

CHAPTER FOUR

Do Practical Matters Affect Whether You Know?

What is the relationship between the propriety of an action based on a belief, on the one hand, and the epistemic justification of the belief itself, on the other? Jeremy Fantl and Matthew McGrath argue that the link is captured by the principle that if your belief that P is epistemically justified, then you're warranted to act in ways suggested by P. And since knowledge requires justification, it follows that if you know that P, then you're warranted to act in ways suggested by P. In short, knowledge is actionable. Fantl and McGrath's view has the consequence – surprising from the perspective of orthodox epistemology – that whether you know, or are justified in believing, P depends in part on how much is at stake in your practical situation. This is called *pragmatic encroachment* or *epistemic impurism*. Justification is not purely a matter of evidence. In his response, Baron Reed argues against the sort of view that Fantl and McGrath defend. Reed argues that there are counterexamples to some of the key principles used to argue for pragmatic encroachment. Reed also argues that pragmatic encroachment overlooks an essential feature of knowledge, that it mischaracterizes practical rationality, and that the relevant intuitions can be explained by alternative means while rejecting pragmatic encroachment.

Practical Matters Affect Whether You Know

Jeremy Fantl and Matthew McGrath

William Clifford (1886) famously argued that “it is wrong always, everywhere, and for any one, to believe anything upon insufficient evidence” (1886: 346). In arguing for this he asked us to imagine that:

Contemporary Debates in Epistemology, Second Edition. Edited by Matthias Steup, John Turri and Ernest Sosa.
© 2014 John Wiley & Sons, Inc. Published 2014 by John Wiley & Sons, Inc.

A shipowner was about to send to sea an emigrant-ship. He knew that she was old, and not over-well built at the first; that she had seen many seas and climes, and often had needed repairs. Doubts had been suggested to him that possibly she was not seaworthy. These doubts preyed upon his mind and made him unhappy; he thought that perhaps he ought to have her thoroughly overhauled and refitted, even though this should put him to great expense. Before the ship sailed, however, he succeeded in overcoming these melancholy reflections. He said to himself that she had gone safely through so many voyages and weathered so many storms that it was idle to suppose she would not come safely home from this trip also. He would put his trust in Providence, which could hardly fail to protect all these unhappy families that were leaving their fatherland to seek for better times elsewhere. He would dismiss from his mind all ungenerous suspicions about the honesty of builders and contractors. In such ways he acquired a sincere and comfortable conviction that his vessel was thoroughly safe and seaworthy; he watched her departure with a light heart, and benevolent wishes for the success of the exiles in their strange new home that was to be; and he got his insurance-money when she went down in mid-ocean and told no tales. (339)

Asks Clifford, "What shall we say of him? Surely this, that he was verily guilty of the death of those men" (339).

Let us agree that Clifford's verdict is correct. The shipowner – sending the ship out as he does while stifling legitimate doubts about the ship's seaworthiness – is guilty of any deaths that result when the ship sinks because it unseaworthy. What, though, is the source of the shipowner's guilt? We might think that the source of the guilt is the mere fact that the ship sank and the passengers died – that it is because the passengers died that the shipowner is culpable. But, says Clifford, that cannot be the whole of the matter. For even if the shipowner lucked out and the ship returned safely, the shipowner would have been just as culpable: the shipowner "would not have been innocent, he would only have been not found out" (340).

Perhaps, then, it is not the results of the act that contribute to the shipowner's guilt. Perhaps the shipowner's guilt is constituted simply by his culpability for the act itself: on such meager evidence, the act of sending the ship out to sea is a culpable act. This, Clifford agrees, must be admitted. Regardless of the shipowner's belief, the shipowner had an obligation to inspect the ship more carefully before sending it out. But, he continues, the culpability of the act cannot exhaust the shipowner's culpability. That is because, in believing as the shipowner does, the shipowner makes it much more likely that the act will follow – by believing as he does, the shipowner increases the likelihood that he will act irresponsibly. In fact, if the shipowner's attitude did not conduce toward the irresponsible action, it would not really amount to a genuine belief:

No man holding a strong belief on one side of a question, or even wishing to hold belief on one side, can investigate it with such fairness and completeness as if he were really in doubt and unbiased; so that the existence of a belief not founded on fair inquiry unfits a man for the performance of this necessary duty.

Nor is that truly a belief at all which has not some influence upon the actions of him who holds it. He who truly believes that which prompts him to action has looked upon the action to lust after it, he committed it already in his heart. (342)

For these reasons, Clifford concludes that there is a link between the normative status of a belief and the normative status of an action suggested by that belief: "it is not possible

so to sever the belief from the action it suggests as to condemn the one without condemning the other” (342). If that’s right, then in condemning the shipowner’s acting upon the belief that the ship is seaworthy, we unavoidably condemn the belief itself.

We think that Clifford’s link is intuitively powerful. It also, when properly spelled out, has some surprising and counterintuitive consequences. The fact that it has such consequences means that the intuitive power of his link needs to be supplemented by argument if it is to be acceptable. We’ll do that here as well. But we don’t want the initial intuitive power of Clifford’s link to be lost by the fact that we are supplementing it with principled argument. For it is clear that when it comes to the condemnation of action, we often feel forced to condemn the belief that suggests the action. James Montmarquet (1993), for example, points out that:

There are times when we want to hold an agent morally to blame for conduct which, from that agent’s *own* point of view, seems quite justified. Cases in point abound, including many of the murderous activities of tyrants, terrorists, racists, and religious fanatics, not to mention the less murderous activities of plain self-righteous hypocrites ... But how are we entitled to cast such moral blame and even fairly punish these individuals unless we can sometimes find these individuals culpable for having these beliefs in the first place? ... If we cannot assign culpability for holding such beliefs, how can we assign culpability for acts premised on their (assumed) truth? (1–2)

If there were no such link of the sort that Clifford alleges, it would be hard to see how the need to condemn beliefs in order to condemn action could be made coherent.

But what precisely does it mean to say that it is not possible so to sever the belief from the action it suggests as to condemn the one without condemning the other? One thing it means is that if an action suggested by some belief is condemnable, then so is the belief itself. Of course, there are many kinds of grounds for condemnation: you might condemn an action because it harms others, or harms the agent, or harms yourself. You might condemn an action because it is uncouth, or immoral, or unseemly. Some criticisms of action do not filter back down to every one of the beliefs that suggest them. For example, a thief might rob a bank because of a rather well-informed belief that the bank has a particularly large supply of cash and a particularly poor security system. The action is morally condemnable, but the belief is not. There may, of course, be other beliefs operant in the thief’s action – that, for example, it is permissible for the thief to rob the bank – and these beliefs may inherit some of the culpability of the thief’s action, but certainly not all the beliefs suggestive of the thief’s actions do so.

Still, we can say this on behalf of Clifford: the reason the shipowner’s action of sending the ship out to sea is condemnable is that the shipowner did not have enough evidence that the ship was seaworthy. This is not so in the case of the thief; the thief’s action is condemnable, but not because the thief lacks evidence that the bank has a poor security system. To introduce a term of art, let’s say that in the shipowner case, the proposition that the ship is seaworthy is not “warranted enough” to justify sending the ship out to sea. On the other hand, in the thief case, that the bank has a poor security system is warranted enough to justify robbing the bank (though, of course, robbing the bank was not justified for other reasons).¹ According to Clifford, then, when a proposition isn’t warranted to justify some action suggested by the proposition, the agent’s belief in that proposition is itself condemnable.

Clifford's paper is entitled "The Ethics of Belief," but the condemnation of the shipowner's belief is not just an ethical condemnation. If there is a moral requirement to trust the word of your closest friends no matter what evidence is against them, then belief in your closest friend's guilt in a murder case might be, in some sense, ethically condemnable, even if decisively supported by evidence. Here the fault of the belief doesn't stem from a lack of evidence. Not so with the shipowner's belief: the grounds of the condemnation are epistemic – had the belief been sufficiently supported by evidence, the action it suggests would *not* have been condemnable and, therefore, neither would have the belief. That means that, when Clifford says that an action is only condemnable if the belief that suggests it is condemnable, to the extent we agree with Clifford, we should think that in the cases he has in mind, if an action is condemnable, then the belief suggestive of it is, not just ethically, but epistemically condemnable; in the terminology of contemporary epistemology, the belief is "unjustified." According to Clifford, then, if a proposition suggestive of some action is not warranted enough to justify that action, then belief in that proposition is unjustified. Or, contraposing:

If your belief that p is justified, then p is warranted enough to justify actions suggested by p .

According to epistemological orthodoxy, agents know that p only if they are justified in believing that p . That means a simple syllogism delivers what we will call "Clifford's Link":

(Clifford's Link) If you know that p , then p is warranted enough to justify actions suggested by p .²

What does it take for an action to be "suggested" by p ? Right now, suppose, you believe that you have the spice cumin at home. Nothing much hinges on this, suppose (you're not planning a dinner that requires cumin). As such, you're perfectly justified in continuing as you are – working, flying your kite, reading, or whatever you happen to be doing. You don't need to rush home and check to see if you indeed have cumin at home. But there are hypothetical situations in which much hinges on whether you have cumin. There are situations in which you are planning a cumin-involving dinner. Worse, there are situations in which you are going to have a relative visit who is deathly allergic to cumin and will fall ill even if there is an unopened bottle in the house. And the stakes might be raised even higher than this. In such situations, your evidence that you have cumin may not be sufficient for you to continue as you are. You may need to go home and check. Does Clifford's Link require that, to know that you have cumin at home, that you have cumin at home must be warranted enough to justify you in not going home in your actual, low-stakes situation? Or does it require that, to know that you have cumin at home, that you have cumin at home must be warranted enough to justify you in not going home in all possible hypothetical situations, no matter how high the stakes? If it requires the latter, then it seems that you don't know you have cumin at home. More generally, for almost any proposition we take ourselves to know, it looks like there are hypothetical cases in which the stakes would be so high that we wouldn't be able to justifiedly perform actions suggested by that proposition. If so, and if Clifford's Link requires p to be warranted enough to justify all actual and hypothetical actions suggested by p , then it looks like no propositions will be known by anyone.

This skeptical conclusion might be a cost some are willing to bear. But we think it is better to understand Clifford's Link not as meaning that if you know that p then p is warranted enough to justify, in any hypothetical situation, actions suggested by p , but rather that if you know that p , then p is warranted enough to justify, in your practical situation, actions suggested by p . This allows you to have all sorts of knowledge of propositions about which there is a reasonable amount at stake in what you do depending on whether the proposition is true. It allows you, for example, to know that you have cumin at home.

However, allowing knowledge of this sort has surprising consequences. It means there is what Jon Kvanvig (2004) has called "pragmatic encroachment" on knowledge. On this view, whether we know something depends on practical features of our context: whether we know can vary depending on what's at stake in whether the proposition is true.³ To see this, consider a variant of Clifford's shipowner example. In this variant, suppose that you are not a shipowner but a visitor at a nautical museum. During your tour of the museum you pass by a full-size three-masted ship. Interested, you read the placard⁴:

The Star of Italy

This iron three-masted ship was built by Harland and Wolff, out of Belfast. Weighing 1784 tons, it was launched in 1877 and taken out of service in 1927. It is still seaworthy.

You nod your head approvingly, think, "Huh! It's still seaworthy. How about that!", and move on. Here, you satisfy the condition for knowing that the ship is seaworthy required by Clifford's Link. And, assuming as we are that knowledge is possible even if there are some hypothetical situations in which doing what is suggested by the proposition is unjustified, we should say that you do know that the ship is seaworthy.

However, evidence of this quality would not have been enough for Clifford's shipowner to be justified in sending the ship out to sea. The say-so of a museum's placard isn't sufficient to risk the hundreds of emigrant-lives the shipowner risks by sending the ship out to sea. Therefore, if all that Clifford's shipowner had was the same evidence you have, by Clifford's Link, the shipowner would not have known that the ship was seaworthy. In short, even with the same evidence, you know, but the shipowner does not, that the ship is seaworthy. And notice that this is not because in the one case it's true that the ship is seaworthy and in the other case it's not: it may very well be true in both cases. As Clifford points out, it's not the fact that the shipowner is wrong that makes him at fault for sending out the ship or believing it's seaworthy. It's the fact that he believed and sent the ship out on insufficient evidence. Nor is belief lacking in one case but not in the other. Both you (the museum visitor) and Clifford's shipowner believe that the ship is seaworthy. What's preventing Clifford's shipowner from knowing is that he doesn't have enough evidence to know, even if he has the same evidence that you do, and we're supposing you do have enough evidence to know.

The result, we think, is not defused by bickering over this specific case – by saying, for example, that you (the museum visitor) can't know on such evidence that the ship is seaworthy. We're allowing – as we must to avoid skepticism – that you can know

that p even if there are some hypothetical cases in which you have the same evidence for and against p and you can't act on p . In such cases, even though you have the same evidence for and against p in the actual and hypothetical cases, you will differ in whether you know that p .

This possibility seems counterintuitive. If two subjects share the same evidence for and against p (and also, as is possible, satisfy all the other standard conditions on knowledge – truth, belief, and the absence of Gettierish features), it seems that one subject knows only if the other does. Knowledge can be lost or gained by gaining or losing information – by acquiring new evidence or forgetting old evidence. But you can't lose knowledge, it might seem, by changing what you care about, your available options, or the expected costs and benefits of acting in various ways. If Clifford's Link is correct, though, and skepticism about knowledge is false, it seems you can.

Such a result might prompt some, despite the initially compelling power of Clifford's Link, to deny Clifford's Link or, at the very least, wonder what more there is to be said in its favor. There is, at least, this: Clifford's Link explains nicely various linguistic habits we have – habits of citing and asking about knowledge to defend, criticize, and explain action. For instance, you might say, after someone criticizes you for taking a left instead of a right to get to the airport, "I know that there's a shortcut to the left." This habit is hard to make sense of unless, in citing knowledge, you're saying that the proposition you're claiming to know is good enough to justify the action suggested by it – namely, taking a left instead of a right. Likewise, with questions: if your spouse tells you that you should get ready to leave for the 7.50 movie, you might well ask, "Do you know that's when it starts?," implying that if the answer is yes, then leaving will indeed be justified. These habits do not vanish when the stakes are high. You might desperately need to get to the airport as soon as possible and defend your taking a left instead of a right by saying "I know this is a shortcut." Again, when the stakes are high, you might ask, "the evidence is strong, but do the scientists know this drug is safe?" But some might think this linguistic data can perhaps be explained in other ways, so for the rest of this paper, then, we'd like to offer a more principled argument in favor of Clifford's Link.

What might seem to be the problematic feature of Clifford's Link is that it makes what should be a purely epistemic concept – knowledge – subject to practical conditions. So we'd like to start our argument for Clifford's Link by arguing for a more mundane-seeming epistemic requirement on knowledge. According to closure principles on knowledge, what you know can justify you in believing all sorts of things that are entailed by what you know. So, if you know that the liquid on the table is white, that can justify you in believing that the liquid isn't red. This seems plausible, but we don't want to assume anything even this strong. One way to weaken this requirement is to talk, not of what justifies what, but about what reasons knowledge can provide. If you know that the liquid spilled on the table is white, for example, then you have at least some reason to believe that the liquid isn't red. And, if you know that the liquid on the table is white, then you even have some reason for thinking that the liquid on the table is milk: you have more reason for thinking that it's milk than that it is orange juice, say. Of course, you might know that the liquid spilled on the table is white but not thereby have a reason to believe other things – for example, that the ship is seaworthy. But, in such cases, the target belief isn't suggested by what you know – that the liquid is white.

So, let's take as our **first premise** in our argument for Clifford's Link this: if you know that p , then p is warranted enough to be reason you have to believe some propositions suggested by p .⁵ If you know that the ship is seaworthy, for example, then that the ship is seaworthy is a reason you have to believe that the ship will return safely from a modest voyage on normal seas, and if it is a reason, then of course it is warranted enough to be such a reason.

But what about other beliefs – beliefs about the potential consequences, costs, or benefits of your own actions? It's one thing to say that what you know can be a reason you have to believe some neutral consequences of what you know. It's another thing to say that what you know can be a reason you have to believe some action-regarding consequences of what you know. Or so it seems. In fact, though, it is quite bizarre to use what you know as a reason to believe some neutral propositions suggested by what you know while refusing to use what you know as a reason to believe some action-regarding propositions suggested by what you know. Consider this dialogue between the shipowner and his son, where the inspection in question is a simple "seaworthiness" inspection that a ship passes (if seaworthy) or fails (if not):

SHIPOWNER: The ship is seaworthy.
SON: So, it'll return safely from a modest voyage on normal seas.
SHIPOWNER: Yes, so it's a waste of time and money to get it inspected.
SON: Wait a second – what reason do you have to believe that?
SHIPOWNER: It's seaworthy!
SON: I granted that; after all, it's your reason for believing it'll return safely from a modest voyage on normal seas. I just don't see what reason you have to believe that it's a waste of time and money to get it inspected.
SHIPOWNER: Look, if – as you grant – it's seaworthy, the inspection will simply reveal that it's seaworthy and I will only have wasted time and money. Right?
SON: Yes.
SHIPOWNER: So that the ship is seaworthy is a reason I have to think it's a waste of time and money to have it inspected.

The son's third speech is absurd to our ears, while the shipowner's response is utterly reasonable. What the shipowner's third remark makes clear is this: that the ship is seaworthy suggests that it's a waste of time and money to have the ship inspected. Further dialogues only reinforce this point. Suppose you're at the edge of a frozen pond of an oval shape. You're at the mid-point of the long side of the oval. It's a long way around but not far across. You want to get to the other side.

YOU: The ice is thick enough to hold me.
YOUR OBVIOUSLY LIGHTER SIBLING: So it's thick enough to hold me, too.
YOU: 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!
SIBLING: Look, I agreed that's your reason for believing it's thick enough to hold you! But I was asking what reason you have for believing that you won't fall through if you walk across.

Here again your sibling's third reaction seems absurd. Just as epistemologists don't restrict closure principles to propositions that are "neutral" regarding action, we shouldn't restrict principles about reasons for belief in this way either.

Thus, we have our **second premise**: if p is warranted enough to be a reason you have to believe some propositions suggested by p , then p is warranted enough to be a reason you have to believe *any* proposition suggested by p .

Our next step links reasons for belief with reasons for action. Here especially we attempt to draw on the insights of Clifford. When we appeal to some consideration as a reason for belief, we tend to treat it as available as a reason for action as well. We do not keep two reasoning streams in our head, one with propositions available as reasons for belief and another with propositions available as reasons for actions. We treat something as available as a reason for action when it is available as a reason for belief. And not only do we do this, it seems absurd to segregate the two. Consider the following dialogue:

- SHIPOWNER: The ship is seaworthy.
SON: Yes, so it's a waste of time and money to get it inspected.
SHIPOWNER: Right, so I'll send it for the emigrant trip without bothering with an inspection first.
SON: Wait a second – what reason do you have to do that?
SHIPOWNER: It's seaworthy!
SON: I granted that: after all, it's your reason for believing that it is a waste of time and money to get it inspected. I just don't see what reason you have to send it out without an inspection.
SHIPOWNER: I don't understand. The ship is seaworthy, and as you admit, if it is seaworthy, then the inspection is a waste of time and money, as the inspection will simply come back "seaworthy" and we will be out the fee for the inspection. So, this is why its being seaworthy is a reason I have to send it out without inspection.

Here, again, the third comment from the son seems bizarrely uncomprehending and the shipowner's final remark clarifies exactly why, if the ship's being seaworthy is a reason the shipowner has to believe it's a waste of time and money to get it inspected before sending it out, then its being seaworthy is also a reason the shipowner has to send it out without inspection. A similar dialogue for the ice case would lead to the same conclusion. It would be bizarre to say or think, "I grant that the ice's being thick enough to hold you is a reason you have to believe you'll be fine and save some time to cross the frozen pond rather than walk the long distance around, but I just don't see what reason you have to walk across." The proper response is: "I just explained why it is a reason I have: since it is a reason I have to believe that I'll be fine and save time to walk around, it is a reason I have to walk across rather than around."

These reflections support our **third premise**: if p is warranted enough to be a reason you have to believe any proposition suggested by p , then p is warranted enough to be a reason you have to perform actions suggested by p .

Reasons to perform a given action provide at least some support for performing that action, but they needn't *justify* you in performing the action. Perhaps the reason is only *pro tanto*, that is, only weighs partially in favor of performing the action. So, if you are prone to dangerous falls on ice, the fact that it is thick enough to hold you (and is shorter across than around) is only a *pro tanto* reason to cross it, and can be

outweighed by a reason not to cross it, that you're prone to disastrous falls on ice. One thought, then, is that knowing something to be true does make it a reason you have for action, but this reason can be defeated if the practical stakes are high enough (and your warrant for the reason is less than absolutely conclusive). So, perhaps when you read that the ship is seaworthy in the nautical museum, you know it is seaworthy and this can be a justifying reason to do various things in that low-stakes situation, such as admire it; whereas, if you were in the shipowner's situation, evidence of that same quality would give you less usable knowledge – you would know that the ship is seaworthy and you would thereby have a reason to send it out on the emigrant voyage but this reason, because of the stakes involved, wouldn't justify you in sending it out, because it would be defeated by something having to do with the chance of error together with the high stakes involved.

To plug this gap, we need to argue for a principle about reasons for action, and this will be our **fourth premise**: if p is warranted enough to be a reason you have for performing an action suggested by p , then it is warranted enough to justify you in performing that action.

The shipowner case seems like a perfect test case. The shipowner, assume, has a reason to send out the ship, namely that the ship is seaworthy. The stakes are high, though. Perhaps this is a case in which this reason is defeated by factors having to do with the fact that there is a chance that the ship isn't seaworthy, and that if it isn't, horrible results might ensue by sending it out on the seas with many emigrants in it. Perhaps, that is, that the ship is seaworthy is outweighed or defeated by the *serious risk* that the ship isn't seaworthy. Suppose that this is so: that the serious risk that it isn't seaworthy outweighs or defeats its being seaworthy, with the result that the shipowner isn't justified in sending out the ship. Given all this, it ought to be perfectly acceptable to weigh the two reasons explicitly and to conclude that the serious risk reason wins out. But consider statements expressing such weighing:

There's a serious risk that the ship isn't seaworthy, so that's a reason I have not to send it out. But the ship is seaworthy, so that's a reason I have to send it out. Which is more important, the serious risk that it isn't seaworthy or the fact that it is seaworthy? The serious risk that it isn't. So, I shouldn't send it out.

This speech sounds absurd, not merely absurd to say but absurd to think. It's not that it is absurd to weigh reasons having to do with facts with reasons having to do with chances. It's fine to weigh the cumbersomeness of the umbrella (a reason having to do only with facts) with the chance that it will rain. What's absurd is to weigh a proposition against the serious risk that that very proposition is false. In fact, it's absurd to weigh propositions *about chances* against the serious risk that those very propositions are false: "On the one hand, p might not be true. On the other hand, p is true. Which is more important?" Whatever the content of the proposition is, it seems absurd to weigh it against the serious risk that it is false.

One might think that this sounds strange because it is so obvious that the serious risk reason wins out. But this won't do, for two reasons. First, it *isn't* obvious that, if you could have both reasons at once, the serious risk reason would win. What we care about fundamentally is what will happen, not what is likely to happen. We care about actual results first and expected results derivatively. Second, even supposing that you could

have both reasons and that the serious risk reason would win, this wouldn't explain why the weighing statements sound absurd. It should rather sound all too obvious – pedantic – to make or think the weighing statements. A pedant isn't being absurd.

If Premise 4 is false, we ought to expect the weighing statements to be perfectly acceptable, at least in high stakes cases, but they are not. They seem absurd. This is grounds for thinking Premise 4 is true.

We want to address a remaining worry one might have about Premise 4. If this premise is true, the weighings should sound absurd, and they do. But why is it, then, that we sometimes find ourselves saying the likes of “it is seaworthy, but I'd better not send it out, just in case” and “the ice is thick enough to hold me, but I'd better not walk on it, just in case”? Don't these statements – call them “yes, but” statements – suggest Premise 4 is false? There is a puzzle here. The absurdity of the weighing statements is reason to think Premise 4 is true, and the fact that we sometimes assert and think “yes, but” statements seems to be a reason to think Premise 4 is true; and *yet* the appropriateness of “yes, but” statements seems to imply the appropriateness of the weighing statements!

We think the puzzle is best resolved by taking the “yes, but” statements to express vacillation or two-mindedness. Notice how odd it would be, and closer to a problematic weighing statement, after saying, “it is seaworthy, but I'd better not send it out, just in case” to follow-up with “and, again, it is seaworthy; I realize that.” Why would we go in for “yes, but” statements, though? What is the point of expressing vacillation in this way? We offer some speculative guesses. For one thing, it might be useful to express our inner conflict to others and to ourselves, and “yes, but” statements allow this. For another, we might not merely wish to express vacillation; rather, we might want to “try out” a possible commitment to the “yes” part, either to reassure ourselves or in figuring out whether it's really so, but then upon finding that it doesn't stick, that it feels wrong, we back out of this commitment by asserting the “but” part. Finally, we might be anticipating a sort of regret for playing it safe and finding out that the precautions weren't necessary, that all would have turned out all right. Later, we might say to ourselves, “I knew it was all going to be fine. I shouldn't have been so worried.” In saying, “the ship is seaworthy, but I'll not send it out just in case,” one might well be anticipating a future regret for playing it safe. These are speculative explanations. What seems to us hard and fast is that if the “yes, but” statements expressed single-mindedness about the issue at hand, the weighing statements would not be absurd, which they clearly are.

Clifford's Link follows from our four premises. Suppose you know that p . Then by Premise 1, p is warranted enough to be a reason you have to believe some propositions suggested by p . By Premise 2, it follows that if you know that p , p is warranted enough to be a reason you have to believe any proposition suggested by p . By Premise 3, which connects reasons for belief with reasons for action, it follows that if you know that p , then p is warranted enough to be a reason you have for actions suggested by p . Finally, given all this, Premise 4, which takes warrant enough to be a reason to be sufficient for warrant enough to justify, enables us to conclude that Clifford's Link is correct: that is, if you know that p , then p is warranted enough to justify you in actions suggested by p .

Supposing that the level of warrant sufficient to justify a p -suggested action can vary with stakes, a plausible fallibilist assumption, then if we accept Clifford's Link we will also have to admit that knowledge varies with stakes. And as we noted above, this is a

counterintuitive claim. There are two possibilities we want to mention, briefly, about how one might cope with this counterintuitiveness. First, one might wish to open up the “philosophy of language toolbox” to explain the counterintuitive results. It is not implausible to think that, while the core relation of knowledge to the practical is stated in Clifford’s Link, it is sometimes useful to apply standards appropriate to one party’s situation more generally. Thus, if what I care about in my situation is using an informant who can give me information I can act on, I might not care about the informant’s stakes; so, I might not call the informant a knower unless I can take what she tells me as a reason for action in my situation. One could then say, if one wished, that my knowledge attribution ‘S knows that p ’ is true in my context of speech only if the warrant for p that S possesses is strong enough to make p a reason for action in my stakes situation. We cannot discuss all the details here of working out such a proposal here.⁶ A second possibility is that, after all is said and done, you just might not be able to accept that knowledge can vary with changes in stakes. This would be a reason for thinking, however surprisingly, that Premise 1 is false. However, Premises 2 through 4 say nothing about knowledge. And they themselves jointly entail, given suitably fallibilist assumptions, that something of real epistemological importance can vary with changes in stakes, namely the status of having a proposition as a reason for belief. If we are right, then, there are compelling grounds for thinking that something of epistemological importance depends on practical matters – if not knowledge, then reasons for belief.⁷

Notes

- 1 This is a case in which something is warranted enough to justify an action but does not justify it, because it isn’t a good reason for the action. In general, being warranted enough to justify a certain action doesn’t suffice for justifying that action, just as being old enough to be the US president doesn’t suffice for being the US president. But if p is warranted enough to justify you in an act ϕ , then any shortcomings in your warrant for p do not stand in the way of p justifying you in ϕ -ing.
- 2 What is it for a belief that p or action ϕ to be suggested by p ? In the case of belief, we can think of it as a matter of there being a good argument of the form “ p , therefore q .” In the case of action, we can think of p being such as, if it were true, to make a good case for ϕ -ing. Thus, the shipowner, whether he thinks he knows or not, might say “well, if it is seaworthy, that’s a reason to send it out.” Here the shipowner is not saying he *has* a good reason to send it out, but only that if a certain fact obtains then there is a reason to send it out.
- 3 Pragmatic encroachment is defended by, among others, Fantl and McGrath (2002, 2009), Hawthorne (2004), Stanley (2005), and Hawthorne and Stanley (2008). Critics include Brown (2008), Nagel (2008), and Reed (2010).
- 4 Information found at [http://commons.wikimedia.org/wiki/Category:Star_of_Italy_\(ship,_1877\)](http://commons.wikimedia.org/wiki/Category:Star_of_Italy_(ship,_1877))
- 5 Except perhaps in cases of reasoning which would be question-begging, such as the famous reasoning, “I have hands, and so I am not a brain in a vat.” When one proposition is warranted enough to be a reason you have to believe another suggested by it, it is a reason you have to believe another suggested by it. Since no worries about question-beggingness arise in all the cases we are considering, we will mostly ignore the difference between “warranted enough to be a reason you have to believe” and “is a reason you have to believe.”
- 6 See Fantl and McGrath (2009, pp. 55–58, 211–212).
- 7 The line of argument developed here is developed in more detail in Fantl and McGrath (forthcoming) and especially Fantl and McGrath (2009, chapter 3).

References

- Brown, J. (2008) Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs* 42: 167–189.
- Clifford, W. (1886) The ethics of belief. In *Lectures and Essays by the Late William Kingdon Clifford, F.R.S.*, ed. L. Stephen and F. Pollock, 339–363. London: Macmillan.
- Fantl, J. and McGrath, M. (2002) Evidence, pragmatics, and justification. *The Philosophical Review* 111: 67–94.
- Fantl, J. and McGrath, M. (2009) *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Fantl, J. and McGrath, M. (forthcoming) Pragmatic encroachment: it's not just about knowledge. *Episteme*.
- Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008) Knowledge and action. *Journal of Philosophy* 105: 571–590.
- Kvanvig, J. (2004) Pragmatic aspects of knowledge? Blog. March 18, 2008. <http://certaindoubts.com/?p=13> (last accessed March 28, 2013).
- Montmarquet, J.A. (1993) *Epistemic Virtue and Doxastic Responsibility*. Lanham, MD: Rowman and Littlefield.
- Nagel, J. (2008) Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy* 86: 279–294.
- Reed, B. (2010) A defense of stable invariantism. *Noûs* 44: 224–244.
- Stanley, J. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Practical Matters Do Not Affect Whether You Know

Baron Reed

One of the central questions that epistemologists have attempted to answer is this: when someone knows something, what explains or grounds that knowledge? Many different answers have been given – for example, clear and distinct perception, evidence, sense experience, reliable belief-producing processes – and the differences between them should not be ignored. But it is interesting to note that all of these answers agree in at least one respect: they are all focused on something that is *truth directed*. For example, evidence is thought to be an indication of what the truth is, reliable belief-forming processes are reliable insofar as they tend to produce true beliefs, and so forth.

Traditionally, this point of agreement was so widely and deeply shared that epistemologists never really thought about it. It has been given a name – *intellectualism* or *purism* – only recently and only by the relatively small number of philosophers who have argued against it.¹ According to these philosophers, knowledge can be only partially explained by (or grounded in) truth-directed things like those mentioned above. A full explanation also has to include the practical stakes of the person in question. In this sense, the pragmatic encroaches on the epistemic.²

The thesis of pragmatic encroachment has been defended in three ways. First, it is thought to be the best explanation of how we use knowledge attributions and knowledge denials in ordinary conversation, when we defend, criticize, or excuse actions. Second, it has been argued that pragmatic encroachment is needed to make sense of the way in which we use knowledge in practical reasoning. And, third, pragmatic encroachment is thought to be supported intuitively by consideration of pairs of cases in which subjects who are identical with respect to truth-directed things (like evidence and reliability) differ in their practical stakes and also differ epistemically – that is, one of the subjects has knowledge that the other lacks.

In what follows, I shall present each of these arguments for pragmatic encroachment. I shall then raise several objections to the view, which will encourage a reconsideration of those arguments. Finally, I will sketch an alternative view that better explains the connection between knowledge and practical reason while remaining within the bounds of traditional, truth-directed epistemology.

1 The Case for Pragmatic Encroachment

Defenders of pragmatic encroachment typically argue for a principle that links knowledge with practical rationality. For example, John Hawthorne and Jason Stanley defend the following Reason-Knowledge Principle (RKP):

Where one's choice is p -dependent, it is appropriate to treat the proposition that p as a reason for acting if and only if you know that p .³

Suppose that I am trying to decide whether to go to the library because I need to meet my friend, Richard. In that sense, my choice of whether to go to the library is dependent on the proposition that Richard will be there. Using this proposition in my deliberation is permissible just in case I know it to be true.

Notice that RKP specifies that knowledge is both *sufficient* and *necessary* for action.⁴ That is, when one has the relevant knowledge, one's epistemic position is good enough to make the action in question rational. (This is compatible with it failing to be rational for some non-epistemic reason – for example, the action isn't optimal for attaining one's goal.) And an action is rational only when it is grounded in knowledge. In other words, if one does not have the relevant knowledge, one does not have the epistemic grounding needed to make the action rational.

We are now in a position to fully appreciate the sense in which the pragmatic is thought to encroach upon the epistemic. If RKP, or something like it, is correct, then knowledge is sufficient for practical rationality: if one knows that p , then one may treat the proposition that p as a reason for acting. There is a logically equivalent way of reading this conditional: one knows that p only if one may treat the proposition that p as a reason for acting. Read in this way, it makes clear that practical rationality is necessary for knowledge. If our concern is to provide the correct analysis of knowledge, then, we have to include in the list of necessary conditions, not only things like truth, belief, and justification, but also suitability for use in practical reasoning. On this view, the practical is embedded in the very nature of knowledge.

1a Practical talk about knowledge

We make claims about knowledge – attributing or denying it, to ourselves and to others – for a variety of reasons. Consider the following examples.

- (1) Maria knows when World War I occurred.
- (2) Ask Ronald – he knows where I parked the car.
- (3) Why did you lock the door? You knew I forgot my key.
- (4) I knew they wouldn't mind sitting on the balcony.
- (5) I didn't know you were allergic to peanuts.

Sometimes, as in (1), our aim is merely to describe reality. Maria's knowing when World War I happened is a fact in just the same way it is a fact that the Amazon is the biggest river in the world. But there are also many uses of knowledge claims that serve a more obviously practical purpose. For example, the speaker of (2) is flagging Ronald as a good source of information, given the hearer's practical needs.

In other instances, the connection between knowledge and action is more intimate still. Some of the most interesting practical uses of knowledge claims are to criticize, defend, or excuse actions, as in (3), (4), and (5), respectively. The thought is that, if RKP were true, we would have an explanation for these natural ways of talking. In (3) and (4), the speaker is attributing knowledge to the subject in question. In doing so, she has indicated that the subject has no epistemic obstacle to performing the relevant action. If the action is also the one that best serves the subject's ends, it is rational for the subject to perform the action – for example, buying tickets for the balcony, as in (4) – and a failure of practical rationality if she does *not* perform the action – like failing to leave the door unlocked, as with (3). In both of these cases, the speaker appears to be relying on the (alleged) fact that knowledge is sufficient for rational action.

A speaker can also criticize a subject for acting with an inadequate epistemic grounding – for example:

- (6) How could you leave the kids home alone? You didn't know I would be home early.

In this sort of case, the subject lacks the knowledge that is (purportedly) necessary for rational action. Finally, a speaker can also excuse behavior – her own or someone else's – by pointing out that the person in question did not have the requisite knowledge. With respect to (5), for example, the speaker is closing out the possibility of criticism – like that made in (3) – by asserting that she did not have the knowledge she would have needed in order to make rational a different action (in this case, offering a different dish to her guest). Not only would the missing knowledge have been sufficient to rationalize a different action, it is also necessary to make rational the choice of a different course of action.

These common ways of talking about knowledge are offered as evidence for pragmatic encroachment because they seem to presuppose the truth of RKP. Why would we focus on what the agent knows or doesn't know? It makes sense to do so on the assumption that knowledge is sufficient for practical rationality.⁵

1b Knowledge as a safe reason

Jeremy Fantl and Matthew McGrath offer an argument for pragmatic encroachment that is grounded in the way we conduct practical reasoning. One of the two premises they employ is *Safe Reasons*:

If p is a reason you have to ϕ , then p is warranted enough to justify you in ϕ -ing, for any ϕ .⁶

The basic idea is that, when we weigh the reasons we have (for acting or for forming beliefs), we regard them as safe to rely upon. We do not take into account how probable they are – or, if we do come to consider their probability, we no longer regard them as reasons we can rely upon. As evidence for this claim, Fantl and McGrath point out that we never weigh against one another reasons like “the horse I bet on won the race,” and “there’s a chance that the newspaper reported the race results incorrectly and the horse didn’t win.” If I am trying to decide whether I can afford to eat a fancy dinner tonight, I simply take it as a fact that I have the money from betting on the winning horse. (Or, if I am genuinely concerned that the race report is incorrect, I don’t treat this as a reason at all.) I would never think, “Well, on the one hand, I have the money from winning at the track. On the other hand, I may not have that money, since there’s a chance that my horse didn’t really win. What should I do?” Instead, “[o]ne reason is defeated by another, and the defeat has nothing to do with how probable either reason is.”⁷ Fantl and McGrath describe this as a “ledger-keeping” picture of having reasons: when an agent has a reason, “it gets put in the ledger with countervailing reasons and *weighed* against them. But the probabilities of these reasons don’t get recorded alongside.”⁸

The other premise in their argument is the *Knowledge-Reasons* Principle, KR:

If you know that p , then p is warranted enough to be a reason you have to ϕ , for any ϕ .⁹

Not just any proposition can serve as a reason. It wouldn’t after all, be practically rational to buy a particular book simply because one has guessed that it will be assigned by one’s instructor. It makes sense to buy it only if one has the proper epistemic relationship to that proposition. What *is* the proper epistemic relationship? KR identifies knowledge as being good enough epistemically to provide us with reasons, both for acting and for forming beliefs.

Together, *Safe Reasons* and KR entail that, if you know that p , then p is warranted enough to justify you in ϕ -ing, for any ϕ .¹⁰ That is, knowledge is sufficient for practical rationality. Again, it is logically equivalent to say that being suitable for use in practical reasoning is necessary for knowledge.

1c High stakes and low stakes

Perhaps the most common way of defending pragmatic encroachment involves comparing two subjects who are identical in all truth-directed ways but differ in their practical situations. For example, suppose that Jamie hears the waiter say that the cake she has just been served does not contain any nuts. “Too bad,” she thinks, as the waiter returns to the kitchen. Jamie has no reason to distrust the waiter, and she comes to

believe what the waiter has said. Most philosophers would agree that the following assertion is true:

- (7) Jamie knows the cake does not contain nuts.¹¹

Sasha, who has been served a piece of the same cake, also overhears the waiter. But her reaction is quite different from Jamie's – Sasha is extremely allergic to nuts and could die if she eats any. Many, perhaps most, people think that Sasha should not simply start eating the cake. Given the risk of dying, the rational thing for her to do is to check with the waiter and have him confirm the recipe with the chef. But surely, the defenders of pragmatic encroachment say, this would be unnecessary if Sasha really *knew* the cake was nut-free. For that reason, they think, most speakers would not assert:

- (8) Sasha knows the cake does not contain nuts.

More strongly, some speakers may even deny (8) or assert its contradictory. Notice, however, that Sasha and Jamie are identical in all truth-directed ways. They have the same evidence that the cake does not contain nuts and, let us suppose, their relevant cognitive faculties – perception, memory, reason, and so on – are equally powerful and reliable. But they differ in that Sasha has much more at stake practically than Jamie does. Practical rationality seems to require Sasha (but not Jamie) to confirm that the cake doesn't have nuts in it. A natural explanation is that Sasha needs to put herself in a stronger position than Jamie's, with respect to truth-directed things like evidence, in order to have the knowledge that Jamie has. It is only when her epistemic position is strong enough, given her practical interests, that her belief counts as knowledge and she is permitted to treat it as a reason to act.

2 Objections to Pragmatic Encroachment

There are a variety of objections that can be raised to the thesis of pragmatic encroachment. Some offer counterexamples to RKP. Others argue that pragmatic encroachment fails to capture something essential about knowledge. And, finally, others argue that the view gives an inadequate or incorrect picture of practical rationality.

2a The variety of ways we criticize, justify, and excuse

Recall that argument 1a for pragmatic encroachment is grounded in our practice of using knowledge attributions to criticize, justify, and excuse actions. The thought was that there must be a unique and special link between knowledge and practical rationality to explain this sort of behavior. In fact, however, criticism, justification, and excusing are connected to a variety of epistemic properties, both weaker and stronger than knowledge. Here are some examples:

- (9) When I wrote my letter of recommendation, I had every reason to believe she would graduate with honors.

- (10) How could you go rock climbing on Widow's Peak without being completely certain that your gear was in good shape?
- (11) I've done a bit of research on this, so I am quite sure that vaccines are not linked to autism.

In some cases, we even defend or excuse our actions by talking about what we believed, without any explicit mention of the epistemic properties of those beliefs:

- (12) I believed him when he said he wasn't married.

Given this variety of ways in which we criticize, justify, and excuse behavior, there is no reason to think that knowledge bears a special relationship to practical rationality.

These data also give us reason to deny RKP, with respect to both the necessity and the sufficiency of knowledge for rational action. Take (9), for example. In defending his letter of recommendation, the professor is letting his interlocutor know that it was practically rational for him to write the letter, even though he didn't know that the student in question would graduate with honors. (He couldn't know this because, as it turned out, it was false.) In other words, his action was rational even though it was grounded in a reason that failed to be knowledge. Hence, knowledge wasn't necessary to make it rational.

Consideration of (10) makes clear that knowledge isn't always sufficient for rational action, either. Consistent with making this assertion, the speaker might grant that the climber knows her gear is in good shape. The problem is that mere knowledge – an epistemic status that falls short of certainty – might not be good enough when someone's life is at risk.

2b Crazy counterfactuals

One of useful things about epistemic concepts is that they allow us to think about how our cognitive position could be different from what it is. For example, if I have heard vague remarks about a presidential candidate's views on immigration, I can acquire knowledge of his views by reading newspaper articles about them. Doing so will allow me to gather relevant evidence. I can also recognize that, if I hadn't been watching television, I wouldn't have had any evidence at all about the candidate's views.

There is nothing at all surprising about this. It should be obvious that changing one's standing with respect to truth-directed things like evidence will also change the epistemic properties of one's beliefs. This is consistent with pragmatic encroachment. However, it is a consequence of RKP that changing one's practical interests will also change the epistemic properties of one's beliefs. The higher the stakes, the more justified one's belief will have to be to count as knowledge. For this reason, assertions like the following are true, given RKP:

- (13) If she really loved me, she wouldn't know that I cheated on her.
- (14) If I didn't have so much money riding on the game, I would know the Bears had won.

Defenders of pragmatic encroachment have acknowledged that counterfactuals like these will be true, if their theory is correct.¹² Being invested in a relationship or gambling

large sums of money make it the case, given RKP, that the subject has to do much more to have knowledge than someone who does not have anything at stake in the situation. This is a serious problem: not only is it extremely implausible that these counterfactuals are true, but it is also clear that people never assert anything like them. If pragmatic encroachment is motivated, in part, by paying careful attention to the way people use attributions of knowledge, then surely we need to consider the way they use counterfactual knowledge attributions, too.

2c Knowledge and high stakes

Perhaps the most straightforward objection to pragmatic encroachment is grounded in counterexamples to principles like RKP. Here is one such case:

I am taking part in a psychological study that tests the effects of stress on memory. I am asked a question: when was Julius Caesar born? If I give the correct answer, I get a jelly bean. If I give an incorrect answer, I am given a horrible electric shock. Nothing happens if I give no answer. I remember that Caesar was born in 100 BC, but I am not so sure of it that it is worth risking electrocution. Nevertheless, I quietly say to myself, "I know it's 100 BC."¹³

Notice that in this scenario it is plausible that I retain my knowledge even though the practical stakes are so high it isn't rational for me to act on it. If this is right, it is a case in which knowledge isn't sufficient for rational action. Notice, too, that it is natural for me to continue to attribute knowledge to myself, even when I recognize that I shouldn't treat it as a reason to act.¹⁴

2d The relativity of knowledge

The defenders of pragmatic encroachment sometimes claim that their view has an advantage in being able to fully capture the value or importance of knowledge.¹⁵ Showing how knowledge has an intimate link to practical rationality explains why it is fitting for knowledge (rather than, say, justification) to be the central focus of epistemology. But the sort of objection raised in section 2c can be extended to show that pragmatic encroachment forces us to view knowledge in a way that actually undermines our sense of its importance:

The psychological study changes so that I am playing two games at the same time. The first is as before: I'll get a jelly bean for a correct answer and a horrible electric shock for one that is incorrect. But, in the second game, I get \$1000 for a correct answer and only a gentle slap on the wrist for a wrong answer. In both games, nothing happens if I fail to answer, and I can take different strategies in the two games. Both games begin when I am asked: when was Julius Caesar born? I give no answer in the first game, but I answer "100 BC" in the second.¹⁶

In this scenario, I am faced with two decisions, both of which are *p*-dependent, where that *p* is the proposition that Julius Caesar was born in 100 BC. But I am rational to treat that *p* as a reason for only one of the decisions. If RKP is true, this means that I both know and don't know that *p* at the same time.¹⁷

The only way an object can have contradictory properties at the same time is if they are relativized – for example, to time or to place. In this case, I both have and lack the same knowledge in the same place and at the same time. It would seem that the only thing left to which they might be relativized is something like *practical situation*. This would be an unwelcome consequence for several reasons. First, we do not ordinarily talk as though knowledge is relativized in this way. Second, our account of how to reason with knowledge would become rather complicated, as we would have to be sure that pieces of knowledge combined in deliberation have not been relativized in incompatible ways. And, finally, knowledge would no longer seem to be the stable accomplishment it is usually thought to be. Not only would it “come and go with ease,” as Hawthorne suggests, it would also be confined to practical situations in a way that would make it difficult to rely upon.¹⁸

2e Safe reasons and the risk of error

Recall that Fantl and McGrath offer what they call a “ledger-keeping” picture of reasoning, according to which we weigh reasons against one another without taking into account their probabilities. This picture is suggested by the fact that we do not weigh that p against the chance that not- p . They conclude that our reasons are safe: they are warranted enough epistemically to justify us in anything we do or believe. They then argue that knowledge provides us with reasons that are safe in this sense.

If fallibilism is correct, however, knowledge is distinguished from certainty in that knowledge carries with it the risk of error.¹⁹ When I see a medium-sized black bird in a nearby tree, that’s good evidence for me to think it is a crow. A fallibilist will go further and say that the evidence is good enough for me to *know* that it’s a crow. Still, the evidence I have is logically compatible with it being something else – an escaped mynah, perhaps, or merely a figment of my imagination.

Ordinarily, we ignore the chance of error because it is quite low when we have knowledge. But there are times when it becomes relevant after all. When we combine many fallibly known premises in a single piece of reasoning, the risk of error rises until we would be irrational to continue ignoring it. This is something that Fantl and McGrath acknowledge; for this reason, they say, “no fallibilist can say that when you have a reason, you get to put it to work in *any* reasoning in which it figures.”²⁰ But this admission does not fit well with the ledger-keeping picture of reasoning. If knowledge really did provide safe reasons, then it *ought* to be something that we could put to work in any reasoning at all.

Fantl and McGrath are right to point out that we do not weigh against one another that p and the chance that not- p , but they draw the wrong conclusion from it. When we begin to worry about the chance that not- p , we stop treating that p as a reason. (We might substitute for it the probability that p , but that is a different fact.) So, it’s not that knowledge provides us with a reason that is always safe – rather, we can use knowledge as a reason only when it is safe to do so. As we saw in section 2c, there can be cases in which knowledge is not good enough to act on because the agent has stakes that are too high for this to be rational. To this we can also add cases in which the stakes are low but the risk of error has simply grown too great through the use of too many fallibly known premises. The upshot in either sort of case is the same: the thesis of pragmatic encroachment does not capture how we reason with knowledge because knowledge is not always good enough for practical rationality.

2f The Dutch book

If RKP were true, it would license some highly irrational behavior. Consider the following case:

I have a broker who is extremely reliable at picking stocks. She tells me that a biotech stock, BXD, is a good long-term investment and that she can move a fourth of my assets into BXD stock. Given her testimony, I know that it will go up in value, so I agree. An hour later, she tells me that she can now move another fourth of my assets into BXD. I know it will go up in value, so I agree again. An hour later, the same thing happens. When she calls me for the fourth time, she offers to move my remaining assets into BXD. But she also points out that I would then have all of my assets tied up in a single stock, which is a very risky thing to do. The stakes have become too high, and so it's not rational for me to buy more shares of BXD. Given RKP, this means that I no longer know that the stock will go up in value. I no longer have the knowledge that would permit me to keep the stock, so I tell my broker to sell all of it in favor of other investments I know to be safe. After an hour, she calls back to remind me that BXD is an excellent long-term investment. Having sold all my shares, this is no longer a high-stakes proposition for me. I reflect that she is reliable in her stock tips, and I again come to know that BXD will go up in value. So, I take her up on her offer to move a fourth of my assets into BXD. And so on.²¹

This is a bad deal for me. Because my broker charges a fee for every transaction, she has set up a Dutch book against me: my money is guaranteed to melt away.

Of course, no one would continue making trades once it has become clear that losses are guaranteed. The problem for pragmatic encroachment, though, is that RKP permits every step and never gives me a reason to stop. I can get out of the Dutch book only by ignoring RKP. That's good reason to think the principle provides us with an incorrect account of knowledge and practical reason.

2g Putting the cart before the horse

If RKP were true, it would mean that knowledge and action could interact with one another so that practical deliberation turns out to be incapable of reaching a stable end point. The agent's knowledge might license action that raises the stakes in such a way that the agent ceases to have the knowledge in question. The Dutch book objection shows this occurring over a sequence of events, but it can also happen in a way that prevents the agent from taking any action at all. Consider the following dialogue:

ME: If BXD will go up in value, I should invest in it. So, should I invest in it?
FRIEND: That depends – do you know it will go up in value?
ME: That depends – should I invest in it?

If I don't invest in BXD, I have nothing at stake relative to knowing whether it will go up in value. So, this is relatively easy for me to know. But, if I invest heavily in BXD, the stakes may go so high that I can't know its value will increase.

Presumably, the defenders of pragmatic encroachment put forward a principle like RKP because they think we can discover what the practically rational thing to do is by using knowledge in deliberation. But, if RKP were true, it would mean that the reverse

is true: we can't tell what our knowledge is until we have first determined what the practically rational thing to do is.²² But this is to put the cart before the horse. Moreover, it is hard to see how knowledge would retain any importance at all – once we have made an independent determination of what the practically rational thing to do is, knowledge has no role left to play.

3 An Alternative Proposal

Even if the above objections show that pragmatic encroachment cannot be correct, it doesn't yet mean that a traditional, truth-directed epistemology can provide a plausible account of knowledge and practical rationality. In particular, we would like to have an explanation of the intuitive data (from section 1c) regarding the different ways we sometimes talk about people who differ only with respect to their practical stakes.

Toward this end, notice that justification comes in degrees, ranging downward from perfect justification. If fallibilism is correct, there is some degree of justification lower than perfect justification that serves as a threshold for knowledge. In other words, knowledge is compatible with having a variety of degrees of justification, ranging from the threshold to perfect justification. Corresponding to these different degrees of justification, we can also say that there are different degrees of knowledge.

If we take knowledge to rest on truth-directed things like evidence or reliability, then Jamie and Sasha, the two subjects we compared in section 1c, both know that the cake doesn't have nuts in it. Because we reject RKP – we don't hold that knowledge is always sufficient for rational action – we can say that only Jamie is rational to act on her knowledge. Sasha needs a higher degree of knowledge, perhaps something approaching certainty, before it is rational for her to eat the cake.

That explains the behavior of both subjects. Advocates of pragmatic encroachment, however, will argue that the intuitive data include, not only how the subjects behave, but what we say about them. Ordinary speakers will assert (7) but not (8) – in fact, some people will even positively deny (8). On the view I have sketched, it seems that denial should be false. So, why would speakers say such a thing?

The simplest explanation is that “knows” is ambiguous: corresponding to the various degrees of knowledge are various meanings of the term. Each meaning is available for use in any conversational context; what the term ends up meaning on any occasion of use is determined only by the speaker's intentions. “Knows” is also strongly associated with a threshold usage, where the threshold can be set either at the minimal threshold to count as knowledge at all or at some higher degree of knowledge.²³ This threshold will also be set by the speaker's intentions alone, where the speaker will usually intend to talk about the degree of knowledge that is currently practically useful. Because this is the usual practice, one's interlocutors will expect one's knowledge attributions to be about only practically relevant degrees of knowledge. So, when a speaker denies (8), her interlocutor will naturally hear the speaker as denying that Sasha has a *high* degree of knowledge – not necessarily as denying that Sasha has knowledge of *any* degree.

It is worth noting that this expectation can be ignored; speakers can choose to talk about practically useless degrees of knowledge, too. Hence, in the psychological study case in section 2c, I could say – both to myself and as an unofficial side remark to the researcher – that I know the answer even though I don't want to make an official

response to the question. This, too, is an intuitive datum about how we think and talk about knowledge, and it is one that the defenders of pragmatic encroachment simply cannot accommodate.²⁴

Of course, much more could be said about these and other, related issues. But it should be clear enough that practical matters do not affect whether you know. The thesis of pragmatic encroachment simply faces too many debilitating objections to be correct. At most, then, practical matters affect what degree of knowledge it is rational to use in deliberation and it is most useful to talk about. Knowledge itself is still best understood in accordance with traditional, truth-directed epistemology.

Acknowledgments

For helpful discussions of these issues, I am grateful to Jessica Brown, Jeremy Fantl, Matthew McGrath, Ernest Sosa, Sarah Wright, and, most especially, Jennifer Lackey.

Notes

- 1 For "intellectualism," see Stanley (2005); for "purism," see Fantl and McGrath (2009). See also Fantl and McGrath (2002), Hawthorne (2004), and Hawthorne and Stanley (2008).
- 2 Fantl and McGrath (2009) characterize the view as *pragmatic encroachment*.
- 3 Hawthorne and Stanley (2008, p. 578).
- 4 Fantl and McGrath (2009, p. 66), defend a somewhat different principle, KJ: If you know that p , then p is warranted enough to justify you in ϕ -ing, for any ϕ . Notice that KJ takes knowledge to be sufficient but not necessary for rational action. Where it is relevant to an argument or objection, I will indicate this difference between RKP and KJ. Also, Fantl and McGrath do not limit their principle to choices that are p -dependent, as Hawthorne and Stanley do with RKP. As Jessica Brown notes, this means that the subject will need to have an extraordinarily high degree of epistemic justification for all of her beliefs; see Brown (2011, p. 170). It is hard to see how KJ, in its unrestricted form, would not lead to almost total skepticism.
- 5 For this sort of argument, see Fantl and McGrath (2007) and (2009, p. 63), and Hawthorne and Stanley (2008, pp. 571–574).
- 6 Fantl and McGrath (2009, p. 77).
- 7 Fantl and McGrath (2009, p. 79).
- 8 Fantl and McGrath (2009, p. 79).
- 9 Fantl and McGrath (2009, p. 69).
- 10 This is KJ (see footnote 4 above). Remember that it does *not* take knowledge to be necessary for practical rationality, unlike Hawthorne's and Stanley's RKP.
- 11 There is a complication here. Contextualists would say that the truth of (7) depends on the epistemic standards in play in the speaker's conversational context. Defenders of pragmatic encroachment have argued forcefully against contextualism and, though this debate is far from over, I shall simply follow them here in holding that the truth of knowledge attributions depends on what is happening with the subject of the attribution, not with the speaker.
- 12 See Hawthorne (2004, p. 177, n.40), and Fantl and McGrath (2009, pp. 208–212). Both Hawthorne and Fantl & McGrath suggest that there are also crazy counterfactuals that are true in virtue of the Gettier problem, so this is not a problem unique to pragmatic encroachment. Hawthorne, for example, says that the counterfactual (said of a person looking at a real barn but in an area where there are many barn facades), "If there weren't so many fake barns

- around, he would know that he's looking at a barn," sounds odd. I am inclined to say that it sounds *less* odd than the pragmatic encroachment counterfactuals. More than that, any discussion of the Gettier problem – couched in counterfactuals or not – will sound odd in ordinary discourse, given that it is a technical problem in epistemology. The same cannot be said about the pragmatic counterfactuals, which one would think ought to play a role in ordinary deliberation, given how frequently they would be useful if true.
- 13 This case is drawn from Reed (2010). See also Brown (2008).
 - 14 Fantl and McGrath concede that it might be natural to continue to attribute knowledge to oneself in this scenario, but they also argue that “it wouldn't raise eyebrows” if the subject denied that he had the knowledge in question (2009, p. 62). That may be true, but it doesn't change the fact that the case as described (in which the subject *can* continue to attribute knowledge to himself) is straightforwardly a counterexample to pragmatic encroachment.
 - 15 See, for example, Stanley (2005, p. vi), and Fantl and McGrath (2009, chapter 6).
 - 16 See Reed (2010).
 - 17 The situation is not quite as clear for Fantl's and McGrath's KJ, given that it does not take knowledge to be necessary for practical rationality. For more on the implications of this objection for KJ, see Reed (2012).
 - 18 Hawthorne (2004, p. 176).
 - 19 More precisely, if fallibilism is correct, knowledge carries with it the risk of being either false or true only by accident. See Reed (forthcoming a) and (2002).
 - 20 Fantl and McGrath (2009, p. 82).
 - 21 See Reed (2012).
 - 22 Recall that, if RKP is true, being suitable for use in practical reasoning is a necessary condition on knowledge.
 - 23 “Knows” has this threshold usage because we often are interested in whether someone knows *well enough*, given the practical circumstances.
 - 24 For a much more detailed presentation of this view, see Reed (forthcoming b).

References

- Brown, J. (2008) Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs* 42: 167–189.
- Brown, J. (2011) Fallibilism and the knowledge norm for assertion and practical reasoning. In *Assertion*, ed. J. Brown and H. Cappelen, 153–174. Oxford: Oxford University Press.
- Fantl, J. and McGrath, M. (2002) Evidence, pragmatics, and justification. *Philosophical Review* 111: 67–94.
- Fantl, J. and McGrath, M. (2007) On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research* 75: 558–589.
- Fantl, J. and McGrath, M. (2009) *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Hawthorne, J. (2004) *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Hawthorne, J. and Stanley, J. (2008) Knowledge and action. *The Journal of Philosophy* 105: 571–590.
- Reed, B. (forthcoming a) Fallibilism. *Philosophy Compass*.
- Reed, B. (forthcoming b) Fallibilism, epistemic possibility, and epistemic agency. *Philosophical Issues*.
- Reed, B. (2012) Resisting encroachment. *Philosophy and Phenomenological Research* 85: 466–473.
- Reed, B. (2010) A defense of stable invariantism. *Noûs* 44: 224–244.
- Reed, B. (2002) How to think about fallibilism. *Philosophical Studies* 107: 143–157.
- Stanley, J. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.