

Unifying Epistemic and Practical Rationality

Mattias Skipper

Penultimate draft, forthcoming in *Mind*

Abstract: Many theories of rational action are predicated on the idea that what it is rational to *do* in a given situation depends, in part, on what it is rational to *believe* in that situation. In short: they treat epistemic rationality as *explanatorily prior* to practical rationality. If they are right in doing so, it follows, on pain of explanatory circularity, that epistemic rationality cannot itself be a *form* of practical rationality. Yet, many epistemologists have defended just such a view of epistemic rationality. According to them, there is no such thing as a distinctively epistemic form of rationality which could be explanatorily prior to practical rationality. Rather, they maintain, there is just *one* form of rationality—practical rationality—of which epistemic rationality is a species. What to make of this conflict? The aim of this paper is to motivate a view about the relationship between epistemic and practical rationality which resolves the conflict in a way that should be attractive to both sides. The central idea is to ground both epistemic and practical rationality in an independently motivated notion of *evidential probability* which is itself to be understood in non-normative terms. Doing so, I argue, allows us to unify epistemic and practical rationality in a way that does justice to the idea that what it is rational to do depends, in part, on what it is rational to believe; and to do so in a way that avoids explanatory circularity.

1. Introduction

Much of contemporary decision theory is predicated on the idea that what it is rational to *do* in a given situation depends, in part, on what it is rational to *believe* in that situation. Whether it is rational for me to cook you dinner depends, in part, on whether it is rational for me to believe that you would appreciate my doing so; whether it is rational for me to accept the job offer depends, in part, on whether it is rational for me to believe that my new colleagues will be nice acquaintances; and so on. Examples like these are easy to come by. And they are naturally taken to show that the correct theory of rational action, whatever it is, must appeal to some notion of rational belief. In short: they are naturally taken to show that epistemic rationality is *explanatorily prior* to practical rationality.

If epistemic rationality is, in fact, explanatorily prior to practical rationality, it follows, on pain of explanatory circularity, that epistemic rationality cannot itself be a

form of practical rationality (more on this circularity in §2.1). Yet, many epistemologists have defended just such a view of epistemic rationality. According to them, there is no such thing as a distinctively epistemic form of rationality which could be explanatorily prior to practical rationality. Rather, they maintain, there is just *one* form of rationality—practical rationality—of which epistemic rationality is a species.

What to make of this conflict?

My aim in this paper is to motivate a view about the relationship between epistemic and practical rationality which resolves the conflict in a way that should be attractive to both sides. Certain concessions will have to be made. Indeed, I will argue that we should abandon the view that epistemic rationality is a form of practical rationality. But I hope to show that there is a way of doing so which allows us to preserve the core motivations that led philosophers to propose such a view of epistemic rationality in the first place. The central idea is to ground both epistemic and practical rationality in an independently motivated notion of *evidential probability* which is itself to be understood in non-normative terms. Doing so, I will argue, allows us to unify epistemic and practical rationality in a way that does justice to the idea that what it is rational to do depends, in part, on what it is rational to believe; and to do so in a way that avoids explanatory circularity.

As we shall see, the central proposal of the paper is not without substantive commitments. For example, it will not appeal to those who deny on general grounds that practical considerations can play any role in determining what it is rational to believe. I will have more to say about this matter in later sections. For now, I just want to emphasize that my ambition here is not to proceed only from premises that everyone will readily accept. My aim is a more modest one: to outline what seems to me to be the most promising way of unifying epistemic and practical rationality. Whether the proposal is ultimately acceptable will depend on one's other theoretical commitments.

Here, then, is the plan: I will begin by describing the central conflict in a bit more detail (§2). Then, I will develop my proposal for how to resolve the conflict (§3). Finally, I will address some potential worries one might have about the proposal (§4), before rounding off (§5).

2. The conflict

2.1 Thesis: epistemic rationality is not a form of practical rationality

Let me begin by rehearsing a familiar picture of rational action. On this picture, what it is rational for an agent to do in a given situation depends on two components: a *belief-like* component and a *desire-like* component. Roughly speaking, the belief-like component is supposed to capture what the agent thinks the world *is* like, whereas the desire-like component is supposed to capture what the agent wishes the world *were* like.

Together, these two components—the agent’s belief-like and desire-like attitudes—are then supposed to combine in some way to determine what it is rational for the agent to do. For example, if I am hungry and believe that there is food in the kitchen, it might be rational for me to get off the couch and go to the kitchen to get food.

There are different ways one might develop this familiar picture of rational action, and it will not matter for present purposes just how this is done. However, it will be useful to have a concrete theory of practical rationality to work with, so I will frame the following discussion against the background of orthodox Expected Utility Theory. According to this theory, acting rationally is a matter of maximizing expected utility. A bit more precisely:

Practical EUM: It is rational for an agent to perform a given action if, and only if, no alternative action has a higher expected utility, calculated with respect to the agent’s utility function and rational credence function.¹

The desire-like component is here given by the agent’s utility function, which is traditionally defined as a function from possible worlds to real numbers, representing the degree to which the agent desires each world. The belief-like component is given by the agent’s rational credence function, which is traditionally defined as a function from propositions to real numbers between 0 and 1, representing the agent’s rational degree of confidence in each proposition. Together, these two functions then determine a unique expected utility of each available action, which, in turn, determines what it is rational for the agent to do.

All of this will be familiar territory for most readers.² But as mentioned, the finer details are will not matter for what follows. The important thing to notice here is that Practical EUM takes epistemic rationality to be explanatorily prior to practical rationality: it says that what it is rational to do depends, in part, on what it is rational to believe.

Of course, to say that epistemic rationality is *explanatorily prior* to practical rationality is not yet to say that epistemic rationality is not a *form* of practical rationality. To arrive at this conclusion, we need to add a further premise. Before doing so, however, I want to clarify a few points about what has been said so far.

First, note that Practical EUM is formulated in terms of an agent’s *rational* credence function rather than the agent’s *actual* credence function. This might give some readers pause. After all, if you act in accordance with what you *think* will produce the best outcome (more precisely: if you perform the action that has the highest expected utility, calculated with respect to your *actual* credences), isn’t there a sense in

¹ Here and throughout I will be assuming that the right-hand-side of the biconditional explains or determines the left-hand-side.

² If not, Thoma (2019) offers a helpful (and freely available) introduction to orthodox decision theory.

which your action is rational, even if your credences are not? I think so. But there is likewise, I think, a recognizable sense in which your action *fails* to be rational in such a case. The following example is adapted from Christensen (2020, p. 4):

Charlie the Soap-Muncher: Charlie, an eternal optimist, has eaten soap many times in the past. Each time, it has tasted terrible. Nonetheless, Charlie continues to believe with undiminished strength that, from now on, soap will taste great. Lucy, as she has many times in the past, offers Charlie a bar of soap to eat. Charlie, confident as he is that the soap will taste great, happily eats it. Of course, it tastes terrible.

I think we can readily agree that there is a sense in which Charlie's action is perfectly rational. After all, he does what he thinks will have the best consequences. But there also seems to be a recognizable sense in which Charlie's action fails to be rational, precisely because it fails to be sensitive to the fact that, on repeated occasions, he has experienced soap tasting terrible. It is this latter sense of rationality which interests me here, and which Practical EUM is supposed to capture.

Second, given that Practical EUM is stated in terms of rational credences rather than actual credences, one might wonder why it is not likewise stated in terms of the agent's rational utility function rather than the agent's actual utility function. The reason for this is that, although many philosophers agree that desire-like attitudes can be evaluated for rationality, there is no need to take a stance on this issue here. Nothing in what follows is going to turn on whether Practical EUM is stated in terms of actual utilities or rational utilities.

Finally, the claim that epistemic rationality is explanatorily prior to practical rationality should not be taken to imply that we cannot in our philosophical theorizing rely on claims about what it is rational to do in order to argue for or against claims about what it is rational to believe. Take, for example, the famous "Dutch Book" argument for Probabilism, which purports to show that an agent's credences should obey the probability axioms, since the agent would otherwise be willing to accept a series of bets that is guaranteed to lose the agent money. This is an argument which seeks to establish a substantive conclusion about what it is rational to believe on the basis of a comparatively uncontentious premise about what it is rational to do. And while there are certainly worries one might have about this style of argument, the argument is, I take it, compatible with the claim that epistemic rationality is explanatorily prior to practical rationality. Directions of argumentation need not always follow directions of explanation. This is true in philosophy, as much as elsewhere.

Back to the main thread. As noted above, to say that epistemic rationality is explanatorily prior to practical rationality is not yet to say that epistemic rationality is not a form of practical rationality. To establish this conclusion, we also need to assume that a theory of rationality should not be *circular*: that is, a theory of epistemic rationality

should not appeal to the notion of epistemic rationality in the explanans, and a theory of practical rationality should likewise not appeal to the notion of practical rationality in the explanans. Such a non-circularity constraint is, of course, often assumed without argument, and although it connects up with some difficult issues in the debate between foundationalist and coherentist theories of justification, I will do the same here.

It will, however, be helpful to consider what the circularity would look like if were to combine the claim that epistemic rationality is explanatorily prior to practical rationality with the claim that epistemic rationality is a form of practical rationality. To this end, let us again take Practical EUM as our background theory of practical rationality, and let us suppose we wanted to treat epistemic rationality as a form of practical rationality. The resulting theory of epistemic rationality would then be obtained by making the relevant substitutions in Practical EUM:

Circular Epistemic EUM: It is rational for an agent to have a given credence function if, and only if, no alternative credence function has a higher expected utility, calculated with respect to the agent's utility function and rational credence function.

Clearly, this norm is circular: the agent's rational credence function figures on both sides of the biconditional. So, if circularity is to be avoided, Practical EUM commits us to the claim that epistemic rationality is not a form of practical rationality.

Let me pause here to clarify a few points. First, it bears mentioning that a circular theory of epistemic rationality need not be "trivial" or "vacuous" in the sense that it places no constraints on what it is rational to believe. This has been emphasized by Rinard (2017, pp. 133-34), who expresses sympathy for a similar norm to Circular Epistemic EUM. As she points out, such a norm forbids combinations of beliefs of the form " p , but not believing p would have a higher expected utility." This already suffices to show that Circular Epistemic EUM is not an "anything goes" view of rational belief. It may, of course, still seem much too weak for a general norm of belief.³ But in any case, it remains circular, which I shall assume renders it untenable.

Second, one might wonder whether it is possible to avoid the circularity by formulating Circular Epistemic EUM in terms of credences in particular propositions rather than entire credence functions. Here is one way of doing so:

Regressive Epistemic EUM: It is rational for an agent to have a credence of x in a proposition p if, and only if, no alternative credence in p has a higher expected utility, calculated with respect to the agent's utility function and the credences it is

³ See Christensen (2020, pp. 10-13) for a discussion of this point.

rational for the agent to have in each of the various possible outcomes of the agent's having such-and-such credence in p .

There is no obvious circularity involved in this norm, at least insofar as we can determine the rational credences in the various possible outcomes of the agent's having such-and-such credence in p without first having to determine the rational credence in p . However, it leads to an infinite regress: in order to determine the rational credences in the various possible outcomes of the agent's having such-and-such credence in p , we first need to determine the rational credences in the various possible outcomes of the agent's having such-and-such credences in the various possible outcomes of the agent's having such-and-such credence in p , and so on *ad infinitum*. One might perhaps try to argue that we should embrace this infinite regress—but it does not strike me as a promising way forward.

Finally, it is worth emphasizing that the circularity exhibited by Circular Epistemic EUM is not an artefact of our having chosen Practical EUM as our background theory of practical rationality. Any theory of practical rationality which treats epistemic rationality as explanatorily prior to practical rationality will lead to the same kind of circularity, when combined with the claim that epistemic rationality is a form of practical rationality (see Figure 1 for an illustration). And since this includes every theory of practical rationality which respects the idea that what it is rational to do depends, in part, on what it is rational to believe, we seem to have a compelling case for the thesis that epistemic rationality is not a form of practical rationality.




Figure 1: *Circular relationship between epistemic and practical rationality.*

2.2 Antithesis: epistemic rationality is a form of practical rationality

I will now turn to the other side of the conflict. There is a venerable tradition in philosophy, going back at least to Quine, of trying to reduce epistemic rationality to practical rationality. In recent decades, versions of this reductive project have been developed by, among others, Foley (1987), Laudan (1990), Kornblith (1993), Nozick (1993, ch. 3), Papineau (1999), Steglich-Petersen (2006; 2018), and Rinard (2017; 2019a). The particular views defended by these philosophers differ in a number of ways, but they all share a basic commitment to the idea that epistemic rationality is a form of practical rationality.

It is important to be clear about what it means to say that epistemic rationality is a form of practical rationality. The claim is not just that practical stakes can make a difference to what it is rational to believe in a given situation. Such “pragmatic encroachment” views in epistemology may be true even if epistemic rationality is not itself a form of practical rationality.⁴ The claim, rather, is that there is no such thing as a distinctively epistemic form of rationality in the first place. Instead, the thought goes, there is just *one* form of rationality—practical rationality—of which epistemic rationality is a species. Let us refer to this view of the relationship between epistemic and practical rationality as *Reductive Pragmatism*.⁵

What has led some philosophers to endorse Reductive Pragmatism? Historically, one of the central motivations has come from an ambition to accommodate epistemic rationality within a naturalistically kosher picture of the world; a task which has often been thought to be more feasible if epistemic rationality can be treated as a species of practical rationality.⁶ The force of such naturalistic concerns is of course up for debate, and I will not place much weight on them here. Instead, I will focus on two different lines of motivation, which I suspect will have at least somewhat broader appeal, and which have in fact played a central role in recent defenses of Reductive Pragmatism.

The first line of motivation has its roots in a different kind of ambition, namely that of providing a *unified* theory of epistemic and practical rationality. Many philosophers have found it difficult to see how we could give a unified treatment of epistemic and practical rationality if they differed fundamentally in kind. Ultimately, I hope to show that there *is* a plausible way to unify epistemic and practical rationality without reducing one to the other (this is the goal of §3). But for now, the important point is that Reductive Pragmatism promises to secure the relevant unification in a straightforward manner, by reducing epistemic rationality to practical rationality. This has often been seen by its proponents as an important advantage of the view.⁷

The second line of motivation concerns the question of why we should care about being epistemically rational in the first place. As has often been pointed out, it is far from clear why we should care about certain purely epistemic properties of our beliefs

⁴ For example, one kind of pragmatic encroachment view is what Worsnip (2021) calls *threshold pragmatism*: the view that, although practical considerations cannot themselves count for or against believing a proposition, they can influence the degree of evidential support needed for rational belief. This kind of view is, I take it, compatible with the thesis that epistemic rationality is not a form of practical rationality.

⁵ Here I am using the term “reduction” in a somewhat broad sense. On a more restricted usage, one might say that those who deny that there is an interesting distinction to be drawn between epistemic and practical rationality do not purport to “reduce” anything to anything else. However, what ultimately matters for present purposes is what is *explanatorily prior* to what. Whether those explanatory priority relations involve reductions or not is less important.

⁶ See, e.g., Quine (1969), Laudan (1990), and Kornblith (1993).

⁷ See, e.g., Foley (2001), Steglich-Petersen (2018), and Rinard (2017; 2019a).

(say, their being true, or their being supported by the evidence, or their being reliably formed, or what have you) for their own sake.⁸ There are, of course, those who maintain that it is intrinsically valuable to have true beliefs.⁹ But I take it that most of us would find it more plausible to say that true beliefs are valuable because of the role they play in helping us to achieve our practical goals in life (which might sometimes just be a matter of satisfying our curiosity about whether a given proposition is true). This is the kind of story which is available to those who view epistemic rationality as a species of practical rationality.

Needless to say, there are doubts one might have about each of these lines of motivation. For example, one might doubt the importance of providing a unified account of epistemic and practical rationality, or one might doubt whether a view that grounds epistemic rationality in an agent's desire-like attitudes can offer a satisfying explanation of why we should care about being epistemically rational. My aim here is not to dispel such doubts. For the purposes of this discussion, I will simply proceed on the assumption that the considerations mentioned above provide at least some motivation to think that epistemic rationality is a form of practical rationality.

There is, however, one line of resistance, which it will be helpful to address in some detail, because it offers an opportunity to clarify the notion of epistemic rationality I want to focus on in this paper. Recently, some epistemologists have argued that we in fact *shouldn't* care about being epistemically rational, unless we have good practical reason to do so.¹⁰ According to them, the requirements of epistemic rationality have a similar normative status to the rules of a game like chess: just as we should care about complying with the rules of chess only insofar as we have good practical reason to do so (say, because we want to become chess champion of the world), so we should care about complying with the requirements of epistemic rationality only insofar as we have good practical reason to do so (say, because it will help us achieve our goals in life). As it is sometimes put, the requirements of epistemic rationality are "thin:" they do not by themselves constrain what we *just plain should* believe, just as the rules of chess do not by themselves constrain what we *just plain should* do.

I think we can readily grant that there is a meaningful notion of epistemic rationality which is "thin" in the sense that it does not capture what we *just plain should* believe. This is no different from granting that there is a thin sense in which you "should" comply with the rules of chess. The important point, as I see it, is that we

⁸ See, e.g., Steglich-Petersen (2011), Friedman (2018), and Nolfi (2019).

⁹ Oft-cited examples include Goldman (1999), Lynch (2004), and Alston (2005). However, all three authors seem to concede at various points that the value of true belief ultimately depends on our goals and interests; see Goldman (1999, §3.4), Lynch (2004, pp. 15-19), and Alston (2005, §2.1). So, it is not entirely clear to me that they ultimately want to say that true belief is intrinsically valuable.

¹⁰ See, e.g., Côté-Bouchard and Littlejohn (2018) and Maguire and Woods (2020).

can intelligibly ask what we *just plain should* believe, just as we can intelligibly ask what we *just plain should* do. In doing so, we are asking about a “thick” normative notion, one that has a kind of normative authority over our beliefs that the rules of chess do not have over our actions. This is the normative notion I want to elucidate in this paper. What we call it is not of primary importance. Rinard (2019a) calls it the *guidance-giving* sense of “should.” Reisner (2018) calls it the *non-domain specific all things considered ought*. Maguire and Woods (2020) talk about what we *just plain ought* to believe. Here, I will continue to use the term “epistemic rationality” to capture what we should, in the relevant thick sense, believe.

We have now seen why some writers have been attracted to the view that epistemic rationality is a form of practical rationality. But there has, of course, also been voices urging that epistemic rationality *cannot* be a form of practical rationality (and not for the reasons discussed in §2.1). For example, some have argued that one can have compelling reason to believe a proposition even if one has no practical reason to do so, in which case Reductive Pragmatism is extensionally inadequate (Kelly 2003); others have criticized Reductive Pragmatism on the grounds that one cannot form beliefs in response to practical considerations, but only epistemic ones (Kelly 2002; Shah 2006); and yet others have argued that epistemic and practical reasons do not interact in the way we should expect if they were of the same kind (Berker 2018).

These are all formidable challenges, worthy of serious consideration. Nonetheless, I shall not address them here, partly because they have been addressed elsewhere in the literature,¹¹ and partly because my goal here is not to argue that epistemic rationality is a form of practical rationality. On the contrary, the view I will develop below *doesn't* treat epistemic rationality as a species of practical rationality. This is what will allow us to avoid explanatory circularity while doing justice to the idea that what it is rational to do depends, in part, on what it is rational to believe. At the same time, the view is going to provide us with a unified picture of epistemic and practical rationality, as well as a simple pragmatist explanation of why we should care about being epistemically rational in the first place. This is what will, I hope, make the view attractive to at least some proponents of Reductive Pragmatism.

Enough preview. Now for the details.

3. Resolving the conflict with evidential probabilities

Let us take a step back, and ask: What do we *want* from a unified theory of epistemic and practical rationality? The answer, I take it, is something like this: A unified theory of epistemic and practical rationality should allow us to locate epistemic and practical

¹¹ For a response to the first challenge, see Leite (2007). For a response to the second challenge, see Rinard (2019b). For a response to the third challenge, see Steglich-Petersen and Skipper (2019; 2020).

rationality within a single, coherent picture of rationality; it should, to use a familiar phrase from Sellars, tell us how epistemic and practical rationality “hang together.”

One way to achieve this goal, as we have seen, is to reduce epistemic rationality to practical rationality, or *vice versa*. But reduction is not the only path to unification. Instead, one might try to unify epistemic and practical rationality by showing how both notions are related to a *third* notion. This kind of strategy is very familiar from the study of *causal* relationships. When we have reason to think that two correlated variables do not cause each other directly, we may instead try to understand their relationship by searching for a *common cause*. In what follows, I will pursue a similar strategy: I will try to understand the relationship between epistemic and practical rationality by identifying a *common ground* of both notions.

The common ground I have in mind turns crucially on the notion of *evidential probability*, so let me begin by introducing that notion. It is a philosophical commonplace that evidential support comes in degrees: a proposition can be more or less strongly supported by a body of evidence. For example, the fact that *it is cloudy outside* might support the proposition that *it will rain* to a fairly high degree, whereas the fact that *it is sunny outside* might support the proposition that *it will rain* to a fairly low degree.¹² Such talk of evidential support is naturally couched in *probabilistic* terms. For example, I might say that there is a fairly high probability that *it will rain* given that *it is cloudy outside*, whereas there is a fairly low probability that *it will rain* given that *it is sunny outside*. In saying so, I am deploying something like the notion of evidential probability I want to invoke here.

In more general terms, we can say that the evidential probability of a proposition, p , given a body of evidence, e , encapsulates the degree to which e “supports” or “tells in favor of” p . Entailment is standardly taken to be a limiting case of this support relation: if e entails p , e supports p to the highest possible degree, and so we say that the evidential probability of p given e is 100%. Logical incompatibility is the opposite limiting case: if e is logically incompatible with p (that is, if e entails not- p), e supports p to the lowest possible degree, and so we say that the evidential probability of p given e is 0%. In-between these two extremes, we find a spectrum of degrees of evidential support, which we represent by letting the evidential probability of p given e take on any value between 0% and 100%.

Beyond these preliminary remarks, I will not say much to characterize the notion of evidential probability in precise terms. In particular, I will not make any substantive assumptions about what grounds, in some sense, evidential probabilities, nor will I assume that evidential probabilities satisfy the probability axioms, though this is usually taken for granted. I *will*, however, assume that the notion of evidential probability is to

¹² Strictly speaking, such claims about evidential support should be relativized to a body of background evidence, but for present purposes we can leave the background evidence implicit.

be understood in *non-normative* terms. That is to say, I will assume that evidential probabilities are not identical to, or otherwise reducible to, rational credences. That way lies explanatory circularity of just the kind that led us to reject Circular Epistemic EUM in the first place.

In saying that the notion of evidential probability is to be understood in non-normative terms, I do not mean to imply that all uses of terms like “evidence” or “evident” are non-normative. A claim like “It is evident that Mr. Smith committed the murder” may well have a normative reading in some contexts. The assumption I am making here is just that we can make sense of a non-normative notion of evidential probability, even if there are legitimate normative uses of terms like “evidence” or “evident.” Of course, even this assumption may be questioned. It may not be immediately obvious that there should *be* a cogent, non-normative notion of evidential probability. Some writers have argued that there *is*;¹³ others have seemed to presuppose it;¹⁴ and yet others have voiced doubts about it.¹⁵ This is not the place to defend the overall cogency of a non-normative notion of evidential probability. Rather, the notion will earn its keep by the work it can do in helping us to resolve the conflict described above.

Consider, then, the following alternative to Circular Epistemic EUM, which is formulated in terms of evidential probabilities rather than rational credences:

Epistemic EUM: It is rational for an agent to have a given credence function if, and only if, no alternative credence function has a higher expected utility, calculated with respect to the agent’s utility function and evidential probability function.

As before, the desire-like component is given by the agent’s utility function, but the belief-like component is now given by the agent’s “evidential probability function,”

¹³ Perhaps the most explicit example is Williamson (2000, ch. 10), who argues that evidential probabilities cannot be reduced to rational credences, and who instead takes the notion of evidential probability as a primitive. Other examples include proponents of what has come to be known as the “logical” interpretation of probability, according to which the notion of probability is to be understood as a generalization of the notion of logical entailment which allows us to capture the degree of support that a piece of evidence confers on a hypothesis (a view which is usually contrasted with “subjective” interpretations of probability, according to which probabilities are to be understood as credences of some sort). Early proponents of the logical interpretation of probability include Keynes (1921), Jeffreys (1939), and Carnap (1950); more recent ones include Hawthorne (2005) and Maher (2006).

¹⁴ For example, participants in the debate over *Evidentialism*—the view, roughly, that your evidence alone determines what you should believe—presumably think (or presuppose) that we can meaningfully distinguish the question of what an agent’s evidence supports from the question of what the agent should believe; otherwise, the debate would be settled trivially in favor of Evidentialism.

¹⁵ For example, Eder (2019) argues, against Williamson, that evidential probabilities can be reduced to rational credences.

whose role it is to specify, for each proposition, how strongly the agent’s total body of evidence supports that proposition. Following Williamson (2000, ch. 10), we may define an agent’s evidential probability function by positing a “prior” evidential probability function—that is, a function specifying the evidential probability of each proposition given *no* evidence—which is then conditionalized on the agent’s total body of evidence. But those finer details will not matter for what follows.

If we combine Epistemic EUM with Practical EUM, what emerges is a novel picture of the relationship between epistemic and practical rationality. On this picture, the most explanatorily basic components are the agent’s utility function and the agent’s evidential probability function. Together, these two functions determine a unique expected utility of each credence function, which, in turn, determines what credence function is rational. Thus determined, the rational credence function then combines in the usual way with the agent’s utility function to determine the expected utility of each available action, which, in turn, determines what action is rational. The result is a view on which epistemic and practical rationality are both grounded in a combination of utilities and evidential probabilities (see Figure 2 for an illustration). I will refer to this view of the relationship between epistemic and practical rationality as *Non-Reductive Pragmatism*.¹⁶




Figure 2: *Non-Reductive Pragmatism*

Why should we find Non-Reductive Pragmatism attractive? Because it ticks all the boxes we have been wanting to tick. First, it treats epistemic rationality as explanatorily prior to practical rationality and hence does justice to the idea that what it is rational to do depends, in part, on what it is rational to believe. Second, it avoids explanatory circularity, since it does not treat epistemic rationality as a form of practical rationality. Third, it offers a unified picture of epistemic and practical rationality: it tells us how

¹⁶ In calling the view “Non-Reductive Pragmatism,” I do not mean to suggest that the view has no reductive ambitions. In fact, the view purports to reduce both epistemic and practical rationality to a combination of utilities and evidential probabilities. However, unlike the view I have called “Reductive Pragmatism,” it does not purport to reduce epistemic rationality to practical rationality. This is what the label “Non-Reductive Pragmatism” is meant to signify.

epistemic and practical rationality fit into a single, coherent picture of rationality. And fourth, it grounds epistemic rationality in broadly pragmatic considerations and hence offers a natural explanation of why we should care about being epistemically rational in the first place.

To my mind, these are weighty reasons in favor of Non-Reductive Pragmatism. But there are also worries one might have about the view. I will turn to these in the next section. Before we proceed, however, I want to emphasize once again that Non-Reductive Pragmatism is intended as a fairly schematic view. For the sake of concreteness, I have used a particular set of norms—Epistemic EUM and Practical EUM—to illustrate the theoretical benefits of Non-Reductive Pragmatism. But any pair of norms which conform to the structure depicted in Figure 2 will provide us with the same theoretical benefits: they will offer a non-circular, unified view of epistemic and practical rationality, which does justice to the idea that what it is rational to do depends on what it is rational to believe, and which offers a natural pragmatist explanation of why we should care about being epistemically rational in the first place.

4. Three Objections

On to objections, then. The aim of this section is to address three potential worries one might have about the view outlined above. But first I want to set aside a cluster of worries that, although interesting and important in their own right, are too general to fall within the purview of this paper. As previously mentioned, epistemologists have levelled a number of objections against Reductive Pragmatism, some of which may be levelled against Non-Reductive Pragmatism as well. For example, those who reject Reductive Pragmatism on the grounds that one cannot form beliefs in response to practical considerations will presumably want to reject Non-Reductive Pragmatism on similar grounds (Kelly 2002; Shah 2006).

More generally, there are worries one might have about *any* view which says that practical considerations can make a difference to what one should believe. I will not try to address such broad worries here. If some (or all) of them should turn out to be sound, I do not think there is any denying that it would be a problem for Non-Reductive Pragmatism. But it would not be a problem that proponents of Reductive Pragmatism did not already face.

By contrast, the worries to be addressed below pertain specifically to Non-Reductive Pragmatism as I have described it. As we will see, I do not think the worries are ultimately convincing. But seeing why this is so will, I hope, help clarify what has been said so far as well as shore up some additional motivation for the central proposal of the paper.

4.1 Why not let practical rationality depend directly upon evidential probabilities?

The first worry I want to discuss starts from a correct observation, namely that Non-Reductive Pragmatism posits a certain kind of asymmetry between, on the one hand, the way in which evidential probabilities help determine what it is rational to *believe*, and, on the other hand, the way in which evidential probabilities help determine what it is rational to *do*. Consider again our two norms Epistemic EUM and Practical EUM. According to Epistemic EUM, the evidential probability function helps determine what it is rational to believe in a fairly straightforward way: it, in combination with the agent's utility function, determines the expected utility of each credence function. By contrast, according to Practical EUM, the evidential probability function only plays an indirect role in determining what it is rational to do: it helps determine the rational credence function, which, in turn, helps determine what action has the highest expected utility.

This asymmetry may seem puzzling. After all, we have already made good use of evidential probabilities in explaining what it is rational to *believe*. Why not put them to similar use in explaining what it is rational to *do*?

To give this worry its full force, it is worth considering what the resulting theory would look like if we were to let practical rationality depend directly upon evidential probabilities. Here is what the relevant alternative to Practical EUM would look like:

Evidence-Based Practical EUM: It is rational for an agent to perform a given action if, and only if, no alternative action has a higher expected utility, calculated with respect to the agent's utility function and evidential probability function.

According Evidence-Based Practical EUM, an agent's utility function and evidential probability function together determine what it is rational for the agent to do in precisely the same way that, according to Epistemic EUM, they determine what it is rational for the agent to believe. Thus, if we combine Evidence-Based Practical EUM with Epistemic EUM, we get a fully symmetric treatment of epistemic and practical rationality: both notions are grounded in exactly the same way in an agent's utility function and evidential probability function (see Figure 3 for an illustration). Let us call this alternative view *Symmetric Non-Reductive Pragmatism*.




Figure 3: *Symmetric Non-Reductive Pragmatism*

The first thing to observe about Symmetric Non-Reductive Pragmatism is that it shares many of the attractions of Non-Reductive Pragmatism: it avoids explanatory circularity; it provides us with a unified view of epistemic and practical rationality; and it grounds epistemic rationality in broadly pragmatic factors and hence offers a natural explanation of why we should care about being epistemically rational in the first place.

However, Symmetric Non-Reductive Pragmatism does not treat epistemic rationality as explanatorily prior to practical rationality and hence does not do justice to the idea that what it is rational to do depends, in part, on what it is rational to believe. Instead, Symmetric Non-Reductive Pragmatism requires us to adopt what I take to be a rather revisionary understanding of the relationship between belief and action: it requires us to abandon the idea that one of the primary functional roles of belief is to guide action. One might perhaps try to argue that such a revisionary move is called for. But those who are drawn to the idea that what it is rational to do depends, in part, on what it is rational to believe should, I think, take this as a reason to favor Non-Reductive Pragmatism over Symmetric Non-Reductive Pragmatism.

4.2 Why not let epistemic rationality depend only on evidential probabilities?

The second worry starts from another correct observation, namely that many of the considerations that led us to favor Reductive Pragmatism over Non-Reductive Pragmatism in the first place might just as well be taken to support a simple evidentialist view, according to which what it is rational to believe depends only on the evidential probabilities. Here is the relevant alternative to Epistemic EUM:

Evidentialism: It is rational for an agent to have a given credence function if, and only if, the credence function is identical to the evidential probability function in the agent's situation.

If we combine Evidentialism with Practical EUM, we get a view of the relationship between epistemic and practical rationality on which the evidential probability function alone determines the rational credence function, which then combines in usual way with the utility function to determine what action is rational (see Figure 4 for an illustration). With a slight abuse of terminology, we may refer to this view of the relationship between epistemic and practical rationality as *Evidentialism*.




Figure 4: *Evidentialism*

Like Symmetric Non-Reductive Pragmatism, Evidentialism shares many of the same attractions as Non-Reductive Pragmatism: it avoids explanatory circularity; it treats epistemic rationality as explanatorily prior to practical rationality; and it provides a unified picture of epistemic and practical rationality.

However, unlike Non-Reductive Pragmatism, Evidentialism does not ground epistemic rationality in pragmatic considerations and hence does not offer a simple pragmatist explanation of why we should care about being epistemically rational in the first place. Of course, a proponent of Evidentialism can still say that evidentially supported beliefs tend to be accurate, and that accurate beliefs tend to help us reach our practical goals in life. But if Evidentialism is correct, this is, in a sense, just a happy coincidence: it does not help explain why we should believe what we should believe. To my mind, this is something of an embarrassment for Evidentialism.¹⁷ At the very least, it leaves us with a challenge of explaining why we should care about whether our beliefs are rational or not. Perhaps this challenge can be met. I do not want to foreclose the possibility of telling a plausible story about why we should care about forming evidentially supported beliefs which does not ultimately bottom out in practical considerations. But I do not myself see how such a story would go. Those who are in the same boat should take this as a reason to prefer Non-Reductive Pragmatism over Evidentialism.

4.3 Does postulating evidential probabilities go against core pragmatist motivations?

I have said that Non-Reductive Pragmatism should be attractive to those who have traditionally been drawn to Reductive Pragmatism. This may seem doubtful in light of the fact that Non-Reductive Pragmatism relies crucially on a notion of evidential probability. “Granted,” our objector might say, “Non-Reductive Pragmatism shares many of the attractions of Reductive Pragmatism; but shouldn’t proponents of Reductive Pragmatism be reluctant to embrace a purely epistemic notion like that of evidential probability?”

Something like this worry has been voiced by Rinard (2017, p. 134) and Christensen (2020, §3), both of whom criticize a norm very similar to Epistemic EUM on the grounds that it relies on a purely epistemic notion.¹⁸ Christensen puts the worry by saying that endorsing such a norm would “go against fundamental motivations of avoiding reliance in our theorizing on a distinctive evidentially-based, non-pragmatic dimension of normativity” (Christensen 2020, p. 13). Similarly, Rinard opposes such a

¹⁷ See Cowie (2014) for a discussion of this point.

¹⁸ I should perhaps insert here that the projects of Rinard and Christensen are otherwise very different: Rinard argues that there is no distinctively epistemic form of rationality, whereas Christensen argues that there is.

norm by reference to the fact that it “involves a constitutive connection between rationality and evidence” (Rinard 2017, p. 134).

I think there is something right about the sentiment expressed here. Someone who is committed to thinking that epistemic rationality must depend only on pragmatic factors should indeed find Epistemic EUM untenable. But I do not think proponents of Reductive Pragmatism need be committed to thinking that epistemic rationality must depend *only* on pragmatic factors. After all, as we have seen, Non-Reductive Pragmatism does justice to many of the same motivations that led philosophers to endorse Reductive Pragmatism in the first place. So, if Non-Reductive Pragmatism goes against core pragmatist motivations, the motivations must be different from the ones I have been considering here. Perhaps such motivations exist. But I do not myself see any reason why proponents of Reductive Pragmatism cannot readily grant the idea that different bodies of evidence can support different propositions to different degrees. Doing so does not commit us to a distinctively epistemic kind of normativity, insofar as the relevant notion of evidential support is a non-normative one. Anyone is of course free to stipulate a sense of “epistemic rationality” in which an agent is epistemically rational if, and only if, the agent’s credences match the evidential probabilities. But such a terminological maneuver should not, it seems to me, worry the Reductive Pragmatist. The resulting notion of epistemic rationality would not be any more normative than the notion of evidential probability with which we began.

A final point before wrapping up: it is worth emphasizing that nothing I have said here prevents epistemologists from paying special attention to questions about what our evidence supports independently of any practical concerns. Indeed, if what I have said is correct, understanding the nature of evidential support is a crucial part of understanding epistemic rationality itself. As such, the central proposal of the paper leaves much of epistemological theorizing unscathed. What it *doesn’t* leave unscathed are certain traditional views of the relationship between epistemic and practical rationality, such as Evidentialism and Reductive Pragmatism.

5. Conclusion

I began this paper with a conflict. On the one hand, many theories of rational action treat epistemic rationality as *explanatorily prior* to practical rationality: they say that what it is rational to *do* in a given situation depends, in part, on what it is rational to *believe* in that situation. On the other hand, many epistemologists have argued that epistemic rationality is just a *form* of practical rationality. Yet, both claims cannot be true, on pain of explanatory circularity.

My aim has been to propose a view of the relationship between epistemic and practical rationality—*Non-Reductive Pragmatism*—which promises to resolve the conflict in a way that should be attractive to both sides. The central idea was to ground both

epistemic and practical rationality in a notion of evidential probability, which is itself to be understood in non-normative terms. Doing so, I argued, allows us to achieve a number of things: to avoid explanatory circularity; to do justice to the idea that what it is rational to do depends, in part, on what it is rational to believe; to provide a unified theory of epistemic and practical rationality; and to offer a natural pragmatist explanation of why we should care about being epistemically rational in the first place.

This all sounds quite compelling to me. But I don't want to overstate the probative force of what has been said. As previously mentioned, many of the criticisms that epistemologists have levelled against Reductive Pragmatism may be levelled against Non-Reductive Pragmatism as well. Thus, when I say that Non-Reductive Pragmatism fares better than Reductive Pragmatism, it is not because Non-Reductive Pragmatism smotheres all doubts one might have about a broadly pragmatist view of epistemic rationality. Rather, it is because Non-Reductive Pragmatism offers what seems to me to be a uniquely attractive solution to a deep tension in our understanding of the relationship between epistemic and practical rationality.

Acknowledgments: An earlier version of this paper was presented at the Goethe University of Frankfurt. I would like to thank the audience on that occasion for constructive feedback. Thanks also to Zach Barnett, Bob Beddor, Ben Blumson, Dominik Kauss, Michael Pelczar, Lavinia Picollo, Abelard Podgorski, Eyal Tal, Weng Hong Tang, David Thorstad, Dan Waxman, and two exceptional referees for *Mind* for helpful comments and suggestions. This research was supported by a Tier 1 research grant from the Ministry of Education, Singapore.

References

- Alston, W. (2005): *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.
- Berker, S. (2018): "A Combinatorial Argument Against Practical Reasons for Belief." In: *Analytic Philosophy* 59, pp. 427-70.
- Carnap, R. (1950): *Logical Foundations of Probability*, first edition, Chicago: University of Chicago Press.
- Christensen, D. (2020): "The Ineliminability of Epistemic Rationality." In: *Philosophy and Phenomenological Research*. Online first.
- Côté-Bouchard, C. and C. Littlejohn (2018): "Knowledge, Reasons, and Errors about Error Theory." In *Metaepistemology: Realism and Anti-Realism*, C. Kyriacou and R. McKenna (eds.), Palgrave Macmillan.
- Cowie, C. (2014): "In Defense of Instrumentalism About Epistemic Normativity." In: *Synthese* 191, pp. 4003-17.
- Eder, A.-M. A. (2019): "Evidential Probabilities and Credences." In: *British Journal for the Philosophy of Science*. Online first.

- Foley, R. (2001): "The Foundational Role of Epistemology in a General Theory of Rationality." In Fairweather and Zagzebski (eds.), *Virtue Epistemology: Essays in Epistemic Virtue and Responsibility*. Oxford: Oxford University Press.
- (1987): *The Theory of Epistemic Rationality*. Cambridge: Harvard University Press.
- Friedman, J. (2018): "Junk Beliefs and Interest-Driven Epistemology." In: *Philosophy and Phenomenological Research* 97, pp. 568-83.
- Grimm, S. (2009). Epistemic Normativity. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford: Oxford University Press
- Goldman, A. (1999): *Knowledge in a Social World*. Oxford: Oxford University Press.
- Hawthorne, James (2005): "Degree-of-Belief and Degree-of-Support: Why Bayesians Need Both Notions." In: *Mind* 114, pp. 277-320.
- Jeffreys, H. (1939/1998): "Theory of Probability," reprinted in *Oxford Classics in the Physical Sciences*. Oxford: Oxford University Press.
- Kelly, T. (2003): "Epistemic Rationality as Instrumental Rationality: A Critique." In: *Philosophy and Phenomenological Research* 66, pp. 612-40.
- (2002): "The Rationality of Belief and Other Propositional Attitudes." In: *Philosophical Studies* 110, pp. 163-96.
- Keynes, J. M. (1921): *A Treatise on Probability*. London: Macmillan and Co.
- Kornblith, H. (1993): "Epistemic Normativity." In: *Synthese* 94, pp. 357-76.
- Laudan, L. (1990): "Normative Naturalism." In: *Philosophy of Science* 57, pp. 44-59.
- Leite, A. (2007): "Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly's 'Epistemic Rationality as Instrumental Rationality: A Critique'." In: *Philosophy and Phenomenological Research* 75, pp. 456-64.
- Lynch, M. (2004): *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.
- Maguire, B. and J. Woods (2020): "The Game of Belief." In: *Philosophical Review* 129, pp. 211-49.
- Maher, P. (2006): "The Concept of Inductive Probability." In: *Erkenntnis* 65, pp. 185–206.
- Nolfi, K. (2019): "Epistemic Norms, All Things Considered." In: *Synthese*. Online first.
- Nozick, R. (1993): *The Nature of Rationality*. Princeton University Press.
- Papineau, D. (1999): "Normativity and Judgement." In: *Proceedings of the Aristotelian Society*, supplementary vol. 73, pp. 17-41.
- Quine, W. V. (1969): "Epistemology Naturalized." In: *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Reisner, A. (2018): "Two Theses about the Distinctness of Practical and Theoretical Normativity." In *Normativity: Epistemic and Practical*, Way, J., C. McHugh, and D. Whiting (eds.). Oxford: Oxford University Press.
- Rinard, S. (2019a): "Equal Treatment." In: *Philosophical Studies* 176, pp. 1923-50.
- (2019b): "Believing for Practical Reasons." In: *Noûs* 53, pp. 763-84.

- (2017): “No Exception for Belief.” In: *Philosophy and Phenomenological Research* 94, pp. 121-43.
- Shah, N. (2006): “A New Argument for Evidentialism.” In: *The Philosophical Quarterly* 56, pp. 481-98.
- Steglich-Petersen, A. (2018): “Epistemic Instrumentalism, Permissibility, and Reasons for Belief.” In *Normativity: Epistemic and Practical*, C. McHugh, J. Way, and D. Whiting (eds.). Oxford: Oxford University Press.
- (2011): “How to Be a Teleologist About Epistemic Reasons.” In A. Steglich-Petersen and A. Reisner (eds.), *Reasons for Belief*. Cambridge University Press.
- (2006): “No Norm Needed: On the Aim of Belief.” In: *The Philosophical Quarterly* 56, pp. 499-516.
- Steglich-Petersen, A. and M. Skipper (2020): “Instrumental Reasons for Belief: Elliptical Talk and Elusive Properties.” In S. Schmidt and G. Ernst (eds.), *The Ethics of Belief and Beyond. Understanding Mental Normativity*. Abingdon: Routledge.
- (2019): “An Instrumentalist Account of How to Weigh Epistemic and Practical Reasons for Belief.” In: *Mind* 129, pp. 1071-94.
- Thoma, J. (2019): “Decision Theory.” In R. Pettigrew and J. Weisberg (eds.), *The Open Handbook of Formal Epistemology*. PhilPapers Foundation, pp. 57-106.
- Williamson, T. (2000): *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Worsnip, A. (2021): “Can Pragmatists Be Moderate?” In: *Philosophy and Phenomenological Research* 102, pp. 531-58.