

Episteme

<http://journals.cambridge.org/EPI>

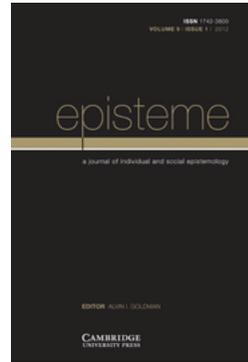
Additional services for *Episteme*:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)



PRAGMATIC ENCROACHMENT: IT'S NOT JUST ABOUT KNOWLEDGE

Jeremy Fantl and Matthew McGrath

Episteme / Volume 9 / Issue 01 / March 2012, pp 27 - 42

DOI: 10.1017/epi.2011.3, Published online: 04 April 2012

Link to this article: http://journals.cambridge.org/abstract_S1742360011000037

How to cite this article:

Jeremy Fantl and Matthew McGrath (2012). PRAGMATIC ENCROACHMENT: IT'S NOT JUST ABOUT KNOWLEDGE. Episteme, 9, pp 27-42 doi:10.1017/epi.2011.3

Request Permissions : [Click here](#)

PRAGMATIC ENCROACHMENT: IT'S NOT JUST ABOUT KNOWLEDGE

JEREMY FANTL AND MATTHEW MCGRATH

fantl@ucalgary.ca
mcgrathma@missouri.edu

ABSTRACT

There is pragmatic encroachment on some epistemic status just in case whether a proposition has that status for a subject depends not only on the subject's epistemic position with respect to the proposition, but also on features of the subject's non-epistemic, practical environment. Discussions of pragmatic encroachment usually focus on knowledge. Here we argue that, barring infallibilism, there is pragmatic encroachment on what is arguably a more fundamental epistemic status – the status a proposition has when it is warranted enough to be a reason one has for believing other things.

Sometimes we believe things for reasons. That the car's lights won't turn on can be the reason you believe that the car's battery is dead. Without engaging in any discussion of a plausible metaphysics of reasons, let us assume that in many cases what reasons we have for our beliefs are propositions. Perhaps these propositions have to be truths or facts. But the question we're ultimately interested in is independent of these issues.

What does it take for a proposition to be a reason you have to believe something? Minimally, the proposition has to be connected in the right way to the belief. That you're now reading a paper on pragmatic encroachment is not a reason you have to believe that your child's birthday is in December. It is a reason you have to believe that you're reading a paper that will cite work by Jason Stanley, because that you're now reading a paper on pragmatic encroachment is connected in the right way to the proposition that you're now reading a paper that will cite work by Jason Stanley. But it takes more than a good connection between a proposition and a belief for the proposition to be a reason you have for the belief. That you're now reading a paper on intellectualism about knowledge-how is connected in the right way to the proposition that you're now reading a paper that will cite work by Jason Stanley, but it is not a reason you have to believe you're now reading a paper that will cite work by Jason Stanley.

It's clear enough *roughly* why. It has something to do with your epistemic relations to the proposition that you're now reading a paper on intellectualism about knowledge-how. Some might think the problem is that the proposition is false – you're not, of course, now reading a paper on intellectualism about knowledge-how – and that prevents it from being your reason. Others might think the problem is instead a lack of justification,¹ or a lack of knowledge.²

1 So we argue in Fantl and McGrath 2009.

2 So argue John Hawthorne and Jason Stanley in Hawthorne and Stanley 2008.

Some terminology is helpful here. Let us say that a proposition p is ‘warranted enough’ to be a reason you have to believe a proposition q iff no weaknesses or shortcomings in your epistemic relationships to p stand in the way of p ’s being a reason you have to believe q . Sometimes we will use convenient shorthand, though, and just speak of propositions as being warranted enough to be a reason for belief, where to be warranted enough to be a reason for belief is just to be warranted enough to be a reason to believe something. Notice that being warranted enough to be a reason to believe q doesn’t entail being a reason to believe q : again, that you’re now reading a paper on pragmatic encroachment is not a reason you have to believe that your child’s birthday is in December, but it is warranted enough.

It is an important question for epistemology under what conditions a proposition has the status of being warranted enough for a person to be a reason for belief. It is this status which is at issue when we ask questions such as, ‘When do we get to count on, rely on, or base beliefs on other propositions?’ and ‘What level of evidence do we need before treating it as settled that p in our inquiries?’ At the very least, we’d like epistemology to yield results that bear on questions like, ‘Can I base belief that q on my belief that p , or do I need more evidence for p first?’ To abandon these kinds of questions, it seems to us, is to abandon a significant part of traditional epistemology.

Our aim here is not to argue for some particular answer to the question of what epistemic relation – knowledge, true belief, justification, etc. – a person must bear to a proposition for it to be warranted enough to be a reason for belief. Rather, we argue that, given facts about the structure of reasons, a surprising conclusion follows concerning warrant sufficient for being a reason one has for belief: either such warrant requires epistemic certainty or it is subject to pragmatic encroachment. This paper, if successful, thus shows why pragmatic encroachment ‘isn’t just about knowledge’. Our conclusions here, together with certain assumptions about knowledge (in particular the assumptions, (a) that knowledge is sufficient for the status of being warranted enough to be a reason for belief, and (b), that knowledge doesn’t require epistemic certainty) guarantee pragmatic encroachment on knowledge. But the action in this paper will not be with knowledge. One might think knowledge goes one way while reasons go another, and we have nothing to say against that possibility here. Our conclusion here is just that if to have a reason for belief doesn’t require certainty, then there is pragmatic encroachment on being warranted enough to be a reason one has for belief.

We will not attempt to define the notion of epistemic certainty, but we assume it has two features. First, if a proposition is epistemically certain for a person, then the person is justified in being psychologically certain of it, in having a credence or confidence of 1 in it. We also assume that epistemic certainty has implications for rational preference between bets: if a proposition is epistemically certain for you then it is rational to gamble on it at any odds and no matter what the stakes.

Epistemic certainty is hard to come by. John Hawthorne (2004: 29) says that he wouldn’t take a gamble on the law of noncontradiction at any odds and counts himself rational on that score. If he’s right and if his claim generalizes to many other people, then not even the law of noncontradiction counts as certain for many people. But even if he’s wrong about this, many other things we ordinarily take ourselves to know aren’t certain in this sense. For example, we take ourselves (at the time of this writing) to know that it’s Tuesday, but there are possible situations in which the costs and benefits of various courses of action would render it irrational to take a gamble on its being

Tuesday. If the world ends if we bet on it being Tuesday and we're wrong, but we only get a euro if we're right, then we're not rational to take that bet. Or so it plausibly seems. And if any probability lower than probability 1 entails that there is a possible situation in which we shouldn't take a gamble on the relevant proposition, then we'll also fail to be certain of such matters as the day of our birth, whether there will be at least one regulation goal scored in the next full NHL season, and (some might argue) even that here's a hand.

So, even if there are propositions that are epistemically certain for us, most of the propositions we bother to say that we know or bother to appeal to as reasons will not be epistemically certain. Therefore, if a proposition must be certain for you if it's to be warranted enough to be a reason you have for belief, the things you appeal to as reasons for belief aren't reasons after all, because they lack sufficient warrant. This shows there is a dramatic skeptical cost to the view that epistemic certainty is required for a proposition to be warranted enough to be a reason you have for belief; it's strong pressure to accept what we will call *fallibilism* about the warrant required to be a reason for belief:

(Fallibilism) A proposition can be warranted enough to be a reason you have for belief even if that proposition isn't epistemically certain for you.

This paper argues that, if fallibilism is true, then there is pragmatic encroachment on the warrant required for a proposition to be a reason you have for belief. There is pragmatic encroachment on some epistemic status just in case whether a proposition has that status for you is not a function solely of your strength of epistemic position with respect to that proposition, where your strength of epistemic position with respect to p is a function of your position on various purely truth-relevant dimensions, for example, perhaps, how probable p is for you, how reliable your relevant belief-forming processes are with respect to p , your evidence for and against p , and even whether p is true. It is not a function of your plans, goals, likes, dislikes, or even your moral duties. On the contrary, if there is pragmatic encroachment on an epistemic status, then whether a proposition has that status for you is at least partly a function of non-epistemic factors – say, what the costs and benefits of various actions are, depending on whether the proposition is true. In other work we have put similar claims as denials of supervenience theses about knowledge and justification, but here our focus is on the warrant required to be a reason for belief. There is pragmatic encroachment on that status just in case the following supervenience thesis – what we will, following our earlier vocabulary, call Purism – is false:

(Purism) Necessarily, if you and I have the same strength of epistemic position with respect to a proposition, then either that proposition is warranted enough to be a reason each of us has to believe something, or is not warranted enough to be a reason either of us has to believe anything.

The cost of accepting fallibilism is accepting pragmatic encroachment. You and I might stand in equally strong epistemic positions with respect to a proposition, but only I can use that proposition as a reason for any further beliefs. We'll build up the argument for this slowly. It will be helpful to have a case to refer to. Suppose it's 1:00 in the morning and you're nestled comfortably in bed after having made some rather sweeping changes to a paper you've been working on that day. The changes would be very difficult to reproduce, and you need the paper – changes intact – first thing tomorrow morning, so you won't

have time in any case. In addition, it is of surpassing importance that you have the paper first thing tomorrow – the fate of the world, suppose, hangs on it. Or at least your career. So it would be a disaster if your computer crashed and you lost all the changes. Fortunately, you have recently purchased backup software that automatically backs up any changes you’ve made and – say – prints out any changed work – at 1:00 every morning.

We want to explore the consequences of assuming that the proposition that your file has been successfully backed up is warranted enough to be a reason you have to believe other things. If it is, then presumably it’s appropriately connected to some of those things as well. And if it’s warranted enough to be a reason you have to believe something to which it’s also appropriately connected, then it *is* a reason you have to believe those things. For example, it’s a reason you have to believe that either the file has been successfully backed up or the moon is made of green cheese. More interestingly, it’s a reason you have to believe that the file will be there in the morning even if the computer crashes.

To make things explicit, call this ‘Premise 1:

(Premise 1) If that the file is successfully backed up is warranted enough to be a reason you have to believe something, then that the file is successfully backed up is a reason you have to believe that it will be there in the morning.

That the file will be there in the morning does not essentially have to do with you, your actions, or your stakes. So perhaps that the file is successfully backed up can be a reason you have to believe it will be there in the morning, even if it’s not warranted enough to be a reason you have to believe other things – propositions that are essentially and explicitly about your actions and/or stakes. For example, might the fact that the file is successfully backed up be a reason you have to believe that the file will be there in the morning, but not be a reason you have to believe that it’s a waste of time to get up and check or that it’s best to just lie in bed and drift comfortably back to sleep? We think not. If it’s a reason you have to believe that the file will be there in the morning, it’s also a reason you have to believe that it’s a waste of time to get up and check and a reason you have to believe that it will have the best results to just lie in bed and drift contentedly back to sleep. We can bring this out by considering a dialogue:

- You: The file is successfully backed up.
- Spouse: So, it will be there when you wake up tomorrow, then, even if the computer crashes.
- You: Yes, and it’s a waste of time to get up and check.
- Spouse: Wait a second – what reason do you have to believe that?
- You: The file is successfully backed up!
- Spouse: I granted that; after all, it’s your reason for believing it’ll be there when you wake up tomorrow. I just don’t see what reason you have to believe that it’s a waste of time to get up and check.
- You: Look, if, as you grant, it is successfully backed up, then it’s a waste of time to get up and check, right? I mean, if it’s backed up and I go check it, I’ll simply find it’s backed up and will only have wasted time. Right?
- Spouse: Yes.

- You: So its being successfully backed up is a reason I have to think getting up and checking is a waste of time.

Your spouse's third reaction is absurd to our ears, while your response is utterly reasonable. In effect, your third remark, and your spouse's agreement, guarantees there is a relevant connection between the two propositions, and this together with the agreement that *the file is successfully backed up* is a reason you have to think it will be there tomorrow simply settles the question of whether it is also a reason you have to think it's a waste of time to get up and check.³ Many other conversations illustrate the same point:

- You: The Yankees will win at least 82 games this year.
- Us: So, they'll finish with at least a .500 record, then.
- You: Yes, and if I bet this 1,000 dollars that they'll win at least 82 games, I'll win some money.
- Us: Not so fast! Why do you think that?
- You: They're going to win at least 82 games this year.
- Us: Well, yeah! Hence they'll finish with at least a .500 record (as we already noted). But the question is why you would believe that if you bet 1,000 dollars that they will, that you'll win some money.
- You: Well, surely *if* they'll win at least 82 games this year, then if I bet the money, I'll win. Right?
- Us: Yes.
- You: So that they'll win at least 82 games is a reason I have to think that if I bet the money I'll win.

Or this:

- You: The ice is thick enough to hold me.
- Your obviously lighter sibling: So it's thick enough to hold me, too.
- You: Yes, and so I won't fall through the ice if I walk across it.
- Sibling: Hold on a second! What reason do you have to believe that you won't fall through if you walk across it?
- You: Well, the ice is thick enough to hold me!

³ An alternative interpretation, here, is that it is the conjunction of the fact that the file is backed up and the fact that if it is backed up it is a waste of time to get up and check, which is the reason to think it's a waste of time to get up and check. Of course, there is nothing special about the fact that there is reference to actions and stakes here. The same interpretation could be given in the case of the reason to think the file will be there in the morning, namely that it is a conjunctive fact consisting of the fact that the file is backed up together with background knowledge concerning the reliability of the computer and the absence of interference & etc. If this 'expanded reason' approach is correct, our arguments below can be reformulated accordingly. So, we could modify Premise 1 to read: If that the file is successfully backed up is warranted enough to be a reason you have for belief, then that the file is successfully backed up is part of a reason you have to believe that it will be there in the morning. And similarly for Premise 2 to follow. These premises would not of course be true of all possible cases (e.g. they won't be true when there are special reasons to doubt the background material making up the second conjuncts of the reasons) but they will be true in the sort of ordinary cases we are discussing. In what follows, we will not make these reformulations, but the 'expanded reason' theorist is free to do so.

- Sibling: Look, I agreed that's your reason for believing it's thick enough to hold me! But I was asking what reason you have for believing that you won't fall through if you walk across.
- You: I'm confused. Obviously, if it's thick enough to hold me – which you say you grant – then if I walk across it I won't fall through it.
- Sibling: Yes.
- You: Well, then, what don't you understand exactly?!

The three examples provide evidence for a more general principle about reasons for belief – that if p is a reason you have to believe a proposition q , then if another proposition r is suitably connected to p , then p is also a reason you have to believe r (even if r , unlike p or q , is about actions or stakes). These examples support this principle in two ways: first, by illustrating the absurdity of granting a proposition when it's used as a reason for believing one thing while – simultaneously with that granting – failing to grant the proposition when it's used as a reason for believing a second, equally well-connected thing, and, second, by illustrating the reasonableness of using propositions as reasons for a well-connected conclusion when mutual agreement has been established about those propositions.

Notice that this conclusion is quite in line with how debates proceed over epistemic closure principles. Consider the widely accepted version deriving from Williamson:

If you know that p and competently deduce q , coming to believe q on that basis, then you know that q .

Some question whether mere competent deduction is enough, or raise issues about the length of the deduction, and so on. What people don't say is this: 'Well, it depends on how much is at stake in whether p and q or whether p or q is about action. Sure, some qs are such that if you know p and competently deduce q , etc. you'll thereby know q . But other qs are such that if you know p and competently deduce them (with an equally simple deduction) – well, there's just too much at stake, so you don't know them.' If knowledge that p is enough to confer knowledge on one deduced q , it's enough to confer knowledge on any deduced q , at least assuming the deductions do not differ in certain ways having nothing to do with their content (e.g. their length). Closure principles are formulated without regard to the contents of the propositions. The principle about reasons for belief that we're endorsing here does no less.

For these reasons, concerning our file case, we now have 'Premise 2':

(Premise 2) If that the file is successfully backed up is a reason you have to believe that it will be there tomorrow morning even if the computer crashes, then that the file is successfully backed up is a reason you have to believe that it's a waste of time to get up and check and a reason you have to believe that it's best to just lie in bed and drift contentedly back to sleep.

Those hoping to reject the conclusion we're ultimately leading to – the existence of pragmatic encroachment in epistemology – will not want to pick Premise 2 as their target. For it is difficult to see how the denial of this premise does not itself constitute a kind of 'encroachment' on the properly epistemic. If the premise is false then there are $p/q/r$ triples

such that while p is warranted enough to be a reason a person has to believe q , p isn't warranted enough to be a reason that person has to believe r , where both q and r are suitably connected 'inferentially' to p , but where only r concerns actions and their results. Why would this be? Why would p be warranted enough to be a reason had for q , but lack this status with respect to r ? The obvious explanation involves encroachment: because of the stakes in whether r , more warrant is needed in a proposition for it to be a reason had to believe r than is needed for that proposition to be a reason had to believe q .⁴

We'll suppose that enemies of pragmatic encroachment will therefore let Premise 2 stand. What next? Well, if that the file is successfully backed up is a reason you have to believe that the file will be there tomorrow morning, then it's a reason you have to believe that it's a waste of time to get up and check and that it will have the best results to just lie in bed and drift contentedly back to sleep. But if that it's successfully backed up is a reason you have to *believe* all those other things, then it's very hard to see how it couldn't be a reason you have to *do* certain things, too. For, consider this conversation:

- You: The file is successfully backed up.
- Spouse: So, it's a waste of time to get up and check, then.
- You: So I'll just lie here in bed and drift contentedly back to sleep.
- Spouse: Wait a second – what reason do you have to do that?
- You: I don't understand. The file is successfully backed up, and as you admit, if it is backed up, then checking is a waste of time and going back to sleep would have better results. So, that's why its being backed up is a reason I have to drift back to sleep.

The absurdity of your spouse's incomprehension and the reasonableness of your explanation here show your spouse's objection to be misplaced. We can illustrate the same point by modifying the other conversations so that they are about action. Here is one concerning the bet:

- You: The Yankees will win at least 82 games this year.
- Us: So, if you bet this 1,000 dollars that they'll win at least 82 games, you'll win some money.
- You: Good idea. I'll do it.
- Us: Not so fast! Why on earth would you do that?
- You: I don't understand. They're going to win at least 82 games this year, as you admit and, and as you also admit, if they do, I'll win the money, and winning money is better than not. So, all this is why their winning at least 82 games is a reason I have to take the bet.

What your conversational partner is attempting to do in these conversations is to convince you to run twin reasoning streams – a 'theoretical' stream for forming beliefs, and a practical one for action. The theoretical stream allows you to draw all sorts of theoretical conclusions – the file will be there tomorrow, it's a waste of time to get up and check, it'll have

4 Another encroachment story might be that a true belief in q has a higher cognitive value than a true belief in r , and thus it takes less warrant for a proposition to be a reason to believe q than it does for it to be a reason to believe r . See our 2009: ch. 7, pp. 194–201, for a discussion of cognitive value encroachment.

better results if I stay in bed than if I don't, etc. But this stream can't be brought to bear on the other practical reasoning stream, or at least it can't be brought to bear when the stakes are high. The above dialogues show just how absurd that suggestion is. If p is a reason you have for belief, and if it is suitably connected to an action, say by virtue of there being a good 'practical syllogism', then p is at least a non-decisive reason you have to perform that action.⁵ Call the instance of this general conclusion manifested in the computer case, 'Premise 3':

(Premise 3) If that the file is successfully backed up is a reason you have to believe that the file will be there in the morning, that it's a waste of time to get up and check, and that it'll have the best results to stay in bed and drift contentedly back to sleep, then that the file is successfully backed up is a reason you have to stay in bed and drift contentedly back to sleep.

So, if that the file is successfully backed up is warranted enough to be a reason you have for belief, not only do you have a reason to believe that the file will be there in the morning and have a reason to believe that it's a waste of time to get up and check and have a reason to believe that it's best to just stay in bed and drift contentedly back to sleep, you also have a reason to just stay in bed and drift contentedly back to sleep.

But, still, the stakes are high! Sure, you have some reason to stay in bed and drift contentedly back to sleep (if that the file is successfully backed up is warranted enough to be a reason you have for belief). But is it ok to just stay in bed and drift contentedly back to sleep? Shouldn't you, rather, be thinking like this: 'Yes, the file is successfully backed up. But, look, it's just not certain that it is, so just in case I'd better get up and check.' This sort of thinking should seem pretty familiar. We seem to engage in it all the time. We say things like, 'I did turn the oven off, but I'd better go back just in case I didn't.' A surgeon, to modify Jessica Brown's (2008) example, might say, 'Yeah, it's the left kidney that we need to take out, but I'll check the chart in case it's not.'

In general, we routinely say things of the form, ' p , but I better check just in case not- p .' Call them *yes-but* statements. Not just any statement of the form, ' p , but I better check' will count as *yes-but* statements in the intended sense. Suppose it is clear that you won't be able to fall asleep if you don't get up and check the file. You might say 'it's backed up but I better get up to check'. This is *not* a *yes-but* statement in the sense we have in mind. You need to check not because of the possibility it isn't backed up but because checking is the only way to end the worrying. You might be like this even in a case in which there was no chance that it wasn't backed up, if it were certain it was. It is important to add 'just in case' to signal that the reason you better get up and check

5 For an extended argument for this conclusion, see our 2009: esp. pp. 71–6. We don't want to commit to what further conditions must be satisfied for p , when warranted enough to be a reason you have to perform an action, to be a reason you have to perform that action. Some might think that p must not only figure in a good bit of practical reasoning in favor of performing the action but that p must also be known to so figure. All we need to claim here is that those conditions can be met when it comes to some p and some actions and that, in the specific case in question, those conditions are met when it comes to the proposition that the file is successfully backed up and the actions of staying in bed and drifting contentedly back to sleep (though see the complications in n. 3).

has to do with the possibility that it isn't backed up. Similarly, in Brown's example, the hospital might have a policy that a surgical patient's chart has to be checked within five minutes of operating. Again, this is not the sort of situation in which a yes-but statement (as we are construing them) would be made.

We've just argued that, in the cases we're concerned with, that the file is successfully backed up is a reason you have *not* to check, if it's a reason you have for anything. So, what should we make of yes-but statements? Do they point to *defeating* reasons? We often have reasons to do things (or not to do things) that are outweighed by other reasons. To use an example from (Fantl and McGrath 2009), suppose you're at the edge of a frozen pond and deciding whether to cross or walk around. You've satisfied yourself completely that the ice is thick enough to hold you: it's 10 below – been that way for weeks – and there are people much heavier than you skating on it. Walking around will take a while, but you also know you're prone to disastrous falls on ice, no matter how thick it is. You seem to have a reason to walk across – that it's shorter than walking around – and a reason to walk around – that there's a risk of slipping and falling. Here, one reason can beat out another, even though both are reasons you have.

Can it work that way with our computer example? If it works the same way, then you have a reason to stay in bed and drift contentedly back to sleep – that the file is successfully backed up. But there is also a competing reason to get up and check. What is that reason? It seems to be this: that there's a pretty decent chance that the file isn't successfully backed up and if it isn't it's disastrous, i.e. there is a *serious risk* that the file isn't successfully backed up. Is it coherent to weigh these two reasons against one another in the way it is coherent to weigh 'Walking across is shorter' against 'There's a chance I'll slip and fall'? Compare how we weigh reasons in uncontroversial cases with how we would be required to weigh reasons in the computer case:

The walk across the ice is shorter, but there's a good chance I'll slip and fall, which would be bad. What's more important, a shorter walk or making sure I don't slip and fall?

Or

There is a 90% chance of rain, so that's a reason I have to take my umbrella, but the umbrella is also really cumbersome. What's more important, the chance it will rain, or the fact it's really cumbersome?

Contrast these examples with:

There's a serious risk the file isn't successfully backed up, so that's a reason I have to get up and check. But the file is also successfully backed up, so that's a reason I have to remain in bed. Which is more important, the serious risk that it isn't backed up or the fact that it is backed up?

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, at least when

the stakes are high. But we find no such thing.⁶ What we find is vacillation: ‘The software is really reliable. Surely the file is backed up. Buuuut . . . there’s a real possibility it didn’t back up. I better not risk it.’ Perhaps even with the right halting tone of voice someone might say, ‘the file’s backed up (isn’t it? *surely* it is, right?). Forget it. I’ll play it safe and get up and check.’⁷ What you don’t find is the likes of, ‘Hmm, the file might not be backed up and that would be bad. That’s one consideration. Another is that it is backed up.’

All this might be puzzling, given that we do make yes-but statements. These statements seem precisely to bring the ‘yes’ part in as a reason but then decide in favor of some factor involving risk against the action that the ‘yes’ part is a reason to perform. What is going on? Why would the very sort of weighing that the yes-but statements seem to rely on be so problematic when made explicit?

Perhaps there is something useful or proper about expressing vacillation in the way we’re committed to saying many yes-but statements do. It certainly doesn’t seem psychologically unrealistic to claim that when we make yes-but statements we are vacillating or are of two minds about whether to embrace the ‘yes’ – that is, the actual-results proposition – as a reason, because of our awareness of potential risks. In making the yes-but statement, one is expressing this two-mindedness, along with one’s inclination to resolve it by refusing to embrace the ‘yes’ part as a reason. If one were to instead say, ‘Maybe I should make sure, but it’s thick enough to hold me’, one would be instead expressing the two-mindedness along with an inclination to resolve it by embracing the ‘yes’ part as a reason. If this explanation were incorrect, if there were only single-mindedness without any vacillation, we would expect to find exactly the absurd weighing of reasons we have mentioned above. We would expect to find ‘the ice is thick enough to hold me’ not dropping out of the subsequent reasoning, but continuing to be present and available for weighing against the serious risk reason. Since we do *not* find this, the most plausible conclusion is that in ordinary life when we make yes-but statements we typically are expressing two-mindedness over whether to embrace the ‘yes’ part, i.e. the actual-results proposition, as a reason.

Why would we express such vacillation instead of simply uttering the part of the statement that we end up embracing? We think there are at least three reasons we might make yes-but statements in high-stakes situations. First, it might be useful for various reasons for our interlocutors to be aware of our internal conflict. Second, we might say something like ‘The file is successfully backed up’, if we were *trying out* how it feels to commit to that proposition. Third, we might say something like ‘The file is successfully backed up’ if we were trying to *reassure ourselves* that it is. In either case, we might find that the

6 Our explanation of the oddity of weighings is that you can’t have both such reasons at the same time for alternative actions. Clayton Littlejohn has helpfully pointed us to another possible explanation. Joseph Raz (1999) distinguishes a kind of reason for φ -ing that also is a reason for not considering reasons not to φ . In Raz’s example, if a superior officer commands an underling to commandeer a civilian’s van, that the superior so commanded is a reason not just to commandeer the van, but to commandeer the van no matter what reasons there are to the contrary (1999: 38). But even if there are such reasons, it doesn’t really fit the case here. That’s because, though the underling perhaps shouldn’t reason by saying, ‘My superior officer commanded it but, on the other hand, I don’t have the right to take a civilian’s van’, there is nothing remotely absurd or senseless about the weighing, nor would we be surprised to find an underling reasoning in this way.

7 Notice the use of ‘surely’. This could be replaced with ‘I’m absolutely sure’, provided we keep the tone of voice the same.

commitment didn't stick. We might find that, in trying out how it feels or in attempting to reassure ourselves, it just rings false or doesn't really command genuine assent. In that case, we might easily follow up that attempted commitment with the expression of an intention to check further. Given the ease and naturalness of this sort of explanation, contrasted with the implausibility of the alternative, we should accept that, in the case under consideration, either that the file is successfully backed up is not a reason you have to remain in bed or that there is a serious risk it isn't backed up is not a reason you have to get up and check.

If that the file is successfully backed up is a reason you have to remain in bed, then it can't be defeated (in the sorts of cases we're discussing here) by reasons having to do with the risk of falsehood. And, in this case, there's nothing else that stands in the way or counts as a potential defeater. (Nor is that the file is successfully backed up merely an enticing reason to stay in bed, in Jonathan Dancy's sense.) So that the file is successfully backed up is a good and undefeated reason you have to stay in bed and drift contentedly back to sleep. (Again, for further discussion, see Fantl and McGrath 2009.) The upshot is what we will call 'Premise 4'.

(Premise 4) If that the file is successfully backed up is a reason you have to stay in bed and drift contentedly back to sleep, then it is a *justifying* reason for you to do this, and therefore you are justified in staying in bed and drifting contentedly back to sleep.

The evidence for Premise 4 is the absurdity of weighing actual-results reasons against reasons referring to the chance that those results don't obtain. The correct explanation for that absurdity, we think, is that we cannot have both such reasons at once. And if we can't have both reasons at once, then if that the file is successfully backed up is a reason you have to stay in bed, then it's an undefeated reason you have to stay in bed and, hence, Premise 4 is true.

One might think, however, that there is an alternative explanation of the oddity of weighing actual-results reasons against reasons referring to the chance that those results don't obtain. Perhaps the statements expressing these weighings are true but improper to assert; perhaps they implicate something clearly false rather than state something clearly false. And perhaps our intuitions of absurdity are tracking the implicature only. If this is the right explanation of the absurdity of the weighings we have in mind, Premise 4 might well be false. We need to consider the possibility of making a 'warranted assertability maneuver' – WAMming Premise 4.

The statements expressing the problematic weighings are statements like the following:

p is a reason I have to A, and the serious risk that not-*p* is a reason I have not to A.

We say that these sorts of statements are absurd because *p* can't be a reason you have to A while not-*p* is a reason you have not to A. That is, we say that the truth of the first conjunct entails the falsity of the second conjunct. But there are other statements that seem similarly absurd even though the truth of the first conjunct does not entail the falsity of the second conjunct:

1. '*p* but there is a significant chance that not-*p*.'
2. '*p* but I don't have any idea whether *p*.'

3. '*p* but I don't believe *p*.'
4. '*p* but I don't know *p*.'
5. '*p* but I don't have good enough evidence to assert *p*.'

If our problematic weighings are like these, then the absurdity of the weighings shouldn't suggest what we think it does.

Of most interest here is the first member in the list: '*p* but there is a significant chance that not-*p*.' For one might think that the absurdity of our problematic weighings can be completely explained by the absurdity of this statement. After all, one might think, the first conjunct in the weighing – '*p* is a reason I have to do A' is either factive (so entails *p*) or otherwise imparts *p* and either way inherits everything that an assertion of '*p*' imparts. And we might further think that 'the serious risk that not-*p* is a reason I have not to do A' entails that there is a significant chance that not-*p* – that the chance that not-*p* is a significant one. If both those claims are true, then the first conjunct of the problematic weighing-statement entails or imparts *p* (the first conjunct of Γ), while the second conjunct of the problematic weighing-statement entails the second conjunct of Γ . Therefore, the problematic weighing-statement either entails or imparts Γ . And so it makes sense that the problematic weighing-statement would inherit the absurdity of Γ . But Γ , while absurd, can well be true. Therefore, we have very little reason to think that problematic weighing statements are necessarily false.

Why is Γ absurd and can the explanation of the absurdity of Γ explain why the problematic weighing statement is absurd? Here's one natural explanation for why Γ would be absurd to say: suppose that an assertion of *p* imparts that the assertion is *proper* and that, furthermore, an assertion of *p* is proper only if the chance that not-*p* is insignificant. If both those things are true, which they plausibly are, then the second conjunct of Γ entails the falsehood of something the first conjunct imparts.

We might think a similar explanation could show why the weighing oddities are odd. According to such an explanation, the first conjunct of the weighing oddities would impart that an assertion of *p* is proper and so would impart that the chance that not-*p* is insignificant, while the second conjunct of the weighing oddity would entail that the chance that not-*p* is significant. However, there is good reason to think that the weighing oddities can't be explained away in this manner. For the oddities remain when we turn to

- third-person statements: '*p* is a reason Leslie has to do A, but the serious risk for Leslie that not-*p* is a reason Leslie has not to do A';
- counterfactual weighings: 'If I had evidence *e*, then *p* would be a reason I had to do A, but the serious risk that not-*p* would be a reason I had not to do A';
- modal statements: 'It's possible for *p* to be a reason I have to do A while the serious risk that not-*p* is a reason I have not to do A';
- past-tense first-person statements: '*p* was a reason I had to do A, while the serious risk that not-*p* was a reason I had not to do A'.

Any imparting of *p* by an utterance of the first conjunct can only impart that the chance that not-*p* is significant for the *speaker*, not the subject. But the second conjunct implies only that the chance that not-*p* is significant for the subject, not the speaker. There is no contradiction generated in this way. Therefore, some other explanation is necessary

for why these weighings are absurd. We think the most natural explanation is that it is impossible for a subject to have both kinds of reasons at once to do contrary things.

Notice how different the situation is when it comes to the analogous variants of \mathbf{r} , above. These are not at all absurd:

- third-person statements: ‘ p , but there is a significant chance for Leslie that not- p ’;
- counterfactual weighings: ‘If I had evidence e , then it would be the case that p while there was a significant chance for me that not- p ’;
- modal statements: ‘It’s possible for it to be the case that p while there is a significant chance that not- p ’;
- past-tense first-person statements: ‘It was the case that p , but there was a significant chance for me that not- p ’.

In these cases the first conjunct imparts only that the chance of error of not- p is insignificant *for the speaker* of the attribution, not for the subject. But the second conjunct entails only that the chance of error is significant *for the subject*, not for the speaker. So no contradiction is generated between what is imparted by the first conjunct and what is entailed by the second.⁸ And, as expected, the conjunctions are not absurd. This is evidence that the absurdity of \mathbf{r} is generated by what’s imparted by the first conjunct. It is evidence that’s lacking in the problematic weighings. It seems unlikely, then, that the odd weighings are odd only because \mathbf{r} is absurd.

There is reason to doubt that the absurdity of \mathbf{r} even explains the oddity of the original first-person present-tense weighings. The first conjunct in these statements – ‘ p is a reason I have to do A ’ – arguably entails, and does not merely impart, that p is assertable. If a proposition is a reason, it is epistemically proper to assert it. But if this is right, then the first conjunct both asserts that p and entails that that assertion is epistemically proper; however the second conjunct entails that the chance that not- p is significant. Given that the proposed norm of assertion is that an assertion that p is epistemically proper just in case the chance that not- p is insignificant, it follows that the weighing statements *are false*. It follows that they are false because they entail both that the assertion of p is proper and that the chance that not- p is significant, which given the assumed norm of assertion is impossible.

Absent a plausible explanation of the oddity of the odd weighings that doesn’t invoke their necessary falsehood, we conclude that Premise 4 is true. Put the Premises 1–4 together and run through the syllogisms: if that the file is successfully backed up is warranted enough to be a reason you have for belief, then it’s a reason you have to believe that it will be there tomorrow when you wake up. And if it’s a reason you have to believe that it will be there tomorrow when you wake up, then it’s a reason you have to believe both that it’s a waste of time to get up and check and that it’s best to just lie in bed and drift contentedly back to sleep. But if it’s a reason you have to believe all of these things, then it’s a reason you have to just lie in bed and drift contentedly back to sleep. Furthermore, there is nothing to defeat this reason. Therefore, if it’s a reason you

8 We should remark that attempts to WAM any of the Premises 1–3 using similar strategies are bound to fail for the same reason. The absurd remarks about reasons for belief or reasons for belief and reasons for action, which we put into the mouths of our imagined speakers, are no less absurd when reformulated in the third person, past-tense first-person, etc.

have to just lie in bed and drift contentedly back to sleep, then you are justified in doing this. Therefore, if that the file is successfully backed up is warranted enough to be a reason you have to believe something – anything at all – then you are justified in lying in bed and drifting contentedly back to sleep. Call this conclusion the *actionability result* concerning the file case:

(The Actionability Result concerning the file case) If that the file is successfully backed up is warranted enough to be a reason you have to believe something, then you are justified in lying in bed and drifting contentedly back to sleep.

These premises have concerned a single case, but each illustrates a more general principle. Analogous premises could be generated for any proposition and suitably connected action.

(The Actionability Result, generalized) If p is warranted enough to be a reason you have to believe something, then p is warranted enough for you to be justified in doing any suitably connected thing.

These results require that either fallibilism or purism is false. Why? Well, suppose for reductio that fallibilism is true and that there is no pragmatic encroachment on being warranted enough to be a reason one has for belief. Well, if fallibilism is true, then propositions that aren't epistemically certain for one can be warranted enough to be reasons one has for belief, and so can be reasons one has for believing at least some things. For example, suppose that one such proposition is the proposition that one's recently purchased backup software has successfully and automatically backed up the changes one has made to one's paper that day. Perhaps the software has a month-long track record of success – whatever it takes for the proposition to be a reason one has for belief, though not yet certain. Now suppose that, though the subject in this scenario doesn't want to lose the changes, it's no big deal if they end up lost – they're easily remedied. It's certainly not worth getting up out of the cozy bed at 1:00 in the morning to go check and back up the file manually. And, anyway, that that the file is successfully backed up is warranted enough to be a reason the subject has for belief. Therefore, according to our instance of the actionability result, the subject is justified in lying in bed and drifting contentedly back to sleep. Of course, since it isn't certain that the file is successfully backed up, there's a possible situation in which, because of the costs and benefits of various actions, the subject shouldn't lie in bed and drift contentedly back to sleep and so isn't justified in doing this, even though in that situation, the subject has the same strength of epistemic position with respect to the proposition.

But in our running example, you're a subject in exactly such a position. You absolutely need that file immediately tomorrow morning. You won't have time to make any changes, and we are supposing a great deal hinges on whether you have it. In such a situation – given the proposition's lack of certainty for you – you should not drift contentedly back to sleep; you should check further. Therefore, that the file is successfully backed up isn't warranted enough to be a reason you have for belief. Since, we may suppose, you and our low-stakes subject are in equally strong epistemic positions with respect to the proposition that the file is successfully backed up, but that the file is successfully backed up is warranted enough to be a reason only the subject has for belief, it follows

that purism is false: there is pragmatic encroachment on the warrant required to be a reason one has for belief. This conclusion follows once fallibilism is assumed. Hence, either fallibilism or purism is false.

The argument is not defused by bickering over this specific case – for instance, by insisting that that the file is successfully backed up can't be warranted enough to be a reason anyone has for anything. The argument trades on general features of the example and can be made perfectly general. If fallibilism is true, then there will be some uncertain proposition that is warranted enough to be a reason someone has for belief. Because the proposition is uncertain, there will be some other person with higher stakes but the same strength of epistemic position with respect to that proposition – an epistemic twin on that proposition – whose epistemic position isn't strong enough to make the proposition warranted enough to be a justifying reason for action. By the actionability result, the proposition isn't warranted enough to be a reason that other person has for belief. In short, if fallibilism and the actionability result are true, there will be cases in which epistemic twins with respect to some proposition, p , differ with respect to whether p is warranted enough to be a reason they have. It follows that purism is false and, hence, that there is pragmatic encroachment on the warrant required to be a reason one has for believing even something else. If fallibilism is true, then you and I can be in equally strong epistemic positions with respect to p , but I can use that proposition as a reason to believe other things (say, p or q) while you can't use that proposition as a reason to believe anything at all.

One central epistemological project is the search for what it takes epistemically for a proposition to be a reason one has for belief. Whatever it takes, either fallibilism is false of that thing, or there is pragmatic encroachment on it. If you think that when you know that p , then p is a reason you have to believe at least something else, then knowledge either requires epistemic certainty, or else there is pragmatic encroachment on knowledge. But you don't escape the dilemma by begging off talk of knowledge. Nor do you get around the dilemma by going contextualist about 'knows'.

Here is one further way that we've heard people respond to the connections we and others have argued for between knowledge and action: they've said that knowledge is not always sufficient to justify action. Sometimes, when the stakes are high, you need to know that you know. And sometimes you need to know that you know that you know. Etc. This response is related to one considered but not taken advantage of by Timothy Williamson (2005: 232): 'in some cases q would be appropriate [as a premise for one's practical reasoning] iff one knew q , in others iff one knew that one knew q , and so on, depending on the stakes.' We've also had it pressed against us regularly in colloquia. But it is inert here. The issue here is what epistemic status a proposition must have to be warranted enough to be a reason to believe other things. And the conclusion is that, if fallibilism is true, the epistemic status a proposition must have to be warranted enough to be a reason one has to believe other things varies with stakes. So, the quasi-Williamsonian solution *is* an impurist view. It says that sometimes in order for p to be a reason you have to believe something else you must know p but, if the stakes are higher, you must know that you know p . That is, in order for p to be a reason you have to believe (p or q), you must sometimes know that p and other times know that you know that p , even if your strength of epistemic position is the same across those times.

Traditionally, it has been of central importance to epistemology what epistemic status a proposition must have for you to get to use it as a reason for believing other things. Knowledge has been at the center of epistemology precisely because it's been thought to

be that status. But if knowledge isn't that status, traditional epistemologists must think that something else is, for that is the traditionally important epistemological status. Therefore, whatever the traditionally important epistemic status is, fallibilist purism is false of it. If traditionalists want to be fallibilists, they have to be pragmatic encroachers. It turns out that the pragmatic didn't need to encroach in the first place. It was already inside.

ACKNOWLEDGEMENTS

Thanks go to Julien Dutant, Clayton Littlejohn, and Jason Stanley for helpful comments on earlier drafts.

REFERENCES

- Brown, Jessica. 2008. 'Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning.' *Nous*, 42: 167–89.
- Fantl, Jeremy, and McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford 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.
- Raz, Joseph. 1999. *Practical Reason and Norms*. Oxford: Oxford University Press.
- Williamson, Timothy. 2005. 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge.' *Philosophical Quarterly*, 55: 213–35.

JEREMY FANTL is Associate Professor of Philosophy at the University of Calgary. Department of Philosophy, University of Calgary, 2500 University Drive NW, Calgary, AB T2N 1N4, Calgary, Canada.

MATTHEW MCGRATH is Professor of Philosophy at the University of Missouri. They are the authors of *Knowledge in an Uncertain World* (Oxford, 2009). Department of Philosophy, 419 Strickland Hall, University of Missouri, Columbia, MO 65211-4160, USA
