

1

Putting Fallibilism to Work

Charity Anderson

1 Introduction

A connection between knowledge and reasons for action is defended in recent literature as follows: when one knows p one is in a good enough epistemic position to treat p as a reason for action (hereafter, *Sufficiency*).¹ *Sufficiency*—or some nuanced version of it—is used to motivate pragmatic encroachment: the view that pragmatic factors are relevant to whether a subject has knowledge.² When combined with fallibilism—the widely accepted idea that knowledge is compatible with an epistemic chance of error—*Sufficiency* results in a rather counterintuitive picture.³ Namely, it results in the rejection of purism, the view that pragmatic factors are irrelevant to knowledge.⁴ Fallibilism, purism, and *Sufficiency* each have substantial *prima facie* intuitive support; and yet the three seem to form an inconsistent triad.

To see why these three have been thought to form a trilemma, consider two agents with the same rather strong epistemic position for a proposition. Suppose the practical details are such that one agent ought to act, but the other ought not

¹ A knowledge-reasons connection has also been defended in terms of a necessary principle: one is appropriate to treat p as a reason for action only if one knows p . This chapter is neutral with respect to this principle.

² See Fantl and McGrath (2002, 2007, 2009), Hawthorne (2004), Stanley (2005), and Hawthorne and Stanley (2008).

³ Fallibilism is a technical term and although there is no agreed-upon definition, there is a core idea that is often expressed in terms of the compatibility of knowledge with a chance of error. In Section 3, I offer a gloss on the notion of epistemic chance of error that provides a way to understand fallibilism while also maintaining a knowledge-first framework.

⁴ Also known as *intellectualism*, this view is sometimes characterized as the idea that two subjects with the same strength of epistemic position for p are in the same position to know p . It is not entirely clear which factors count as ‘epistemic’ and which do not, but one point of agreement that tends to guide discussion is that practical factors, such as the cost of being wrong about p , are not epistemic.

to act, for some action (often this is done by positing that one agent is in a ‘high-stakes’ setting and the other in a ‘low-stakes’ setting). Assuming fallibilism, it’s natural to think the agent in the ‘low-stakes’ setting knows. Given this assumption, we can see why *Sufficiency* and purism cannot both be affirmed: if *Sufficiency* is true, then the agent that ought not act does not know the relevant proposition. Since the difference between the agent that knows and the agent that doesn’t know is due to some non-epistemic factor (the agent’s practical situation), purism is violated. (If we instead start by holding purism fixed, we get the result that *Sufficiency* is violated.)⁵

For better or worse (in my opinion, for better), fallibilism has earned the status of an indispensable commitment of contemporary epistemology. From this vantage point, the problem could be stated as a dilemma for fallibilists: fallibilists must choose between purism and *Sufficiency*.⁶ One goal of this chapter is to show that there is a fallibilist option that can avoid the trilemma.

My primary goal, though, is to argue against pragmatic encroachment by arguing against *Sufficiency*. The structure of this chapter is as follows: in Section 2, I challenge the account of reasons that underlies one prominent way of arguing for *Sufficiency* and then propose an alternative account of the relationship between knowledge and reasons for action. Section 3 examines the trilemma between purism, fallibilism, and *Sufficiency* and delineates a position that can maintain all three. Thus, there is a way out of the trilemma. I conclude with a final consideration in support of my preferred resolution of the trilemma: rejection of *Sufficiency*.

2 Against *Safe Reasons* and KJ

2.1 Unpacking the principles

The knowledge-action links have been formulated in a variety of nuanced ways. I will focus here on the following conception of *Sufficiency*, defended at length by Jeremy Fantl and Matthew McGrath (2009):

(KJ) If you know that p , then p is warranted enough to justify you in ϕ -ing, for any ϕ .⁷

⁵ See Fantl and McGrath (2009: 84–6) for a detailed explanation of the trilemma.

⁶ Fantl and McGrath (2009: last chapter) discuss infallibilism as a tenable option, though they do not go as far as to recommend that we reject fallibilism; on the contrary, they seem to strongly favor maintaining fallibilism.

⁷ Alternative wordings of KJ replace “ p is warranted enough to justify you in ϕ -ing, for any ϕ ” with other phrases, such as “one is in a good enough epistemic position to rely on p in practical

The central argument Fantl and McGrath offer to motivate KJ involves a framework that connects knowledge and rational action via reasons. Their case for KJ is based on the following two principles about reasons:

(KR) If you know that p , then p is warranted enough to be a reason you have to ϕ , for any ϕ . (2009: 69)

(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 . (2009: 77)⁸

In what follows I will argue that Safe Reasons is false. If my argument succeeds, one motivation for KJ is undermined. The same line of reasoning I offer against Safe Reasons can be advanced against KJ directly, showing KJ to be false.

A few preliminary points are in order. First, there are multiple ways a proposition could fail to justify you in ϕ -ing; epistemic propriety is just one dimension along which we can evaluate reasons for action. These principles speak only to the epistemic propriety of p for justified action. To fix on how epistemic propriety comes apart from other kinds of propriety, consider a setting where free doughnuts are being given out just outside your office, but you have no idea that this is the case. You fail to be in the right kind of epistemic position towards the proposition *there are free doughnuts outside your office* that is required for you to be justified in treating this proposition as your reason for stepping outside.

Alternatively, a proposition could fail to justify you in ϕ -ing because, despite having strong warrant for the proposition, the proposition fails to be connected in the right way to ϕ -ing. For example, even if I have top-grade warrant for the proposition *I exist*, this proposition does not justify me in walking my dog because it fails to be connected in the right way to the action. Nevertheless, the proposition *I exist* is *warranted enough* to justify me in walking my dog. In this case, something other than epistemic propriety prevents *I exist* from justifying me in walking my dog.

reasoning” (Brown 2012: 155) and “it is appropriate to treat p as a reason for acting” (Hawthorne and Stanley 2008: 578). The differences in wording will not be significant for my purposes: the difficulties I raise for KJ can be extended to these versions of the principle. In Section 3, I discuss a way of understanding Hawthorne and Stanley’s principle that does not equate it with KJ.

⁸ The argument depends on a further connecting principle:

(Connecting Principle): If p is warranted enough to be a reason you have for ϕ -ing, then p is warranted enough to justify you in ϕ -ing, for any ϕ .

I will focus my discussion on Safe Reasons rather than the Connecting Principle for two reasons: first, it is the principle that plays a central role in Fantl and McGrath’s discussion; and second, the locution of a ‘reason one has’ maps onto ordinary language more easily than the ‘warranted enough to be a reason one has.’ But note that the argument I level against Safe Reasons is also a counterexample to the Connecting Principle.

The discussion in this chapter concerns when weakness in one's epistemic position makes it inappropriate to treat a proposition as a reason for action. In Fantl and McGrath's preferred terminology, these are situations where weakness in epistemic position 'stands in the way' of a proposition justifying action.⁹ To identify when weakness in epistemic position stands in the way of a proposition justifying action, I introduce the following test:

Strengthening test: what S ought to do is the same as what S ought to do conditional on *p*.

When what one ought to do differs from what one ought to do conditional on some proposition, one fails the strengthening test for that proposition. When one fails the test, epistemic weakness stands in the way of appropriately treating that proposition as a reason to act.¹⁰ The strengthening test is naturally applied in a decision theory framework, where what one ought to do is the top ranked action on an ordering given by one's epistemic position and utilities.¹¹ What one ought to do conditional on some proposition is the action that ranks top conditional on that proposition.¹²

The strengthening test plays an important role in the principles above. Passing the strengthening test is a necessary condition on a proposition being a reason one has to ϕ (by Safe Reasons). In the same way, KJ makes passing the strengthening test a necessary condition on knowledge. I will argue against both of these principles. My discussion leaves KR untouched as a plausible connection between knowledge and reasons for action. The picture I am inclined towards is one whereby when one knows a proposition the proposition is a reason one *has* to act, but where being a reason one has does not entail that one is appropriate to treat

⁹ See Fantl and McGrath (2009: 67–8).

¹⁰ Passing the strengthening test is not a sufficient condition for one's epistemic position to be warranted enough to justify action because one might pass the test even though one's epistemic position is very weak. Plausibly there is some minimal threshold required, in addition to passing the strengthening test, in order for *p* to be warranted enough to justify action.

¹¹ The notion 'epistemic position' is flexible with respect to various interpretations; as I use it here, it is subject to a few constraints: first, it excludes non-epistemic features of one's situation; second, it is externalist—one's epistemic position does not supervene on one's internal mental state. Third, one's epistemic position is not to be construed as including one's total knowledge, as on such a picture one will automatically pass the strengthening test for each proposition known. See Section 3 for further discussion and for a relaxing of this last constraint.

¹² Some advocates of the knowledge-action principles may object to a test that relies on decision theory. For my purposes here, the helpfulness of the test relies on only basic and relatively harmless features of decision theory. The strengthening test bears strong similarity to the slogan test that Fantl and McGrath offer (2009: 68), which is as follows: "if merely strengthening your epistemic position can make a difference as to whether *p* justifies you in ϕ -ing, then weaknesses in your epistemic position stand in the way of its so justifying you."

the proposition as a reason to act—thus, being a reason one has to act does not require passing the strengthening test.¹³ Our epistemic position for reasons we have is not always strong enough to pass the strengthening test. (This should be unsurprising, given fallibilism.)¹⁴

Finally, we need to make it clear that, according to Safe Reasons, reasons you have are able to justify *any* available action. As Fantl and McGrath (2009: 77) state: if p is a “reason you have, then the chance that [p] is false can’t stand in the way of [p] justifying you in doing anything.”¹⁵ Safe Reasons claims that when p is a reason one has to ϕ , p is warranted enough to justify you in *any* available action.¹⁶ Hence, the strengthening test is an apt test, as it tests all available actions simultaneously. If one ought to x and y , but conditional on p one ought to x and $\sim y$, then one fails the strengthening test. One cannot pass the strengthening test for a particular available action while failing for other available actions: failing for one action is failing for all.

By requiring that one pass the strengthening test, Safe Reasons places a stringent requirement on being a reason one has. KJ likewise places a stringent test on knowledge. Whether p is a reason one has to ϕ and whether one knows p , on this picture, will depend on what actions are available in any given scenario. In situations where there are multiple available actions, for p to be warranted enough to justify S in ϕ -ing, it must be that for all actions available to S , what S ought to do does not differ from what S ought to do conditional on p . P is not a reason S has to ϕ if p is not warranted enough to justify S in ψ -ing. This may strike some as a particularly demanding requirement on reasons one has. In Section 2.2, we will see how this feature plays out in counterintuitive ways.

¹³ There is an expansive literature discussing ‘reasons one has.’ The discussion in this chapter is limited to the epistemic position required for p to be a reason one has—it is neutral with respect to other requirements, and in particular it is neutral with respect to the conditions under which p is an objective reason to ϕ . Throughout the chapter I will speak as though knowledge is sufficient to make p a reason one has, but of course knowing p is not sufficient for p to meet the objective requirements on being a reason one has, if there are such requirements. See Schroeder (2008) for a helpful discussion of objective and subjective reasons.

¹⁴ Although I find the knowledge-reasons connection represented by KR attractive, the argument offered in this chapter does not commit one to KR. Rather, it leaves open the option to reject KR. See Brown (2011) for discussion of this route.

¹⁵ An explicit formulation of Safe Reasons, as Fantl and McGrath intend it to be understood, is thereby as follows: for all S , p , ϕ , if p is a reason S has to ϕ , then for all ψ , p is warranted enough to justify S in ψ -ing.

¹⁶ Although Fantl and McGrath do not qualify ‘any action’ as any *available* action, as I do here, this seems to be their intended meaning. In (2009: 67) they use specifications such as ‘in a certain situation’ or ‘merely strengthening your epistemic position’ holding fixed other factors (such as the agent’s stakes). In holding fixed a subject’s stakes, I take it that we hold fixed what actions are available to the agent (that the subject is not facing a high-stakes bet, etc.). The addition of ‘available’ merely makes this evident. See also Fantl and McGrath (2009: 224–9) for relevant discussion.

2.2 *Against safe reasons*

Consider the following case:

Dinner

Alli tells her husband Tim that she is going out for the evening and won't be home until late. On the basis of her testimony, Tim's epistemic position for the proposition *Alli will be home late* (hereafter HOMELATE) is very strong. Tim decides to make pizza for dinner. Alli hates pizza, so he only has pizza for dinner when she is out. He considers inviting his brother over to play chess, but decides not to since Alli recently had a huge disagreement with his brother and she made it clear to Tim that she does not want to see him for a while.¹⁷

Tim's epistemic position for HOMELATE is very strong but not top level. Given his epistemic position, he is justified in making pizza but he shouldn't invite his brother over. Although Alli almost never comes home early when she goes out for the evening, if she did she would be very upset to run into Tim's brother. We can imagine Tim saying to his brother on the phone: "I'm making pizza for dinner tonight—Alli isn't going to be home until late. I'd invite you over to play chess, but if she came home early that'd be a disaster; I wouldn't want you guys to run into each other." Tim ought not invite his brother over, even though strengthening his epistemic position for HOMELATE makes it such that Tim ought to invite his brother over—that is, conditional on HOMELATE, Tim ought to invite his brother over. Tim fails the strengthening test.

It is natural to think that HOMELATE is a reason Tim has to do certain actions, and in particular, it is natural to think HOMELATE is a reason Tim has to make pizza. But HOMELATE is not warranted enough to justify Tim in *any* available action. Sometimes one's epistemic position is strong enough to justify one action while simultaneously not strong enough to justify another action.

Consider the following statements:

- (1) HOMELATE is a reason Tim has to make pizza.
- (2) HOMELATE is a reason Tim has to invite his brother over.
- (3) Tim ought to make pizza and ought not invite his brother over.
- (4) HOMELATE is warranted enough to justify Tim in making pizza.
- (5) HOMELATE is not warranted enough to justify Tim in inviting his brother over.
- (6) Conditional on HOMELATE, Tim ought to invite his brother over.

¹⁷ This case first appears in Anderson (2015).

Each of (1)–(6) are plausible things to say about this case. But (1), (5), and (6) are jointly incompatible with Safe Reasons. Two points deserve emphasis here. First, (1) is *prima facie* very plausible: Tim’s epistemic position for the proposition is very strong and making pizza is the rational thing for him to do. In this way, denial of (1) is a cost. But according to Safe Reasons, (1) is false. Second, the reason the proponent of Safe Reasons must deny (1) is surprising: the advocate of Safe Reasons must deny (1) *because* (5) and (6) are true. But intuitively, whether HOMELATE is a reason Tim has to make pizza should not depend on whether HOMELATE is warranted enough to justify Tim in inviting his brother over. Safe Reasons forces this counterintuitive result.

It might be tempting to think that the problem could be avoided by introducing probabilities. The envisioned strategy claims that the proposition available to function as Tim’s reason to make pizza is *Probably Alli is coming home late* (hereafter, HOMELATE*). The idea is that if HOMELATE* is a reason Tim has for action, rather than HOMELATE, then HOMELATE* can explain why Tim is justified in doing one action but not the other.

The first thing to note about this strategy is that if Alli comes home late and asks Tim why he made pizza for dinner, he won’t give as his reason that she probably was coming home late. His reason for making pizza is that she was coming home late. Propositions about probabilities do not always provide plausible alternatives when it comes to the reasons that actually motivate us.

Furthermore, in at least some cases, the shift to HOMELATE* will merely push the problem back a step. Consider that according to Safe Reasons, if HOMELATE* is a reason Tim has to ϕ , then HOMELATE* is warranted enough to justify Tim in ϕ -ing, for any ϕ . But assume that Tim has a very strong but less than top-level epistemic position for HOMELATE*. We can construct the details of the case in such a way that Tim is rational to make pizza, but is not rational to invite his brother over, even though conditional on HOMELATE* he is rational to invite his brother over. When HOMELATE* fails the strengthening test, by Safe Reasons, HOMELATE* is not a reason Tim has to do any action. There is nothing special about propositions about probabilities that make them resistant to counterexamples of the kind in *Dinner*.¹⁸

¹⁸ The strategy can be repeated at higher levels. The advocate of Safe Reasons might suggest that Tim’s reason is only *probably* HOMELATE*. But there is no reason in principle why the counterexample cannot be constructed around this alternative proposition. For each iteration there will be cases where, given fallible knowledge of the relevant proposition and the right details of the case, the replacement proposition will justify one action even though weakness in epistemic position for the proposition stands in the way of another action. Of course, the more iterations that are added, the less plausible it will be that we actually treat the suggested proposition as our reason for acting.

More importantly, the availability of HOMELATE* or other propositions that Tim *could have* treated as his reason to make pizza for dinner does not change the fact that it is plausible to think that HOMELATE is a reason Tim has to make pizza for dinner. (This applies equally to alternatives such as *Alli said she was coming home late.*) If *as a matter of fact* Tim treats HOMELATE as his reason to make pizza, the framework under consideration delivers the result that he does something inappropriate. This strikes me as the wrong result.¹⁹

The reasoning in *Dinner* can be advanced directly against KJ by positing that Tim knows HOMELATE. Given certain assumptions about Alli's reliability as a testifier, it is highly plausible that Tim comes to know HOMELATE on the basis of her testimony. If Tim knows HOMELATE, and HOMELATE is not warranted enough to justify Tim in making pizza for dinner and inviting his brother over, then KJ is false. We can construct a similar line of reasoning for HOMELATE*. It is plausible that Tim knows HOMELATE, and even more plausible that he knows HOMELATE*. In this way, *Dinner* is a counterexample to both Safe Reasons and to KJ.²⁰

We should reject Safe Reasons and KJ. Passing the strengthening test for all available actions is too strong a requirement for a proposition to be a reason one has and too strong a requirement for knowledge.

In light of these difficulties, one might be inclined to formulate Safe Reasons and KJ such that neither require passing the strengthening test for all available actions. That is, when ϕ and ψ are available actions, one might hold that one knows p (or, that p is a reason one has to ϕ) if one passes the strengthening test with respect to ϕ -ing even if one does not pass the strengthening test with respect to ψ -ing. Weakening KJ in this way is not a very promising option, since knowing p with respect to one action when one does not know p with respect to a different action will involve a kind of compartmentalization of knowledge that advocates of pragmatic encroachment are not likely to embrace. Weakening Safe Reasons is

¹⁹ In more recent work, Fantl and McGrath (manuscript) state that the action need only be available to the subject in order for failing the strengthening test for that action to deprive the subject of p as a reason S has (and also to deprive the subject of knowledge). The result is that Tim fails to have HOMELATE as a reason to make pizza for dinner when inviting his brother over is an available action even in a setting where Tim does not consider inviting his brother over. In my opinion, this adds further implausibility to the overall picture.

²⁰ It's worth considering how we might respond if some action that prevents us from passing the strengthening test were always available—for example, if there were a genie offering us a high-stakes bet for every proposition we were prepared to treat as a reason. I'd wager that we would hold fast to the idea that these propositions are reasons we have to act, for all sorts of actions, and conclude that our inability to treat them as reasons to accept the genie's bets is irrelevant to whether we can permissibly treat the propositions as reasons to act in other ways.

likewise not an appealing option for the KJ theorist because Safe Reasons will no longer be able to motivate KJ.

At this point it should be clear that KJ relies on a thought that is unnatural given fallibilism. Given fallibilism, there is no reason to expect that when we know p , we are justified in acting on p for any available action. In fact, fallibilism predicts exactly the opposite. That is, fallibilism predicts that sometimes when we know p , the difference between our actual epistemic position for p and a stronger position with respect to p will make a difference to whether p is warranted enough to justify some available action. In this way, KJ (and the framework of reasons underlying KJ) involves a picture that does not fit well with fallibilism.

3 The Trilemma

I turn now to examine the trilemma between fallibilism, purism, and *Sufficiency*. A central term in *Sufficiency* involves an ambiguity which, when elucidated, provides a way out of the trilemma. Yet even in light of a position that can maintain all three commitments, I contend that we have reason to reject *Sufficiency*.

Here are the principles involved in the trilemma, as they are commonly glossed:

Purism: pragmatic factors do not affect whether one knows.

Fallibilism: knowledge is compatible with an epistemic chance of error.

Sufficiency: when S knows p , S's epistemic position for p is such that p is warranted enough to justify S in ϕ -ing, for any ϕ .

The notion of strength of epistemic position is notoriously vague. Without offering a complete account, I will draw a distinction between two disparate ways of thinking about epistemic position. The first construes one's epistemic position in such a way that it aligns naturally with the notion of epistemic chance used to articulate fallibilism. The second departs from this alignment. In what follows, I will paint in broad brushstrokes. I do not intend to give an account of fallibilism, but rather to offer a general picture that many if not all fallibilists agree upon, despite significant disagreement regarding exactly how to articulate fallibilism.

An important feature of the notion of epistemic chance as it is used in statements of fallibilism is that it concerns something other than the facts of success in a particular case—it represents something more like an approximation of one's modal success rate across similar worlds.²¹ Even when we get it right in

²¹ The modal success rate here involves the assumption that the relevant sets of worlds are finite. Things are more complicated when it comes to infinite worlds. For the most part, these complexities

the actual world, there are similar worlds where we form a false belief (in either the same proposition or a similar proposition).^{22, 23} Fallibilists agree that although our modal success rate is less than perfect, nevertheless, we can possess knowledge.²⁴ For example, suppose you come to know that Paul is coming to the party because your friend tells you. Many cases of testimony are such that there is some similar world where the testifier is misinformed or misrepresenting information but where you form a belief on the basis of the testimony nevertheless.²⁵ Such worlds negatively affect an agent's modal profile with respect to particular cases of believing on testimony.²⁶ When considering an unrestricted set of similar worlds, fallibilists can agree about one thing: an imperfect modal success rate is compatible with knowledge.

A further choice point concerns whether knowledge is compatible with an imperfect epistemic position relative to some restricted subset of the similar worlds. It's fairly standard to think that not all similar worlds are relevant to whether you know, and thus to restrict the worlds relevant to knowledge to some more narrow subset. There are similar worlds, after all, where extremely unlikely events occur—the bank is closed because terrorists attack, or an earthquake destroys the building. Nevertheless, it's common to think that in the actual world (where these events do not occur), we can know the bank is open. With respect to this restriction, there is further division between fallibilists concerning whether knowledge requires a perfect modal success rate within this restricted set. (This choice point lines up nicely with a distinction sometimes drawn between strong and weak safety.) Some fallibilists allow that one can know even if one believes *p* falsely within the aforementioned set, others require perfection within the restricted set. It's important to note that this is an 'in-house' disagreement among fallibilists of various stripes.

are not relevant to the picture I present, though it's worth noting that notions like 'epistemic chance' are problematic in the infinite setting and are not interchangeable with 'epistemic possibility.'

²² Where by 'similar world' we hold fixed that you believe using the same or a similar method.

²³ A standard way to account for fallible knowledge of necessary truths is to include in the set of relevant worlds those where one forms beliefs about similar propositions.

²⁴ How strong your success rate must be to be compatible with knowledge is a contentious issue. Since I do not intend here to give an account of fallibilism, I have no need to articulate a threshold.

²⁵ A fully general account of fallibilism will inevitably want to provide a principled reason for drawing the line of similar/relevant worlds. Various options are available here—some may be inclined to use something like the nearby worlds compatible with one's mental states. I leave open what constraints are needed.

²⁶ It's worth noting that because the modal success rate is an externalist notion, one's epistemic position is less than perfect even if one is not aware of one's own fallibility—one is not always in a position to know one's modal profile.

While fallibilists disagree on this point, what they have in common is that one's *unrestricted* modal success rate is imperfect. I'll call this notion of epistemic position (derived from consideration of imperfection in the unrestricted set) one's E_I epistemic position.²⁷ It is this notion that lies behind the gloss on fallibilism above.

An alternative conception of epistemic position can be construed by looking only at the restricted set of similar worlds. The view that one's modal success rate within the restricted set must be perfect in order for one to know results in a corresponding picture of epistemic position whereby when one knows, there is no chance of believing falsely. On this picture, all the worlds where you mess up are outside the relevant range and are thereby irrelevant to this construal of epistemic position. I will refer to this gloss on epistemic position (which requires perfection within the relevant set of similar worlds) as one's E_P epistemic position.

There is an important choice point regarding whether the former or latter notion of epistemic position is relevant to justified action, and thus to *Sufficiency*. Some but not all versions of *Sufficiency* involve thinking of one's epistemic position using the notion of one's E_I chance of error. Given the various construals of epistemic position available, fallibilists face a choice point with respect to which gloss on one's strength of epistemic position is relevant to what one ought to do: one's E_I or one's E_P . There is room for disagreement among fallibilists concerning the relevancy of E_I to justified action. Fallibilism is a commitment about the compatibility of knowledge with an epistemic chance of error—a chance that I have suggested can be understood with respect to one's E_I epistemic position. It is a further (though somewhat natural) commitment for the fallibilist to affirm that the same notion of epistemic position according to which one has a chance of error is the notion of epistemic position relevant to determining what one ought to do.²⁸ It is open to the fallibilist to hold that one's E_I epistemic position is irrelevant to rational action and instead to think of one's E_P epistemic position as the relevant notion.

We are now in a position to see the role that the notion of one's E_I plays in the trilemma. Here is the trilemma again, where the notion of one's E_I is made explicit:

²⁷ One could also speak of the epistemic probability of p relative to one's E_I position, if one holds that epistemic position is the kind of thing that satisfies the probability axioms. Since there is some doubt as to whether and how epistemic position can be construed in terms of probability, I leave it to the reader to make the application.

²⁸ One might try to build into the notion of 'fallibilism' the idea that what one ought to do depends on one's E_I epistemic position. But although many fallibilists do, in fact, think one's E_I position is relevant to justified action, this constraint is not built into the definition of fallibilism. There is space for a fallibilist position that rejects this picture.

Purism: pragmatic factors do not make a difference to whether one knows p .
*Fallibilism**: for some p , S knows p and there is an E_I chance for S that $\sim p$.
*Sufficiency**: when S knows p , S's E_I epistemic position for p is such that p is warranted enough to justify S in ϕ -ing, for any ϕ .

This version of *Sufficiency* gives rise to the trilemma. But fallibilists who hold that one's E_P epistemic position determines what one ought to do are likely to construe sufficiency instead as follows:

*Sufficiency***: when S knows p , S's E_P epistemic position is such that p is warranted enough to justify S in ϕ -ing, for any ϕ .

In virtue of maintaining that one's E_P determines what one ought to do, on this view one will always pass the strengthening test when one knows p (since what one ought to do will always be the same as what one ought to do conditional on p). In this way, *Sufficiency*** is compatible with fallibilist* purism. Fallibilists who hold that one's E_P determines what one ought to do, and thereby hold *Sufficiency*** instead of *Sufficiency**, avoid the trilemma.^{29, 30}

We'll call views that affirm fallibilism* and *Sufficiency*** *chance-irrelevant fallibilism*. According to chance-irrelevant fallibilism, when S knows p , and we consider what S is rational to do, there is no need to consider S's E_I position for p . Since what one ought to do is not given by one's E_I position, it is irrelevant on this view whether one passes or fails the strengthening test with respect to one's E_I epistemic position.

Alternatively, there is *chance-relevant fallibilism*. Chance-relevant fallibilists hold that one's E_I epistemic position is relevant to rational action. Thus, on this view, sometimes the difference between an E_I position of 0.97 for p and an E_I position of 0.99 will make a difference to what one ought to do. (I assume here an idealization of epistemic position such that it can be represented using numbers between 0 and 1. Those who find this objectionable may substitute 'strong epistemic position' and 'stronger epistemic position' for the relevant values.)

²⁹ On one reading of Hawthorne and Stanley (2008) their position is an example of a view that uses a subject's E_P to decide what one ought to do. Their position denies purism, despite the fact that *Sufficiency*** does not force the denial of purism. So denial of purism is a superfluous feature of the view. See also Dutant (manuscript) for relevant discussion.

³⁰ There are other options I haven't considered here. One could try to construe epistemic position using a scale such that one's epistemic position is represented as perfect only when one knows that one knows. When wedded to decision theory, this view will favor a version of *Sufficiency* along the following lines: when S knows that she knows p , S is warranted enough to act as if p . One who holds this principle will also avoid the trilemma. Further iterations of knowledge to set the top threshold of the scale of the epistemic notion relevant to determining what one ought to do are also an option. Exploration of these various positions would take us too far afield.

There is a further divide between chance-relevant fallibilists. According to chance-relevant fallibilists who affirm *Sufficiency**, when the difference between an E_1 position of 0.97 and an E_1 position of 0.99 makes a difference to what S ought to do, S does not know p . Hence, not only is the difference between a strong E_1 position and a top-level epistemic position relevant to justified action, it is also relevant to knowledge. This results in pragmatic encroachment. An alternative chance-relevant position is *stable fallibilism*. Stable fallibilists hold that one's E_1 position is relevant to justified action, and one can know p even when weakness in one's E_1 stands in the way of acting (that is, even when one's E_1 position is such that one fails the strengthening test). Knowledge is stable on this view—whether one has knowledge does not depend on one's practical decision setting.

In closing, I advance a consideration in favor of stable fallibilism that has gone unappreciated. The consideration is this: only stable fallibilism allows us to put fallibilism to work. A key fallibilist insight is that knowing p is compatible with a small E_1 chance that $\sim p$. Stable fallibilism is the only view among the three that allows this insight to be widely relied upon. Chance-irrelevant fallibilism, although it is committed to the compatibility of knowledge and a small E_1 chance that $\sim p$, does not allow these chances to make a difference to rational action because E_1 chances are always irrelevant to what one ought to do. Chance-relevant fallibilism tied to *Sufficiency** also does not allow the fallibilist insight to be put to work because when the E_1 chance that $\sim p$ is such that one fails the strengthening test, one fails to know p . Consequently, one's fallible knowledge cannot be relied on for anything in these settings. Both of these views hold that you know p only when it is rational for you to act just as you ought to act conditional on p (that is, conditional on there being no E_1 chance of error). In this way, both views make the key fallibilist insight—the compatibility of knowledge and a chance of error—idle.

In conclusion, I've argued that KJ and Safe Reasons ought to be rejected and that affirmation of *Sufficiency* involves fallibilists in a kind of double-mindedness. I've delineated two notions of epistemic position relevant to unpacking the principles in the trilemma, and shown how *Sufficiency* can be understood in a way consistent with purism and fallibilism. Finally, I've argued that stable fallibilism has an unappreciated advantage over the competing two views: it is the only position among the three that allows us to put a key fallibilist insight to work.³¹

³¹ I am particularly grateful to Julien Dutant and John Hawthorne for providing detailed comments on drafts of this chapter. Thanks also to Trent Dougherty, John Greco, Jeffrey Russell, Matthew McGrath, and the editors of this volume for helpful comments and discussion.

References

- Anderson, C. (2015). On the intimate relationship of knowledge and action. *Episteme*, 12(3): 343–53.
- Brown, J. (2011). Fallibilism and the knowledge norm. In *Assertion: New Philosophical Essays*. Edited by J. Brown and H. Cappelen. Oxford: Oxford University Press.
- Brown, J. (2012). Assertion and practical reasoning: Common or divergent epistemic standards? *Philosophy and Phenomenological Research*, 84(1): 123–57.
- Dutant, J. Normative sceptical paradoxes. Unpublished manuscript.
- Fantl, J. and McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1): 67–94.
- Fantl, J. and McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75(3): 558–89.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Fantl, J. and McGrath, M. (manuscript). On two ultimately unsuccessful objections to pragmatic encroachment. Presented at Oxford University's Workshop on Contextualism, Pragmatic Encroachment, and Religious Epistemology.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, J. and Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, CV (10): 571–90.
- Schroeder, M. (2008). Having reasons. *Philosophical Studies*, 139: 57–71.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.