Knowledge and Practical Risk

Christoph Kelp
christoph.kelp@glasgow.ac.uk

Abstract. This paper develops a novel version of pragmatic encroachment. In a nutshell, the starting idea is that knowledge is about eliminating risk of error. While the idea that knowledge excludes epistemic risk already enjoys a fair amount of popularity in the literature, the present pragmatic encroachment theory argues that practical risk of error must be eliminated also. It is argued that the resulting version of pragmatic encroachment compares favourably with the perhaps two main rival views in the literature which develop pragmatic encroachment either by appeal to stakes-sensitivity or else a knowledge norm of reasons for action.

Keywords: Epistemology; knowledge; pragmatic encroachment; risk

1. Introduction

Purism about knowledge is the view that whether or not one knows that *p* does not depend on practical factors. It is the default view in epistemology. For the longest time, the mere thought that purism might be false would have been received with little more than ridicule.

However, recent epistemology has witnessed an important challenge to purism. In fact, there has been a growing number of researchers who hold that purism is false and that knowledge depends at least in part on practical factors.[1] In what follows, I will refer to these foes of purism as 'pragmatic encroachers' and as their view as 'pragmatic encroachment'.

The lion share of the literature on pragmatic encroachment focuses on the question whether some version of pragmatic encroachment is true rather than purism.[2] This paper does not make a contribution to this part of the literature. Given

[1] Examples include (Fantl and McGrath 2002; 2009; Hawthorne 2004; Hawthorne and Stanley 2008; Stanley 2005; Weatherson 2011; 2012).

[2] For arguments for pragmatic encroachment see (Fantl and McGrath 2002; 2009; Hawthorne 2004;

that I am a champion of purism myself, this may not come as much of a surprise. What I want to address an issue that has received less attention, to wit, what the best version of pragmatic encroachment looks like. Accordingly, for the purposes of this paper, I will simply assume that pragmatic encroachment is true (henceforth also 'the guiding assumption'), and focus on developing a novel version of it, which places pragmatic encroachment in the ambit of anti-risk epistemology (Section 1). I will also provide reason to think that this new kind of pragmatic encroachment compares favourably with extant versions of the view (Section 3).

2. Anti-Risk Pragmatic Encroachment

2.1 The View

According to anti-risk epistemology, knowledge is, in essence, non-risky true belief.[3] According to orthodox version of the view, what this idea amounts to is that knowledge is true belief that does not run too high a risk of *error*.

Hawthorne and Stanley 2008; Stanley 2005; Weatherson 2011; 2012). Important challenges to the view include (Brown 2008; 2018; Eaton and Pickavence 2015; Ichikawa, Jarvis, and Rubin 2012; Reed 2010; Roeber 2018; Zweber 2016). For a nice overview over the arguments for and against pragmatic encroachment see (Kim 2017).

[3] The idea that knowledge excludes luck is widely accepted in the literature and has been taken up in the pragmatic encroachment literature. For instance, (Kim 2019) develops a version of pragmatic encroachment that addresses issues surrounding epistemic luck and Gettier cases. Exactly how the present proposal relates to this kind of view is a fascinating question. Unfortunately, I will have to leave it for another occasion. For present purposes, I will rest content with pointing out that while the notions of luck and risk are closely connected, there is reason to think that they are different. Moreover, (Pritchard 2016) makes a nice case for parting with orthodoxy and move from anti-luck to anti-risk epistemology. The present paper adds to this case, especially for those who want to embrace pragmatic encroachment. (In a nutshell, this is because a key feature of risk, captured in Alternatives below, is not plausibly a feature of luck. At the same time it plays a key role in the present anti-risk version of pragmatic encroachment.) Finally, I'd like to emphasise that I will leave open whether there are further conditions on knowledge, besides truth, belief and the anti-risk condition (Pritchard Forthcoming). To keep things simple, I will make the idealising assumption that the anti-risk

One promising strategy for anti-risk epistemologists is to develop an account of risk and put it to use to derive an account of knowledge. I will not pursue this strategy here. Rather, what I rest content with giving a simple and informal modal model of when something is too risky according to which X puts one at too high a risk of Y if and only if there is a relevant possible world at which both X and Y obtain. For instance, going out into the rain (=X) puts one at too high a risk of catching a cold (=Y) if and only if there is a relevant possible world at which one goes out into the rain (=X) and catches a cold (=Y). Most importantly for present purposes, believing that p puts one at too high a risk of error if and only if there is a relevant possible world at which one falsely believes that p.

This gives us the following relevant alternatives model for knowledge:

> RA. One knows that p if and only if one truly believes that p and there is no relevant possible world at which one falsely believes that p (or, alternatively, one avoids error at all relevant possible worlds).[4]

The view that I am interested in exploring in this paper combines anti-risk epistemology with pragmatic encroachment about knowledge. According to pragmatic encroachment about knowledge, whether one knows that p depends at least in part on practical factors. To get there, note first that traditional (purist) versions of anti-risk epistemology take the risk of error that one only needs to eliminate in order to know to be a distinctively *epistemic* risk. In contrast, the view I want to explore here has it that in order to know one needs to eliminate a distinctively *practical* type of risk of error.

One key difference between epistemic and practical risk of error is that practical risk of error is analysed in terms of the risk of unsuccessful action in a way in which

---

condition is the only condition on knowledge, besides truth and belief.

[4] For classic defences of the relevant alternatives accounts of knowledge, see (Dretske 1970; Goldman 1976; Stine 1976). Note that I do not mean to advance the above as an *account* of knowledge. I want to leave open whether RA captures the ultimate nature of knowledge. All I want to commit to here is that RA provides a useful *model* for knowledge, which allows us to understand certain properties of knowledge, no matter whether or not our understanding thereby reaches all the way down to the ultimate nature of knowledge.

epistemic risk isn't: one's belief that p is running too high a practical risk of error if and only if acting on p puts one at too high a practical risk of failure because of error.

To capture the difference between the two kinds of risk in the model, let's first distinguish between two relevancy relations, epistemic and practical (henceforth also 'e-relevance' and 'p-relevance' for short). Here is the resulting 'pragmatic' relevant alternatives model for knowledge:

> PRA. One knows that p if and only if one truly believes that p and there is no e-relevant and no p-relevant possible world at which one falsely believes that p (or, alternatively, one avoids error at all e- and p-relevant possible worlds).

One important question that pragmatic encroachers will need to address concerns the relation between e-relevant and p-relevant possible worlds. According to the view that I want to explore here, these two are independent. In particular, whether or not one avoids error at e-relevant worlds does not directly depend on practical factors in the way in which whether or not one avoids error at p-relevant worlds does. Colloquially speaking, what this amounts to is that, according to the view I am interested in, there are separate epistemic and practical normative requirements on knowledge. In this way, knowledge turns out to be a normative hybrid.

I have said something about how the set of p-relevant worlds is fixed and I have noted that the set of e-relevant worlds is independent of the set of p-relevant worlds. One question that remains is how the set of e-relevant worlds is fixed; another, why we should think that there are these two independent conditions on knowledge, i.e. that knowledge is a normative hybrid. These are difficult questions. In fact, if no answer to them can be given, there is reason to doubt that the view I am interested is viable at all. After all, it is prima facie implausible that knowledge should have a hybrid structure. And this implausibility will only be compounded if it turns out that there is no satisfactory account of one of its key conditions. In what follows, I will try to show how both questions can be answered in one fell swoop, thereby addressing a key worry about the viability of the view developed here.

2.2 Hybridity

To begin with, I'd like to take a look at functional things. Many things have functions. Passwords, for instance, have a function, *viz.* to protect from unauthorised access. Functional things may be better or worse and they have good-making properties in virtue of which they are better or worse. There are better and worse passwords and

they have good-making properties in virtue of which they are better/worse. To keep things simple, let's make the somewhat idealised assumption that the quality of a password is proportionate to its degree of complexity.

With these points in play, let's turn to the question as to what a good password may look like. Of course, a good password is one that it complex enough. But how complex is complex enough?

One plausible answer is that how complex your password should be will depend on and vary with properties of the things it protects. To keep things simple once more, let's assume that the relevant property is importance of the things protected. Accordingly, it will make sense for you to have a more complex password for your bank account than for the dropbox folder containing your holiday photos. Let's call this *the variable condition*.

At the same time, many services that have a password protection feature have rules for passwords, which set a minimal standard for password quality. You may be required to use upper and lower case letters, numbers, special characters, a minimum number of characters, etc. Sometimes, these requirements are overkill for you in the sense that they require a much higher degree of complexity than your purposes require. Let's call this *the invariable condition*.

Why does it make sense to operate this minimal standard for password quality? One promising answer is that the password protected service can be used by a variety of different users who want to protect things with varying degrees of importance. And, of course, it makes perfect sense to operate a requirement on good passwords that is fixed by the requirements of what we may call the normal user.[5]

Finally, note there may well be cases in which your specific practical situation will require you to choose a password that exceeds the minimal standards imposed by the service provider. What you are storing is so important that a good password is one that is more complex than that.

What comes to light is that there are two independent conditions on good passwords. One condition is variable, the other is invariable. It makes sense to operate the invariable one when it makes sense to have a minimal standard for good passwords. This happens when a number of different users use the service provider. Moreover, I suggested that it makes sense to fix the minimal threshold by the

[5] I am deliberately leaving open the question as to what determines these requirements. All I need for present purposes is that normal userhood, whatever it may be, calls the shots for the minimal standards on admissible passwords.

requirements of the normal user. At the same time, this does not mean that we should now abandon the variable condition. On the contrary, it makes perfect sense to keep this one as well. After all, there may be cases in which what you are trying to protect is so important that a good password will be more complex than what's required by the minimal requirements of the service provider.

Now, what I want to suggest is a parallel story about belief and knowledge. Just as with passwords, beliefs have a function, i.e. to guide us to successful action. They can be better or worse at this and have good-making features. In line with the kind of anti-risk epistemology that I am defending, I will also suppose that non-riskiness is one such good-making property of beliefs. To keep things simple, let's assume it's the only good-making property of beliefs. The less risky your belief, the better. Or, in terms of our relevant alternatives model, the more possible worlds you avoid error at, the better your belief.

With these points in play, let's look at what a good belief looks like. A good belief is one that is sufficiently non-risky. In terms of the relevant alternatives model, a good belief is such that you avoid error at enough possible worlds, i.e. at the relevant possible worlds. But what does this amount to?

One plausible answer for pragmatic encroachers about knowledge is that which possible worlds are relevant will depend on and vary with practical factors relating to the actions your belief may guide. Since these practical factors may vary, this will give us a variable condition on good belief. It is captured by the p-relevant possible worlds in the relevant alternatives model.

What about the invariant condition on good belief that the kind of view I am developing here also countenances? Why does it make sense to operate such a condition? The answer, I suggest, has again to do with social factors. We want to share our good beliefs with others. Just as with passwords, in order to achieve this, it makes sense to operate a minimal standard for belief quality (cf. Craig 1990). In terms of the relevant alternatives model, it makes sense to have an invariant set of relevant possible worlds. And in further analogy with the password case, I'd like to suggest that it is the practical requirements of what we may call the normal believer that fixes the minimal threshold.[6] In terms of the model, it is the practical requirements of the normal believer that fixes the invariant set of relevant alternatives. This gives us an

[6] For a similar proposal see (Sosa 2015). Again, I am leaving open what exactly these requirements amount to.

invariable condition on good belief.[7] It is captured by the e-relevant possible worlds in the relevant alternatives model.

These considerations serve to explain how the set of e-relevant possible worlds is fixed, i.e. in accordance with the practical requirements of normal believers. They also offer one answer to the difficult question of why we should think that good belief has a hybrid structure, with independent epistemic and practical conditions. And, last but not least, I'd like to suggest to identify good belief and knowledge.[8] In this way, we also have one answer to the difficult questions of why we should think knowledge has a hybrid structure and how the set of e-relevant possible worlds are fixed.

2.3 Risk

In what follows, I will further investigate risks. In particular, I want to investigate certain properties of X putting one at too high a risk of Y. My ultimate aim is to get a clearer account of how the requirements for knowledge given by the variable condition vary; or, in other words, how the set of p-relevant worlds expands and contracts.

To begin with, I'd like to characterise a risk for one as a *looming bad* for one. I want this to be a pretty lightweight characterisation of what a risk is which allows for a number of different ways of fleshing out what a *looming bad* amounts to, which, in turn, will lead to more substantive accounts of what a risk is.[9]

---

[7] Note also that the motivation of the invariant condition fits nicely with the general practicalist agenda.

[8] For an argument of this claim, see [Author 2016].

[9] The orthodox account construes risk as expected disvalue, where this is unpacked along standard decision theoretic lines. Recent literature features some noteworthy alternatives to the orthodox account, including the modal account (Pritchard 2015) and the normic account (Ebert, Smith, and Durbach Forthcoming). Note also that that there are different types of risks: moral, practical, etc. On the present proposal, risks are typed in accordance with the type of looming bad. For instance, if the looming bad is a practical bad, then the risk is a practical risk; if the bad is a moral bad, then the risk is a moral risk and so on. Note also that the risk I am mostly concerned with here is plausibly a distinctively *practical* risk. As a result, I will henceforth take it as read that these principles are interpreted accordingly. (In general, I will take 'risk' to mean 'practical risk' unless otherwise noted.)

With this point in play, let's move on to what it is involved in X putting one at too high a risk of Y. Here are two conditions. I hasten to add that there may well be others.

First, whether X puts one at too high a risk of Y may vary with degree of significance or 'what's at stake' for one. Say that we are both considering telling the third person at the table that she is holding the fork in the wrong way. Say that if we go ahead and tell her, we will be running the risk of thereby offending her. This risk may be significant for you but not for me, say because the person is your boss but not mine. Whether telling her is too risky may vary accordingly: it may put you at too great a risk of offending her because the risk is very significant (and so a lot is at stake) for you, but not for me. Whether X puts one at too great a risk of Y depends on *the degree of significance* of avoiding Y.

Second, there may be alternatives to X available to one. You are at a point where you have to decide whether or not to cross an old rickety bridge over a ravine. Now consider two cases: in the first, there is a new and very solid bridge just a few minutes down the trail; in the second, your only alternative is attempting an eight-meter jump. Whether crossing the dodgy bridge puts you at too high a risk of dying may vary with what alternative courses of action are available to you and how risky they are: it may be that in the first but not in the second case crossing the rickety bridge is too risky. Whether X puts one at too great a risk of Y depends on what alternatives to X are available to one and the degree to which they put one at risk of Y.

Here is how these properties of too risky X can be (partially) captured in the relevant alternatives model of risk. How wide the range of relevant possible worlds is depends on and varies with (i) what's at stake for one and (ii) what alternatives one has and how risky they are. Or, to be a little more precise:

> STAKES. If all else is equal[10] between two cases, α and β, and if the stakes are lower in α than in β, the set of relevant possible worlds in α is a proper subset of the set of relevant possible worlds in β.

[10] The all-else-equal-proviso is important. It tells us to look at pairs of cases in which *only* the stakes differ and everything else is the same. For instance, consider two cases, A and B, such that in A you have the flu and in B you are tortured. Although the stakes of having the flu in A are lower than the stakes of being tortured in B, STAKES does not entail that the set of possible worlds at which I'm

ALTERNATIVES. If all else is equal between two cases, α and β, and if in α there is an alternative to X available to one that puts one at a lower risk of Y, whereas in β there isn't, then the set of relevant possible worlds in β is a proper subset of the set of relevant possible worlds in α.

3. The Competition

3.1 Stakes-Sensitive Pragmatic Encroachment

The form of pragmatic encroachment I have sketched above develops the view within the ambit of an anti-risk epistemology. One consequence of this is that there are two dimensions along which practical factors encroach on knowledge, captured in STAKES and ALTERNATIVES. In this way, the view differs from versions of pragmatic encroachment that only countenance stakes-sensitivity (henceforth also 'stakes-sensitive pragmatic encroachment'), which we may think of as embracing PRA and STAKES but not ALTERNATIVES. Why is the anti-risk alternative is better?

To answer this question, let's first consider one important motivation for pragmatic encroachment, which is provided by cases like the following:

LOW. DeRose and his wife are driving home on a Friday afternoon. They are considering whether to stop by the bank on the way home in order to deposit their paycheques. It is not important that they do so. As they approach the bank, they notice that the lines are long and DeRose's wife wonders whether the bank will be open tomorrow and whether it might be an option to go then. Recalling his visit to the bank two weeks ago, he says, 'I know that the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning.'

HIGH COST. DeRose and his wife are driving home on a Friday afternoon. They are considering whether to stop by the bank on the way home in order to deposit their paycheques. As they approach the bank and notice the long lines, DeRose's wife receives an email stating that, due to a bureaucratic quirk, they will have to pay a fine unless they deposit their paycheques before Monday. When DeRose notes that he was at the bank on a Saturday only two weeks ago, his wife points

tortured includes a set of worlds at which I have the flu. This is because it's not only the stakes that differ between A and B. Rather, in A you have the flu and in B you are tortured.

> out that banks sometimes change their hours. In response, DeRose says, 'You're right. I don't know that the bank will be open tomorrow.'[11]

Both DeRose's attribution of knowledge that the bank will be open tomorrow in LOW and his denial of knowledge in HIGH COST are natural and intuitively correct. In this way, we have evidence for their truth (DeRose 2009).

An attractive feature of pragmatic encroachment is that it has the resources to accommodate this result. To see how, let's return to the modal model. Roughly, the idea is that, in LOW, DeRose has a true belief that the bank will be open tomorrow and avoids error at all e- and p-relevant worlds. By PRA, DeRose knows that the bank will be open tomorrow in LOW. At the same time, since the stakes in HIGH COST are higher, by STAKES, the set of p-relevant possible worlds expands. The thought is that it includes a world at which the bank changed its opening hours. If so, DeRose doesn't avoid error at all p-relevant possible worlds in HIGH COST and so doesn't know.

The cases can be tidied up such that the modal model entails straightforwardly that DeRose knows in LOW but not in HIGH COST. The following assumptions will do the trick. First, DeRose has what I will call *minimal knowledge* that the bank is open in LOW. This means that while he avoids error at all possible worlds that are e- and p-relevant in LOW, these possible worlds are *the only* possible worlds at which DeRose avoids error. Second, DeRose has what I will call *minimal ignorance* that the bank is open in HIGH COST. This means that while he does not avoid error at all possible worlds that are (e- and) p-relevant in HIGH COST, he avoids error at *every proper subset* of the possible worlds that (e- and) p-relevant.

Given that DeRose has minimal knowledge that the bank is open in LOW, by stakes-sensitive pragmatic encroachment, in any case that differs from LOW in that the stakes are higher, DeRose doesn't know that the bank will be open tomorrow. After all, by STAKES, the set of p-relevant possible worlds in LOW will be a proper subset of the set of p-relevant possible worlds in HIGH COST. At the same time, since DeRose has minimal knowledge that the bank will be open tomorrow in LOW, it follows that there will be some relevant possible world at which DeRose doesn't avoid error. Since HIGH COST differs from LOW in that the stakes are higher, stakes-

[11] For the original statement of the bank cases, see (DeRose 1992). (Stanley 2005) famously uses versions of bank cases to argue for pragmatic encroachment. For some recent experimental work on bank cases see (Buckwalter 2010; Buckwalter and Schaffer 2015; May et al. 2010; Pinillos 2012; Sripada and Stanley 2012).

sensitive pragmatic encroachment will deliver the result that DeRose doesn't know that the bank will be open tomorrow in HIGH COST. Of course, since what is key to the pragmatic encroachment account of this pair of cases is STAKES, LOW and HIGH COST can be handled equally well by stakes-sensitive and anti-risk pragmatic encroachment.

Consider next the following additional case:

> HIGH BENEFIT. DeRose and his wife are driving home on a Friday afternoon. They are considering whether to stop by the bank on the way home in order to deposit their paycheques. As they approach the bank and notice the long lines, DeRose's wife receives an email stating that, due to a promotion, clients that who are fortunate enough to have a local branch that is open on Saturdays and who will deposit a paycheque this Saturday at it will receive a cash prize. Naturally DeRose's wife wonders whether their local branch is open tomorrow. Recalling his visit to the bank two weeks ago, he says, 'I know that the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning.'[12]

Recall that in LOW, DeRose has minimal knowledge that the bank will be open tomorrow. Moreover, since, in HIGH COST, the stakes are higher than in LOW, DeRose doesn't know in HIGH COST. But now note that the stakes in HIGH BENEFIT are also higher than in LOW. As a result, according to stakes-sensitive pragmatic encroachment, DeRose doesn't know in HIGH BENEFIT either. Crucially, however, knowledge can naturally and intuitively correctly be *attributed* in HIGH BENEFIT. Stakes-sensitive pragmatic encroachment cannot accommodate this intuition.

One striking difference between the cases is that in HIGH COST a threat is looming, whereas in HIGH BENEFIT a reward is in the offing. This raises the question as to whether it is a difference between rewards and threats that explains why stakes-sensitive pragmatic encroachment falls short. In fact, it might be thought that at least one of the most promising versions of stakes-sensitive pragmatic encroachment encodes exactly this distinction, albeit without explicitly considering cases like HIGH COST and HIGH BENEFIT. Jason Stanley explicitly says that what determines how hard it is to know is "the cost of being wrong" (2005: 93). In terms of our modal model, it might be thought that STAKES is to be replaced by the following:

[12] To be maximally clear: this case is meant to be like LOW with an added benefit if the bank is open on Saturday. For a structurally similar case see (Anderson and Hawthorne 2019).

> COST OF ERROR. If all else is equal between two cases, α and β, and if the cost of being wrong are lower in α than in β, the set of relevant possible worlds in α is a proper subset of the set of relevant possible worlds in β.

Unlike STAKES, COST OF ERROR can accommodate the intuitive presence of knowledge in HIGH BENEFIT. After all, the cost of being wrong remains unchanged between LOW and HIGH BENEFIT. Accordingly, there is no reason to think that there is any difference in the set of p-relevant possible worlds between the two cases. But if there isn't, since DeRose knows in LOW, he will also know in HIGH BENEFIT.

Even if this line may look appealing at first glance, there is reason to think that, on reflection, it won't work. To see why not, consider LOW alongside the following case:

> HIGH COST*. DeRose and his wife are driving home on a Friday afternoon. They are considering whether to stop by the bank on the way home in order to deposit their paycheques. As they approach the bank and notice the long lines, DeRose's wife receives an email stating that, due to a bureaucratic quirk, they will have to pay a fine unless they do so this Saturday. The email also states that if their local branch is not open this Saturday, there is no way to avoid the fine. Naturally DeRose's wife wonders whether their local branch is open tomorrow. Recalling his visit to the bank two weeks ago, he says, 'I know that the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning.'

To begin with, note that we find a natural and intuitively correct *attribution* of knowledge in HIGH COST*. At the same time, recall that DeRose has minimal knowledge that the bank will be open tomorrow in LOW. Since in HIGH COST* the stakes are higher than in LOW and they are higher in virtue of the cost of error being higher, the set of p-relevant possible worlds expands, no matter whether we adopt STAKES or COST OF ERROR. Thus, stakes sensitive pragmatic encroachment makes the wrong prediction here. As a result, there is little reason to think that the difference between DeRose's original cases and the above variations is explained by a difference in reward vs. threat.

Can anti-risk pragmatic encroachment do better? There is reason to think that the answer is yes. Since the story for LOW and HIGH COST is the same as for stakes-sensitive invariantism, I will not return to it. Instead, I want to look at the story for HIGH COST, HIGH BENEFIT and HIGH COST*, focusing in particular on HIGH COST and HIGH COST*.

Recall, first, the definition of practical risk of error: one's belief that p is running too high a practical risk of error if and only if acting on p puts one at too high a practical risk of failure because of error. In HIGH COST and HIGH COST*, DeRose aims to avoid having to pay a fine and acting on his belief that the bank will be open tomorrow fails if, as a result of so acting, he ends up having to pay a fine. Note, next, that, in HIGH COST, there is alternative course of action available to DeRose with a substantially lower risk of failure because of error, i.e. going today instead. The same does not hold for HIGH COST*. After all, here any alternative to acting on his belief that the bank will be open tomorrow, including e.g. going today instead, will guarantee that DeRose will be stuck with a fine. By ALTERNATIVES, then, the set of p-relevant worlds in HIGH COST* is a proper subset of the set of p-relevant alternatives in HIGH COST. But now recall that we are assuming that DeRose has minimal ignorance in HIGH COST. That is to say that while he does not avoid error at all possible worlds that are (e- and) p-relevant in HIGH COST, he avoids error at every proper subset of the possible worlds that (e- and) p-relevant. It follows that, in HIGH COST*, DeRose avoids error at all e- and p-relevant worlds and so knows that the bank will be open tomorrow. In this way, anti-risk pragmatic encroachment can improve on stakes-sensitive pragmatic encroachment in that it has the resources to accommodate a broader range of linguistic data.

3.2 Knowledge Norm Pragmatic Encroachment

Another form of pragmatic encroachment connects knowledge with norms of action or practical reasoning. I do not mean to settle the issue of what the best way of stating the relevant norm is here. Rather, I will work with the following version of the norm from Hawthorne and Stanley:

> KN. Where one's choice is p-dependent, it is practically[13] permissible for one to treat p as a reason for acting if and only if one knows that p. (Hawthorne and Stanley 2008, 578)

[13] Early statements of the norm, including Hawthorne and Stanley's do not make explicit that the norm is a practical norm. More recently, there has been more clarity on this (e.g. Anderson and Hawthorne 2019). But, being a *knowledge* norm, isn't KN an *epistemic* norm really? Sometimes the term 'epistemic norm' is used such that KN qualifies as an epistemic norm. However, what is meant by 'epistemic norm', then, is *norm with epistemic content*. This leaves open what kind of norm (practical, moral,

Note that 'p-dependent' is a technical term, which is unpacked as follows:

> "[A] choice between options x1...xn is p dependent iff the most preferable of x1....xn conditional on the proposition that p is not the same as the most preferable of x1...xn conditional on the proposition that not-p." (Hawthorne and Stanley 2008, 578)

It is easy to see that KN supports pragmatic encroachment. After all, whether one is practically permitted to treat p as a reason for acting depends on and may vary only with differences in practical factors. By KN, whether one knows that p will depend on and may vary only with differences in practical factors. In this way, whether one knows depends on practical factors, which is just what the pragmatic encroacher would have us think.

Knowledge norm pragmatic encroachment can deal nicely with cases like LOW and HIGH COST. While in LOW, it is practically entirely permissible for DeRose to treat the proposition that the bank will be open tomorrow as a reason for action, in HIGH COST, it isn't. Since the choice between going today and going tomorrow is p-dependent on this proposition, knowledge norm pragmatic encroachment predicts, correctly, that DeRose knows in LOW but not in HIGH COST. Note that what is going on here is that differences in DeRose's practical situation drive up the epistemic requirements for DeRose to practically permissibly treat the proposition that the bank

---

epistemic, etc.) we are talking about. It is important to see that there is a difference between epistemic norms on the one hand and moral, practical, epistemic etc. norms with epistemic content on the other (Simion 2018). Once this distinction is clear, it is hard to deny that KN must be a practical norm (with epistemic content) rather than an epistemic norm (with epistemic content). There are at least two reasons for this. First, if it were an epistemic norm, it is difficult to see how the norm could lend support to pragmatic encroachment without begging the question. Second, as (Simion 2018) convincingly argues, action does not have a characteristic epistemic aim and, as a result, there can be no epistemic norm for action. It is easy to see that the same holds for treating something as a reason for action. For present purposes, I take these considerations to provide sufficient evidence that we will be well advised not to take KN to be an epistemic norm. Accordingly, I will set any such construals aside in what follows and assume, in line with more recent statements, that KN is a practical norm.

is open tomorrow as a reason for action. By KN, differences in DeRose's practical situation drive up the epistemic requirements for knowledge. In this way, differences in one's practical situation can make it harder for one to know, as it were. That's the good news.

Let's move on to the bad news, then. Given that changes in one's practical situation can make it harder to know, shouldn't we expect that changes in one's practical situation can also make it *easier* to know? In particular, in LOW and HIGH COST, it is plausibly the increased practical cost of treating the target proposition as a reason for action and failing that drives up the epistemic requirements for practically permissibly treating the target proposition as a reason for action and so makes it harder to know. But then shouldn't we expect that if we increase the practical benefit of treating p as a reason for action and succeeding, this will drive down the epistemic requirements for practically permissibly treating the target proposition as a reason for action and so makes it easier to know? In fact, shouldn't we expect that we could increase the practical benefits to the point that treating p as a reason for action is practically permissible even when all one has is a gettiered or perhaps even a mere true belief that p? If so, by KN, it follows that one can know in virtue of having a gettiered or even a mere true belief. But, of course, we know that this couldn't be the case. There is a minimal epistemic requirement on knowledge and it's considerably more demanding than this. Champions of knowledge norm pragmatic encroachment need to explain how this can be on their view. In this way, the view faces an important challenge.[14]

It is easy to see that anti-risk pragmatic encroachment rise to this challenge. While anti-risk pragmatic encroachment features a practical normative requirement on knowledge, it also countenances an additional epistemic normative requirement. Of course, it is exactly this requirement that rules out the possibility of knowing in virtue of having a mere true belief. And, more generally, this requirement places just the kind of minimal epistemic condition on knowledge that will prevent practical factors from making knowledge too easy to come by.

[14] Don't Hawthorne and Stanley have a response to this? After all, they claim that in cases in which it may look as if additional benefits make it easier to practically permissibly treat p as a reason for action, there are other propositions, about the chance of p, that one knows to be true and can practically permissibly treat as a reason for action instead. Ultimately, as I will argue in due course (fn.16), this won't work.

Another problem for knowledge norm pragmatic encroachment is that the normativity of practical permissibility and impermissibility differs importantly form the normativity of knowledge. In particular, whether something is practically permissible or impermissible is sensitive to its consequence profile in a way in which whether or not one knows isn't. Since just how tightly KN ties knowledge to practical permissibility, it will come as no surprise that KN runs into trouble. And, of course, this means further bad news for knowledge norm pragmatic encroachers.

The perhaps easiest way to bring this normative difference into sharp relief is by looking at some cases.[15] Consider, for instance, the following bank case:

> HIGH BENEFIT*. DeRose and his wife prefer to get their paycheques deposited before the weekend. Not doing so will lead to a minor inconvenience. They have to choose between going to the bank today, a Friday, or tomorrow. DeRose is certain that the bank is open today and has a mere true belief that the bank will be open tomorrow. Thanks to the intervention of a capricious god, he and his wife will get a place in heaven if and only if (i) they deposit their paycheques tomorrow and (ii) they treat the proposition that the bank will be open tomorrow as a reason for action.[16]

[15] See (Brown 2008; Reed 2010) for further cases that problematise KN and (Fantl and McGrath 2009) for a response. There are two important differences between Brown and Reed's argument and the present one. The first is that the weight of the present argument is carried by the theoretical point about the differences in normative profiles between knowledge and practical permissibility. The cases are merely meant to illustrate these differences. The second is that the bad consequences attach to treating the proposition as a premise in practical reasoning rather than to being wrong. As a result, a response along the lines suggested by Fantl and McGrath is less promising for the present cases. In a nutshell, they argue that modification of irrelevant features of the case lead to different intuitions about the presence of knowledge. While there is reason for thinking that this may indeed work when the cases turn on the bad consequences that attach to being wrong, it is far from clear that it will be equally plausible if the bad consequences attach to treating a proposition as a reason for action.

[16] A version of this case also serves to show why Hawthorne and Stanley cannot hope to meet the challenge of explain why KN doesn't allow for knowledge to come to cheaply by appealing to knowledge of chances (fn.14). Suppose that the capricious god will not only reward DeRose and his

Note, first, that there can be no question that it is practically permissible for DeRose to treat the proposition that the bank will be open tomorrow as a reason for action. Crucially, the reason for this has to do with the consequence profile of treating the proposition that the bank will be open tomorrow as a reason for action: the enormous benefits (a place in heaven) and comparatively little cost (a minor inconvenience) that are attached to this.

At the same time, it is widely recognised that knowledge is not sensitive to the consequence profile of believing in a similar way: whether or not one knows is not sensitive to benefits or costs (epistemic or other) down the line in this way. For instance, the fact that a belief, even a true one, that the bank will be open tomorrow has enormous practical benefits and very little costs can never be enough to turn it into knowledge.

Given these important differences in normative profiles, it is easy to enough to generate trouble for KN. To see this, note that, in HIGH BENEFIT*, DeRose's choice is p-dependent: conditional on the bank's not being open tomorrow, DeRose and his wife will prefer to go today. After all, they prefer to get the paycheques deposited before the weekend and there will be no chance to secure a place in heaven. But since DeRose's choice is p-dependent, champions of KN run into a dilemma: they can either grant that it is practically permissible for DeRose to treat the proposition that the bank will be open tomorrow as a reason for action. In that case, however, they must also accept that DeRose's knows that the bank will be open tomorrow, despite the fact that all he has is a mere true belief. Alternatively, they can grant that DeRose doesn't know that the bank will be open tomorrow. In that case, they must also accept that it is not practically permissible for DeRose to treat the proposition that the bank will be open tomorrow as a reason for action, despite the enormous benefits and little costs that attach to so doing.

---

wife for treating the proposition that the bank will be open tomorrow as a reason for action, he will also punish them for treating any other proposition as a reason for action. In this case, there is no other proposition that one can practically permissibly treat as a reason for action. Still, it's clearly practically permissible for DeRose and his wife to treat the proposition that the bank will be open tomorrow as a reason for action. As a result the attempt to explain why KN doesn't allow for knowledge to come to cheaply by appealing to knowledge of chances is bound to fail.

Cases like HIGH BENEFIT* cause trouble for the left-to-right direction of KN. It is worth noting that similar worries can be raise for the converse.

> HIGH COST**. Consider proposition that you are absolutely certain of, that you exist, say. Now suppose that you are facing some choice or other. At the same time, there is a capricious demon who will condemn you to an eternity of hell if you treat the proposition that you exist as a reason for acting.

Clearly, it is not practically permissible for you to treat the proposition that you exist as a reason for acting. Again, this has to do with the consequence profile of so doing: it is the enormous costs and few benefits attached to treating the proposition that you exist as a reason for action that make it impermissible.

At the same time, knowledge is not sensitive to the consequence profile in this way. That believing that you exist comes with devastating costs does not mean that you don't know it to be true. On the contrary, that you could not know what you are absolutely certain is no less plausible than that you could know what you merely truly believe.

It is easy enough to see that the differences in normative profiles once again cause trouble. To see how, note first that any choice you make is p-dependent on the proposition that you exist. After all, conditional on you not existing, you have no preferences. Hence, for any choice, the most preferable option conditional on you existing is not the same as the most preferable option conditional on you not existing. But since your choice is p-dependent, champions of KN once again face a dilemma: they can either grant that it isn't practically permissible to treat the proposition that you exist as a reason for action and accept that you don't know what you are absolutely certain of. Or they can grant that you know that you exist and accept that it is practically permissible for you to act on this, despite the fact this means an eternity of hell and precious little in the way of benefits for you.

It comes to light that KN runs into trouble. Given that KN is the central thesis of knowledge norm pragmatic encroachment, so does this version of the view. Fortunately, there is reason to think that anti-risk pragmatic encroachment can do better here, too. Cases like HIGH BENEFIT*, which put pressure on the left-to-right direction of KN, are cases in which knowledge threatens to come too easily. Anti-risk pragmatic encroachment can avoid this worry because it operates an independent epistemic normative requirement on knowledge, which ensures that knowledge cannot be too easy to come by. This leaves cases that threaten the right-to-left direction of KN such as HIGH COST**. Note that here you are absolutely certain that you exist

and so you don't run any practical risk of error. At the same time, due to the enormous looming costs, it is not permissible for you to treat the proposition that you exist as a reason for action. But then, of course, practical risk of error cannot be sensitive to consequences in the way that practical permissibility is. And, of course, that's once more good news for anti-risk pragmatic encroachment.

References

Brown, J. 2008. "Subject-Sensitive Invariantism and the Knowledge Norm of Practical Reasoning." *Noûs* 42: 167–89.

———. 2018. *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press.

Buckwalter, W., and J. Schaffer. 2015. "Knowledge, Stakes and Mistakes." *Noûs* 49: 201–34.

DeRose, K. 2009. *The Case for Contextualism*. Oxford: Clarendon Press.

Dretske, F. 1970. "Epistemic Operators." *Journal of Philosophy* 67: 1007–23.

Eaton, D., and T. Pickavence. 2015. "Evidence against Pragmatic Encroachment." *Philosophical Studies* 172: 3135–43.

Ebert, P., M. Smith, and I. Durbach. Forthcoming. "Varieties of Risk." *Philosophy and Phenomenological Research*.

Fantl, J., and M. McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Goldman, A. 1976. "Discrimination and Perceptual Knowledge." *Journal of Philosophy* 73: 771–91.

Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hawthorne, J., and J. Stanley. 2008. "Knowledge and Action." *The Journal of Philosophy* 105: 571–90.

Ichikawa, J., B. Jarvis, and K. Rubin. 2012. "Pragmatic Encroachment and Belief-Desire Psychology." *Analytic Philosophy* 53: 327–43.

Kim, B. 2017. "Pragmatic Encroachment in Epistemology." *Philosophy Compass* 12: e12415.

———. 2019. "An Externalist Decision Theory for a Pragmatic Epistemology." In *Pragmatic Encroachment in Epistemology*, edited by B. Kim and M. McGrath. London: Routledge.

Pinillos, A. 2012. "Knowledge, Experiments and Practical Interests." In *Knowledge Ascriptions*, edited by J. Brown and M. Gerken. Oxford: Oxford University Press.

Pritchard, D. 2015. "Risk." *Metaphilosophy* 46: 436–61.

———. 2016. "Epistemic Risk." *Journal of Philosophy* 113: 550–71.

Reed, B. 2010. "A Defense of Stable Invariantism." *Noûs* 44: 224–44.

Roeber, B. 2018. "The Pragmatic Encroachment Debate." *Noûs* 52: 171–95.

Simion, M. 2018. "No Epistemic Norm for Action." *American Philosophical Quarterly* 55: 231–38.

Sripada, C., and J. Stanley. 2012. "Empirical Tests of Interest-Relative Invariantism." *Episteme* 9: 3–26.

Stanley, J. 2005. *Knowledge and Practical Interest*. Oxford: Oxford University Press.

Weatherson, B. 2011. "Defending Interest-relative Invariantism." *Logos and Episteme* 2: 591–604.

Zweber, A. 2016. "Fallibilism, Closure, and Pragmatic Encroachment." *Philosophical Studies* 173: 2745–57.