard not to answer: *well, O!* That O will have the best outcome of all your options seems to be a decisive reason to do O, not anything else. Think of what you will know about the other options: that they will have worse outcomes than O. So, if Keith knows that waiting until tomorrow to deposit the check will have the best outcome, he will know that going in to check further will have a worse outcome. Should he do what he knows to be worse than O? That sounds absurd. (Remember here: it's not just that he knows that going in to check further is worse than *some* other actions; he knows which action it is worse than.)

Given that he can't be certain that the bank will be open tomorrow, then assuming the stakes are high enough, and that it is important enough to be right, shouldn't Keith play it safe? Shouldn't he go in and check further, and if necessary, wait in line to deposit the check today? This seems like sound advice, and is likely the advice most decision-theorists will give, at least assuming the stakes are high enough. So, it seems in Case B, Keith isn't justified in waiting until tomorrow to deposit the check; he should rather go in and check further. But we saw that it was

absurd to think one should take an option one knows to be worse than a given option. So, our assumption that Keith knows that the bank will be open tomorrow in Case B must be false.⁸

What is crucial to this argument are not any intuitions specific to the Bank Cases, but rather three general claims: 1) the fallibilist thesis that knowledge doesn't require certainty; 2) the assumption that if a subject lacks certainty concerning a proposition relevant to the question of what to do the lack of certainty can make a difference as to what the subject is justified in doing, if the stakes are high enough, and 3) the assumption that if you know that an option O will have the best outcome of all your available acts then you are justified in doing O. Given these claims, we can show that there is a pair of cases, like the Bank Cases, in which whether a subject knows something varies with mere variations in pragmatic factors.

We take it that assumption 2 is pretty safe. Perhaps in the end 1 should go, though there are serious questions about the skeptical implications of giving it up. There are also hard questions, which we are ignoring here, about how precisely to understand the sort of certainty involved (clearly it is epistemic rather than psychological⁹). Assumption 3 is doing a lot of the work here. One might hope to deny it while explaining away the apparent absurdity of 'I know O will be better than P but I should do P' in terms of a Gricean implicature. We think the prospects for this gambit are poor, but we cannot discuss the matter further here.¹⁰

⁸ Our argument against what we called evidentialism in our 2002 takes this roughly this form (using our train cases). See also our argument against epistemological purism in our 2007. We should stress: the argument given does not pump intuitions about the bank cases in particular. All we need is some case of knowledge without certainty, in which what is known is not irrelevant to the question of what to do.

⁹ In our forthcoming, we argue the relevant notion of certainty should be formulated in terms of epistemic probability: to be epistemically certain that p is for p to have epistemic probability 1 for you.

¹⁰ We discuss the matter further in our (2007).

The most serious worry about 3, to our mind, is whether it is just an isolated intuition, or instead can be fitted into a plausible account of the relation between knowledge and what knowledge can justify. So, let's ask: if 3 is right, why might it be? The natural answer is that the proposition that option O will have the best consequences of all your available acts is a decisive reason for doing O. Of course, not all reasons that there are for doing something will justify you in doing it. That there is petrol in the glass is an excellent reason for refraining from drinking what's in the glass. But it only justifies you doing so if you *have* that reason.

What epistemic relation do you have to bear on a reason in order to have it? Our claim is that knowledge of a reason satisfies a sufficient condition on the epistemic relation you have to bear to a reason in order to have it. If you know that r , and r is a good reason for ϕ -ing, then r is a good reason you have for ϕ -ing.

Surely, this is the case when ϕ -ing is restricted to believing: if you know that r , and r is a good reason for believing q , then you have that good reason for believing q . Suppose you know that Clinton has already been U.S. president twice and that U.S. presidents can only serve two terms. This is a good reason for believing that Clinton won't be president again. Then you have that good reason for believing that Clinton won't be president again. But knowledge of r satisfies an epistemic condition sufficient for having r , even when what r is a reason for is not restricted to belief.

When trying to determine what is true – that is, in forming beliefs – we draw conclusions from the reasons we have. The same goes for trying to decide what to do. Here, too, we draw conclusions about what to do – we form intentions – from the reasons we have. We bring reasons into our reasoning knowing that we might draw all sorts of conclusions from them along

the way, some practical and some theoretical.¹¹ Suppose your sister calls you on the phone to tell you about plans for her upcoming visit to see you. She tells you, and you thereby come to know, that she'll be arriving at the airport at 8:00 am and will need a ride to your place. You might well include this proposition in your reasoning and at some point draw a practical conclusion from it, e.g. 'I'll be there a little after 8 am with my car', but you might also draw along the way any number of theoretical conclusions as well, e.g. that she'll be ready to be picked up a little after 8 am, that she'll be tired when she arrives, that you'll not be able to drop the kids off at preschool, and so on. The bottom line is that we don't segregate reasons by whether they are available for drawing practical or for theoretical conclusions. But if knowledge that r gave us r as a reason for forming beliefs but not for performing actions, we'd expect some degree of segregation, if not always, at least when something significant is at stake.

This is not what we find. Even when the stakes are high there is no segregation; rather, when the stakes are high, we are more careful about drawing theoretical conclusions – as careful as we are about drawing practical conclusions. Do you walk across or walk around the frozen pond? Walking around will take a while, but you don't want to fall through the ice.¹² How do you decide? Presumably, your decision will depend on whether you think the ice is thick enough to hold you. So, you'll start trying to make your decision about what to do by trying to figure out whether the ice is thick enough. Suppose you do some checking and on the basis of your information, you come to know that the ice is thick enough. So *the ice is thick enough* becomes

¹¹ Here it is not crucial precisely what a practical conclusion is. Whether it is an action, an intention, or plan, or even an 'ought' judgment, still, we draw practical conclusions from the same premises from which we draw theoretical conclusions.

¹² Thanks to Mark Migotti for suggesting this example.

a reason you have to draw theoretical conclusions on its basis (e.g., that it would be perfectly safe to cross it). It would then be very odd not allow this knowledge into your practical reasoning. Why did you try to figure out whether the ice was thick enough in the first place? For fun? No, you tried to figure it out because you were (correctly) under the impression that figuring out whether the ice was thick enough would help you decide what to do. When r becomes available as a basis for theoretical conclusions, it is 'barmy' (to use an expression suggested by one of our informants) to just ignore p in one's decision-making and planning.

So, it looks like whatever we treat as epistemic status sufficient to have r for use in reasoning – as a reason – for beliefs, we also treat as sufficient to have r for use in reasoning – as a reason – for intention and action. Similar remarks can also be applied to emotional states, desires, hopes, and reactive attitudes generally (here the emphasis on reasoning would have to be toned down with regard to some of these states, though there clearly is such a thing as responding to reasons in these cases). These observations support the claim that if knowledge that r qualifies r to be a reason you have for belief, it qualifies r to be a reason you have for any ϕ .

Clearly, that option O will have the best results of all your available actions is a reason for taking option O. Therefore, when you know this reason to be true, it's a reason you have for taking option O. If it's a reason you have for taking option O, what could stand in the way of it justifying you in taking option O? Well, perhaps it's defeated by some contrary reason – e.g. that there's a chance that option O is not the best option and, if it's not, the consequences will be disastrous. This reason might be true consistent with your knowing that option O is the best option, provided that knowledge does not require certainty. And if it is, perhaps it is natural to reason this way:

Yes, I do know that waiting until tomorrow to go to the bank is best. Of course, if it's not best, then that means that the bank isn't open tomorrow, and so if I do wait until tomorrow, the consequences will be disastrous. And there's a chance that waiting until tomorrow isn't best. It's just too risky. I'll wait on line now.

Does this reflect the way we reason? In reasoning, do we find ourselves weighing p – a reason for ϕ -ing – against *there is a serious risk that not- p* – a reason for not ϕ -ing? Not plausibly. For here's what the weighing of competing reasons feels like in uncontroversial cases:

Ice cream tastes good, and that's a reason I have to eat it, but it also is unhealthy, so that's a reason I have not to. Which is more important, taste or health?

The rain will make me wet, so that's a reason I have to bring my umbrella, but the umbrella's also really heavy, so that's a reason I have not to. Which is more important, staying dry or being unencumbered?

His paper didn't have an argument, so that's a reason I have to give him a bad grade, but also he did work really hard, so that's a reason I have to give him a good grade. Which is more important, quality of work or quality of effort?

Contrast these examples with:

There's a serious risk that waiting until tomorrow isn't best, so that's a reason I have to wait on line now. But waiting until tomorrow is also best, so that's a reason I have to wait until tomorrow. Which is more important, the serious risk that waiting until tomorrow isn't best, or the fact that waiting until tomorrow is best?

People don't weigh these kinds of reasons in the way we'd expect them to if people could have both of them at once. We'd expect to find people explicitly weighing up reasons concerning actual results against conflicting reasons concerning expected results. We find no such thing. People do vacillate: 'The ice is very thick. Surely it will hold me. But . . . there's a real possibility it won't. I better not risk it.' Perhaps even with the right halting tone of voice someone might say, 'The ice will hold me (won't it? *surely* it will, right?). Forget it. I'll play it safe and walk around.' What you don't find is the likes of, 'Hmm, the ice might not hold me. That's one consideration. Another is that it will hold me.'

So, if you have a reason, r , for ϕ -ing, it seems that the reason can't be defeated by some further reason to the effect that r might be false. And this is for the simple reason that we never have both of these reasons at the same time. Of course, even if we did have *O is best* and *O might not be best* at the same time as reasons for doing contrary things, it's not at all clear that *O might not be best* would be the winning reason. We care about actual results, not expected results, and if we have *O is best* that seems like it should beat out *O might not be best*, regardless of the consequences. But, the fact of the matter is that we don't ever have both reasons at once, and this is reflected in the fact that it is absurd to picture us reasoning by weighing the fact that p against the possibility that not- p .

So, if you know that O is best, then that's a reason you have to do O . And if that's a reason you have to do O , then it can't be defeated by considerations about the chance that O isn't best or any other epistemic considerations with respect to the reason that O is best. And it doesn't seem like, in this case, any other obstacles could stand in the way of your being justified in doing O . Therefore, if you know that O is best, you are justified in doing O . And this is just premise 3 in the above argument.

Premise 3, then, is not just an arbitrary premise. It is grounded in deep principles about knowledge and the having of reasons, in particular:

- A) If you know that r , then if r is a reason for ϕ -ing, r is a reason you have for ϕ -ing.
- B) If r is a reason you have for ϕ -ing, then r can't be defeated by any consideration to the effect that r might be false.

Suppose you know that O is best. That O is best is a reason to do O . By A) O is best is a reason you have to do O . By B), this reason isn't defeated by any consideration to the effect that O might not be best. The only plausible obstacles to O is best being a justifying reason you have to do O are considerations about the chance that O might not be best. So, that O is best is a justifying reason you have to do O . If you have a justifying reason to do something then you are justified in doing it. Thus, if you know that O is best, you are justified in doing O .

Once knowledge doesn't require certainty, then, commitment to the second grade of pragmatic relevance – pragmatic encroachment – doesn't just get motivated by our intuitive reactions to hypothetical cases. It is grounded in deep principles about knowledge and reasons.

Bibliography

- Cohen, Stewart. (1999). 'Contextualism, Skepticism, and the Structure of Reasons'. In Tomberlin, J. (Ed.) *Philosophical Perspectives*. Cambridge: Blackwell: 57-89.
- Comesaña, Juan. (2008). 'Epistemic Pragmatism: An Argument Against Moderation'. *Midwest Epistemology Workshop*. University of Nebraska-Lincoln.
- DeRose, Keith. (1992). 'Contextualism and Knowledge Attributions'. *Philosophy and Phenomenological Research* 52: 913-29.
- (2004). 'The Problem with Subject-Sensitive Invariantism'. *Philosophy and Phenomenological Research* 68: 346-50.
- (forthcoming). *The Case for Contextualism*. Oxford: Oxford University Press.
- Fantl, Jeremy & McGrath, Matthew. (2002). 'Evidence, Pragmatics, and Justification'. *The Philosophical Review* 111: 67-94.
- (2007). 'On Pragmatic Encroachment in Epistemology'. *Philosophy and Phenomenological Research* 75: 558-89.
- Hawthorne, John. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John & Stanley, Jason. (forthcoming). 'Knowledge and Action'. *Journal of Philosophy*.

Rudner, Richard. (1953). 'The Scientist *Qua* Scientist Makes Value Judgments'. *Philosophy of Science* 20: 1-6.

Stanley, Jason. (2005). 'Knowledge and Practical Interests'. Oxford, Oxford University Press.

Williamson, Timothy. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.