

Belief, Credence, and Pragmatic Encroachment¹

JACOB ROSS

University of Southern California

MARK SCHROEDER

University of Southern California

This paper compares two alternative explanations of pragmatic encroachment on knowledge (i.e., the claim that whether an agent knows that p can depend on pragmatic factors). After reviewing the evidence for such pragmatic encroachment, we ask how it is best explained, assuming it obtains. Several authors have recently argued that the best explanation is provided by a particular account of belief, which we call pragmatic credal reductivism. On this view, what it is for an agent to believe a proposition is for her credence in this proposition to be above a certain threshold, a threshold that varies depending on pragmatic factors. We show that while this account of belief can provide an elegant explanation of pragmatic encroachment on knowledge, it is not alone in doing so, for an alternative account of belief, which we call the reasoning disposition account, can do so as well. And the latter account, we argue, is far more plausible than pragmatic credal reductivism, since it accords far better with a number of claims about belief that are very hard to deny.

Recently, a quandary has arisen in the literature on pragmatic encroachment. On the one hand, it has been forcefully argued that there is pragmatic encroachment on *knowledge*, or in other words that whether an agent knows that p can depend on pragmatic factors such

¹ Thanks to Lara Buchak, Kenny Easwaran, Richard Holton, Shieva Kleinschmidt, Matthew McGrath, Julia Staffel, Brian Weatherson, Ralph Wedgwood, Jonathan Weisberg, an audience at Northwestern University, participants in the 2011 Formal Epistemology Workshop, and participants in Jake Ross's spring 2011 graduate seminar, for helpful discussion and/or comments.

as the costs of acting as if it were true that p when it is not true. And it has been argued that the best way to explain pragmatic encroachment on knowledge is by assuming that pragmatic factors affect *justified belief*. On the other hand, it is widely held that there is no pragmatic encroachment on justified *degrees of belief*, or levels of confidence, as the latter, it is maintained, should be strictly apportioned to the evidence.² But this creates a problem. How can pragmatic factors affect whether we are justified in believing a proposition (as opposed to disbelieving it or withholding judgment concerning it) without affecting the degree of belief or level of confidence that we are justified in having in it? In other words, how can pragmatic factors affect our justification for *outright belief* without affecting our justification for *degrees of belief*?

A number of authors, including Brian Weatherson [2005], Dorit Ganson [2008], and, on a natural reading, Jeremy Fantl and Matthew McGrath [2010], have proposed a solution to this problem. According to their proposal, believing a proposition simply consists in having a sufficiently high level of confidence in it, but the level of confidence one must have in a proposition to count as believing it depends on pragmatic factors. On this view, pragmatic factors can affect whether an agent is justified in believing a proposition, not by affecting what level of confidence she is justified in having in it, but rather by affecting whether the agent's having this level of justified confidence is sufficient, under the circumstances, for the agent to count as believing the proposition. We will call this view *Pragmatic Credal Reductivism* (PCR)—'credal reductivism' because it understands beliefs in terms of credences or levels of confidence, and 'pragmatic' because it implies that the level of confidence that an agent must have in a given proposition to count as believing it depends on pragmatic factors.³

In this paper, we will argue against PCR, and we will propose an alternative explanation of the data it was meant to explain. The paper consists of two main parts. In the first, we set out two alternative accounts of belief: the pragmatic credal reductivist account, and an alternative account that we call the *reasoning disposition account*. We show how each of these accounts can explain the possibility of pragmatic encroachment on knowledge. We then argue, in the second part, that there is strong reason to prefer the reasoning disposition account of belief. We do so by presenting four claims that we take to be central to our ordinary notion of belief, and arguing that the

² For a dissenting opinion, however, see Stanley [2005], esp. pp. 88–89.

³ See Fantl and McGrath [2002] for a similar argument.

reasoning disposition account helps to explain the truth of each of these claims, whereas PCR conflicts with each of them.

1. Two Explanations of Pragmatic Encroachment

In the first two sections of part 1, we review some reasons for holding that there is pragmatic encroachment on knowledge, and we show how PCR can explain such pragmatic encroachment. And in the remaining three sections, we present an alternative account of belief, the reasoning disposition account, and we show how it, too, can explain pragmatic encroachment on knowledge, but in a different way.

1.1. *Pragmatic Encroachment and the Knowledge Action Principle*

Consider the following pair of cases.⁴

Low: Five minutes ago, Hannah made three sandwiches and placed them in the refrigerator. She told Sarah that she placed the peanut butter sandwich on the left, the tuna sandwich in the middle, and the almond butter sandwich on the right. Hannah then departed just as Sarah's friend Almira arrived for lunch. Sarah knows that Almira has no allergies. Almira says: "I'd love an almond butter sandwich." And so Sarah opens the refrigerator door, points to the sandwich on the right, and says: "The sandwich on the right is an almond butter sandwich. You can have it."

High: This case is just like Low, except here it is Sarah's nephew Algernon who is visiting for lunch, and he has a severe peanut allergy. He asks Sarah for a sandwich. Sarah knows that the peanut butter sandwich would be fatal to Algernon, but that the almond butter sandwich would be harmless. She also knows that he would slightly prefer the almond butter sandwich to the tuna sandwich. When Sarah goes to the fridge, she can tell, by visual inspection, which is the tuna sandwich, but she cannot tell, by visual inspection, which is the peanut butter sandwich and which is the almond butter sandwich. So she gives him the tuna sandwich.

Let r be the proposition that the sandwich on the right is the almond butter sandwich. In Low, it seems appropriate for Sarah to assert that r to Almira, since Sarah knows that r .⁵ This seems appropriate because

⁴ Similar cases can be found in DeRose [1992], Cohen [1999], Stanley [2005] and Fantl and McGrath [2010].

⁵ Throughout the paper, we will use italicized lower-case letters as schematic sentence letters, and use invisible 'propositional quotes' to turn a sentence into a referring term denoting the proposition expressed by utterances of that sentence. Please note the invisible propositional quotes wherever appropriate, as in the first appearance of ' r ' in the main text.

it seems Sarah knows that r on the basis of Hanah's testimony. In High, however, the situation seems rather different. Here we are inclined to say that it is appropriate for Sarah to give Algernon the middle sandwich rather than the sandwich on the right, since she knows that the middle sandwich is the tuna sandwich, but she doesn't know that the sandwich on the right is the almond butter sandwich—it might instead be the fatal peanut butter sandwich. And so, in Low, we are inclined to say that Sarah knows that r , whereas in High we are inclined to say the opposite. And yet Sarah's evidence bearing on r appears to be the same in both cases.

Here's a natural explanation of the difference between these two cases.

Knowledge Action Principle: For any agent S and proposition p , if S is in a choice situation in which S could not rationally act as if p , then S does not know that p (where to *act as if* p is to act in the manner that would be rationally optimal on the supposition that p is true).⁶

In Low, acting as if r (by affirming that r) would be perfectly rational, given Sarah's strong evidence for r . By contrast in High, acting as if r (by giving Algernon the sandwich on the right) would be irrational, given the disastrous consequences of doing so if r is false. Thus, the Knowledge Action Principle is consistent with Sarah's knowing that r in Low, but it is inconsistent with her knowing that r in High.

Moreover, even independently of our intuitions about such cases as High and Low, the Knowledge Action Principle has considerable plausibility. For it is very plausible that if an agent knows that p , then she could rationally rely on p by reasoning on its basis. And it likewise seems plausible that if an agent can rationally rely on p in her reasoning, then she can rationally act as if p , since that's how she would choose to act were she to reason rationally while relying on p . And so it is very plausible that if an agent knows that p , then she can rationally act as if p .

We will assume, for the sake of argument, that the Knowledge Action Principle is true, and turn our attention to the question of how its truth can best be explained.⁷ One way to explain the Knowledge Action Principle, as we will see in the next section, is on the basis of PCR.

⁶ Similar principles are endorsed in Williamson [1999], Fantl and McGrath [2002] and [2010], Hawthorne [2004], and Hawthorne and Stanley [2008].

⁷ Note that DeRose [1992] and Cohen [1999], among others, take such cases to instead be explained by the thesis that knowledge-ascriptions are context-dependent. For an extensive discussion of the relative merits of these two kinds of explanation, see Stanley [2005].

1.2. How Pragmatic Credal Reductivism can Explain Pragmatic Encroachment

Here's one possible explanation of the Knowledge Action Principle. In order to know that p , an agent must be justified in believing that p . But to believe that p is simply to have a sufficiently high credence in p to count as believing that p under the circumstances. Hence, in order to know that p , one must be *justified* in having a sufficiently high credence in p to count as believing that p under the circumstances. Accordingly, in order to explain the Knowledge Action Principle, all we need is an account of belief on which the minimum level of credence an agent must have to count as believing that p under given circumstances is at least as high as the minimum level of credence the agent would need in order for it to be rational for her to act as if p under those circumstances. For, given such an account of belief, it will follow from the above argument that in order to know that p , one must be justified in having sufficient credence in p to make it rational for one to act as if p . Hence it will follow that one will know that p only if it is rational for one to act as if p .

Weatherson [2005], Ganson [2008], and on a natural reading, Fantl and McGrath [2010] each propose an account of belief that allows for such an explanation of the Knowledge Action Principle.⁸ And the accounts they propose all share a common form. To a first approximation, they all maintain that what it is to believe that p is to have a sufficiently high credence in p to rationalize acting as if p when choosing among relevant actions under relevant circumstances—where the relevant circumstances and actions include, but may not be limited to, the agent's actual circumstances and the actions available therein. Thus, Ganson [2008] suggests that what is required for “outright believing that p ” is that one be “willing to act as if p in all or most circumstances,” and she adds that “for those circumstances where our degree of belief isn't high enough [to rationalize acting as if p], we simply fail to count as believing that p in those circumstances” (p. 451). Similarly, Fantl and McGrath [2010] propose a conception of belief according to which “you believe that p iff you have enough credence in p for p to be your motivating reason for ϕ -ing, for all relevant ϕ ” (p. 160).

⁸ Some of Fantl and McGrath's remarks look clearly like an endorsement of pragmatic credal reductivism. For instance, they say, “you believe that p iff you have enough credence in p for p to be your motivating reason for ϕ -ing, for all relevant ϕ ” (p. 160). Other remarks, as in their claim that if you believe that p “then you are prepared to put p to work as a basis for what you do, believe, etc” (p. 143) look more like glosses on the alternative, reasoning disposition, account that we will offer in this paper, and in conversation, Matthew McGrath has suggested sympathies with our alternative account.

Weatherson proposes a conception of belief that he expresses rather differently, but it amounts to a view of the same general type. He identifies believing that p with being such that conditionalizing on p wouldn't affect one's conditional preferences over relevant actions. More precisely, on Weatherson's view, S believes that p just in case, for any actions ϕ and ψ in the appropriate class of actions, and any proposition q in the appropriate class of propositions, S prefers ϕ to ψ conditional on q if and only if S prefers ϕ to ψ conditional on the conjunction of p and q (p. 422). Note however that, given the rationality assumption that Weatherson makes throughout, this is equivalent to saying that S believes that p just in case S 's credence in p is sufficiently high to rationalize acting as if p in any circumstance in which S is choosing between some pair of actions ϕ and ψ (belonging to the appropriate class of actions) having learned some proposition q (belonging to the appropriate class of propositions). Thus, Weatherson's account of belief, like those of Ganson and of Fantl and McGrath, can be understood as one according to which to believe that p is to have sufficiently high credence in p to rationalize acting as if p in the relevant choice situations (i.e., when choosing among the relevant actions under the relevant circumstances).⁹ Henceforth, we use the expression 'pragmatic credal reductivism' (PCR) to refer to this kind of view.

1.3. *The Reasoning Disposition Account of Belief*

While PCR can explain the Knowledge Action Principle, it is not the only account of belief that can do so. Here we will propose another. We will begin by discussing the relation of *treating a proposition as true in reasoning*, and then we'll propose an account of belief in terms of this relation. We will thus be in a position to examine the explanatory potential of this account of belief.

Consider, first, what it is to treat a proposition as true in practical reasoning, since this kind of reasoning will be our main focus in this paper. In the context of practical reasoning, we may say that an agent *treats* a given proposition p as true just in case she evaluates her alternatives by the same procedure by which she would evaluate them conditional on p . As an illustration, consider the following case. Suppose Renzo has rented a DVD from a store on Canal St., and the DVD is due before the store closes twenty minutes hence. Renzo is deciding

⁹ Weatherson [Unpublished] proposes a revised version of the account offered in Weatherson [2005]. See Weatherson [2011] and Ross [2011] for a debate concerning of whether this revised account solves the problems for the original account that we present below.

which train to go to the store by, the Broadway train or the Canal St. Express. He reasons as follows:

If I take the Canal St. Express, it will cost me \$3, but I'll get to the store on time and so I won't be fined. Thus, I'll be out \$3. If, on the other hand, I take the Broadway train, it will cost me \$2. If it stops at Canal St., I'll get to the store on time and won't be fined, so I'll be out only \$2. But if the Broadway train doesn't stop at Canal St., I won't get to the store on time, and I'll be fined \$5, so I'll be out \$7. Since it's as likely as not that the Broadway train won't stop at Canal St., it isn't worth risking the fine to save \$1 on the train, so I'll take the Canal St. Express.

So described, Renzo's decision problem can be represented by the following decision matrix:

TABLE 1

	Broadway train stops at Canal St.	Broadway train doesn't stop at Canal St.
Take the Broadway train	Pay \$2 and avoid fine; out \$2.	Pay \$2 and incur fine; out \$7.
Take the Canal St. Express	Pay \$3 and avoid fine; out \$3.	Pay \$3 and avoid fine; out \$3.

In this table, the rows represent acts, the columns represent the states of nature (i.e., the relevantly different ways the world might be independently of which act is chosen), and the cell at the intersection of a given act and a given state of nature represents the outcome of taking this act on the condition that this state of nature obtains. If Renzo decides among his options in the manner represented in Table 1, then for any act and state of nature in the table, Renzo will be *treating it as true* that performing the action in question in the state of nature in question would result in the outcome specified in the table. Thus, he will be treating as true the proposition that, if he takes the Broadway train and it stops at Canal St., then he will pay \$2 and avoid the fine, and so he'll be out \$2 (call this proposition *b*). For he will be evaluating his acts in the same manner in which he would evaluate them conditional on *b*.

Note, however, that if Renzo is a reasonable human being with ordinary evidence, then he won't be absolutely certain that *b* is true. For he will have nonzero credence in a number of alternative possibilities, such as the possibility that the price of taking the Broadway train has changed, the possibility that the Broadway train is running too late to get him to the store on time even if it stops at Canal St, the possibility that if he takes the Broadway train he'll be mugged and lose all his money, etc. Nonetheless, it is not unreasonable for Renzo to ignore

these possibilities, and instead treat b as true in deciding how to act. After all, if he were to take into account every relevant possibility in which he has nonzero credence, then his decision problem would be unmanageable, and the store would be closed long before he could decide between his options.

The need to treat uncertain propositions as true has long been acknowledged by decision theorists. Leonard Savage makes this point very clearly. According to his formulation of decision theory, decision problems are defined by the kind of table illustrated above, in which ‘consequences’ (i.e., outcomes) are assigned to pairs consisting of acts and states of nature. Thus, the agent can calculate the *expected utility* of her possible acts (about whose actual utility she is uncertain) in terms of the *actual utilities* of the possible consequences of these acts, multiplied by the probabilities of the states of nature in which these acts would have these consequences. The general problem, of which we observed one instance above, is that for any given act, an agent will typically have nonzero credence in vastly many possible consequences of this act. And so if she were to associate a given consequence with a given act-state pair only if she were certain that the act-state pair would have this consequence, then she would need to employ a vast partition of ultrafine-grained states of nature, and the resulting computational task would be unmanageable. Thus, Savage concludes that, inevitably, act-state pairs with “actually uncertain consequences [must] play the role of sure consequences” Savage [1972, p. 84]. Similar considerations apply, as James Joyce [1999] has shown, on other formulations of decision theory.¹⁰

In theoretical reasoning, too, it is often necessary to treat uncertain propositions as true, in the sense of evaluating possibilities in the same manner in which we would evaluate them conditional on these uncertain propositions. As an illustration, suppose Rainer knows that the LA Dodgers will be playing against the Seattle Mariners, and he wants to determine the probability of the following proposition (call it g): that it will rain during the game. Suppose, however, that Rainer doesn’t know which team will have home field advantage. In this case, he might arrive at the probability of g by taking the weighted average of the probability of g conditional on the game being in LA and the probability of g conditional on the game being in Seattle. In so doing, however, he would be treating as true, in the sense just defined, the proposition that the game will be held in either LA or Seattle. And yet, if he is rational and has ordinary evidence, then he will have nonzero credence in a multitude of alternative possibilities.

¹⁰ See especially sections 2.6, 4.2, 5.5 and 7.1 of Joyce [1999].

It seems, therefore, that in virtue of our limited cognitive resources, we cannot avoid the heuristic of treating as true propositions about which we are uncertain. Nor can we, in every instance, first *reason* about the employment of this heuristic before employing it. For any agent who always reasoned before employing this heuristic would be faced with a dilemma. On the one hand, if she always employed the heuristic, then she would require an infinite regress of meta-reasoning processes. On the other hand, if this regress terminates in an initial reasoning process in which she treats no uncertain propositions as true, but instead takes every relevant possibility into account, then she would have simply replaced one cognitively intractable reasoning task with a sequence of reasoning tasks beginning with one that is cognitively intractable.¹¹

Since we must treat uncertain propositions as true, and since we must, at least sometimes, do so without first reasoning about whether to do so, it seems we must have automatic dispositions to treat some uncertain propositions as true in our reasoning. It would not make sense, however, for us to have *indefeasible* dispositions to treat these propositions as true in our reasoning. For if an agent had an indefeasible disposition to treat a proposition *p* as true, then she would act as if *p* even in a choice situation such as High, in which she has an enormous amount to lose in acting as if *p* if *p* is false, and little to gain in acting as if *p* if *p* is true. Thus, having an indefeasible disposition to treat *p* as true would make one vulnerable to acting in dangerously irrational ways.

What we should expect, therefore, is that for some propositions we would have a *defeasible* or *default* disposition to treat them as true in our reasoning—a disposition that can be overridden under circumstances where the cost of mistakenly acting as if these propositions are true is particularly salient. And this expectation is confirmed by our experience. We do indeed seem to treat some uncertain propositions as true in our reasoning; we do indeed seem to treat them as true automatically, without first weighing the costs and benefits of so treating them; and yet in contexts such as High where the costs of mistakenly treating them as true are salient, our natural tendency to treat these propositions as true often seems to be overridden, and instead we treat them as merely probable.

But if we concede that we have such defeasible dispositions to treat particular propositions as true in our reasoning, then a hypothesis naturally arises, namely, that beliefs consist in or involve such dispositions. More precisely, at least part of the functional role of belief is that believing that *p* defeasibly disposes the believer to treat *p* as true in her

¹¹ Similar points are made in Resnik [1987], pp. 10–11, and Joyce [1999], pp. 70–77.

reasoning. Let us call this hypothesis the *reasoning disposition account* of belief.¹²

1.4. *The Anti-Reductionist Implications of the Reasoning Disposition Account*

The proponents of pragmatic credal reductivism might insist that the reasoning disposition account of belief is not a genuine alternative to their account. For they all insist that their account, too, is one on which believing a proposition involves treating it as true. Thus, Weatherston [2005] states that his view “start[s] with the functionalist idea that to believe that *p* is to treat *p* as true for the purposes of practical reasoning” (p. 421). Similarly, Fantl and McGrath [2010] claim that on the kind of pragmatic view of belief they endorse, if you believe that *p* “then you are prepared to put *p* to work as a basis for what you do, believe, etc.” (p. 143). And Ganson states that her view is similar to Williamson’s view according to which “one believes *p* outright when one is prepared to use *p* as a premise in practical reasoning.” Her view accounts for this feature of belief, she says, because “our degree of willingness to use *p* as a premise in practical reasoning is a direct product of our degree of belief that *p* and our fundamental preferences.”¹³ And so the proponents of PCR might claim that their view is simply a more precise version of the reasoning disposition account, since it specifies what kind of mental state plays the functional role assigned to belief by the latter account.

But this would be a mistake. PCR is not a view on which believing that *p* necessarily involves a disposition to *treat p as true in one’s reasoning*; rather it is a view on which believing that *p* involves having sufficient credence in *p* to rationalize *acting as if p*. But there’s an important difference between these relations, as one can act as if *p* without treating *p* as true in one’s reasoning. As an illustration, consider the train example from the previous section. Suppose Renzo were to reason as follows:

Regardless of which train I take, and regardless of whether the Broadway train stops at Canal St., there is a .0001 chance that the

¹² Here we are indebted to Richard Holton, who suggested, in conversation, that the account of acceptance in Ross [2006] could be extended to provide an account of outright belief. For a somewhat similar view of belief, see Frankish [2005] and [2009]. Frankish, however, accepts many of the tenets of credal reductivism, and as a result his view faces serious problems, one of which we discuss in footnote 15, below. For a view of the relation between outright belief and credence that is much closer to our own, and for a fascinating discussion of the interrelated roles that these two kinds of mental state play in our reasoning, see Weisberg [unpublished].

¹³ Ganson [2008], p. 451, n. 9. The Williamson quote is from his [2000], p. 99.

scanner will reject my train ticket, and hence that I'll need to purchase a second ticket. If I take the Canal St. Express, I'll get to the store on time, so I won't be fined. Thus, if the scanner accepts my ticket, I'll just be out the \$3 I paid for my ticket. But if it rejects my ticket, I'll have to pay twice, so I'll be out \$6. On the other hand, if I take the Broadway train, the train will cost me either \$2 or \$4, depending on whether the scanner accepts or rejects my ticket. Hence, if the train stops at Canal St. and so I avoid the fine, then I'll be out just the \$2 or \$4 I paid for the train. But if the train doesn't stop at Canal St., I'll be out the \$2 or \$4 in addition to a \$5 late fee.

So described, Renzo's decision problem can be represented by the following matrix:

TABLE 2

	Ticket accepted; Broadway train stops at Canal St.	Ticket accepted; Broadway train doesn't stop.	Ticket rejected; Broadway train stops at Canal St.	Ticket rejected; Broadway train doesn't stop.
Take the Broadway train	Pay \$2 and avoid fine; out \$2.	Pay \$2 and incur fine; out \$7.	Pay \$4 and avoid fine; out \$4.	Pay \$4 and incur fine; out \$9.
Take the Canal St. Express	Pay \$3 and avoid fine; out \$3.	Pay \$3 and avoid fine; out \$3.	Pay \$6 and avoid fine; out \$6.	Pay \$6 and avoid fine; out \$6.

If Renzo were to evaluate his options according to Table 2, then the expected utility he would assign to taking the Broadway train would be the average of the utilities of the outcomes in the first row, weighted by the probabilities of the corresponding states of nature, and so it would be $.49995 * (-\$2) + .49995 * (-\$7) + .00005 * (-\$4) + .00005 * (-\$9) = -\$4.5002$. And the expected utility he would assign to taking the Canal St. Express would be $-\$3.0003$. Renzo would therefore choose to take the Canal St. Express. And this is also the option that is rationally optimal conditional on the following supposition (call it q) *that the scanner will accept his ticket*. Thus, if Renzo were to evaluate his options according to Table 2, then he'd end up acting as if q . But evaluating his options in this manner would not involve treating q as true in his reasoning. Rather, it would involve treating q as highly probable, while at the same time taking into account the possibility that not q .

It seems, therefore, that the advocates of PCR are wrong to think that their view captures the idea that if someone believes a proposition then she is disposed to treat it as true in her reasoning. For an agent might have sufficient credence in q to rationalize acting as if q —and hence she might count as believing that q according to PCR—and yet she may be disposed never to treat q as true in her reasoning, but

instead to always take into account the possibility that not q , in the manner illustrated in Table 2.¹⁴

We have argued that the idea that believing a proposition essentially involves an automatic disposition to treat it as true in reasoning is inconsistent with *pragmatic* credal reductivism. We will now argue, more generally, that this idea is inconsistent with *any* form of credal reductivism, that is, with any view on which to believe a proposition is to have sufficient credence in it. Why this is so can be seen if we consider an agent, Quirin, who has the following quirk. Quirin is so constituted that if his credence in a proposition is between .99 and .999999, then he will be disposed to treat this proposition as true in his reasoning, but if his credence in a proposition is greater than .999999 but less than 1, then he will not be disposed to treat it as true in reasoning, but will instead be disposed to treat it as extremely probable. Presumably it's consistent with what we've said that there is a proposition p and a circumstance C such that Quirin believes that p in C and Quirin's credence in p is between .99 and .999999. So suppose this is true. Assuming credal reductivism, *what it is* for Quirin to believe that p in C is for Quirin's credence in p in circumstance C to be at least as great as some minimal level, a minimal level that can be no greater than .999999. It follows that Quirin would count as believing that p in C if Quirin's credence in p were *above* .999999. And we have stipulated that if Quirin's credence in p were above .999999, then he would not be disposed to treat it as true in reasoning. And so it follows that if credal reductivism is true, then Quirin's believing that p does not essentially involve his being disposed to treat p as true in reasoning.

We will conclude this section with one final generalization. The basic idea of the reasoning disposition account—namely, that believing that p essentially involves having an automatic but defeasible disposition to treat p as true in reasoning—is not compatible with any view that reduced the property of believing that p to the Bayesian substratum of credences and preferences. For, where B is any property defined in terms of credences and preferences, if having property B is consistent

¹⁴ The defender of PCR might attempt to circumvent this objection by defining 'acting' sufficiently broadly to include *treating as true in reasoning* as a kind of acting. If she maintains that treating p as true is a way of acting as if p , then her view will imply that believing that p requires having sufficient credence in p to rationalize treating p as true in reasoning. Even so, however, her view still won't imply that believing that p involves having an *automatic disposition* to treat p as true in reasoning. For the manner in which we are automatically disposed to reason may differ from the manner in which it would be most rational for us to reason, given our credences and preferences. Thus, on the resulting view, an irrational agent may count as believing that p even if she has no disposition to treat p as true in reasoning, but is instead disposed to treat p as merely probable.

with having a credence of less than one in p , then it will be possible for someone with property B to be disposed to treat p , in reasoning, not as *true* but rather as (at most) *extremely probable*. And if having property B precludes having a credence of less than one in p , then having property B will involve being *indefeasibly* disposed to treat p as true in reasoning.¹⁵

1.5. How the Reasoning Disposition Account can Explain Pragmatic Encroachment

It is generally agreed that knowledge requires justification, and, in particular, that it requires a level of justification that is at least high enough to justify believing that p . We may distinguish, however, between *having the belief* that p and *occurently believing* that p . Accordingly, we may distinguish between the level of justification required to justify having the belief that p , and the level of justification required to justify occurently believing that p . Plausibly, knowing that p requires a level of justification that is at least high enough to justify occurently believing that p . Or, to put the same point in simpler (if less grammatical) terms, it is plausible that anyone who knows that p must be *justified to occurently believe that p* .

But if this is right, then the reasoning disposition account of belief can help explain the Knowledge Action Principle. To provide such an explanation, we will need the following additional assumption.

Justification Condition on Occurrent Attitudes: If having attitude A essentially involves being disposed to ϕ under circumstance C , then an agent S is justified to occurently have attitude A in C only if it is rationally permissible for S to ϕ in C .

¹⁵ The account of belief offered in Frankish [2009] may appear to be a counterexample to the thesis of this paragraph. For Frankish claims that believing that p involves a disposition to treat p as true, or, as he puts it, a ‘policy of premising that p .’ But he nonetheless maintains that outright beliefs are reducible to credences and preferences. For he claims that ‘one will count as having a flat out belief with content p if one is highly confident that one has embarked on a course of premising that p and attaches high desirability to adhering to it’ (p. 85). These two claims are compatible, on his view, because having such a combination of credences and preferences would *constitute* having a policy of premising that p . But this seems wrong: for one can be *confident* that one has a policy of reasoning in a certain way, and *strongly prefer* that one adheres to such a policy, without actually having such a policy. Moreover, Frankish’s account of belief has highly counterintuitive implications. It implies, for example, that one can rationally be certain that p is false while at the same time believing that p . For there are circumstances in which it would be rational to be certain that p is false, while at the same time *being confident* that one has embarked on a policy of premising that p and *strongly preferring* that one adheres to such policy.

This appears to be a plausible principle. It implies, for example, that if preferring X to Y essentially involves being disposed to choose X when faced with a choice between X and Y , then an agent is justified to *occurrently prefer* X to Y when faced with such a choice only if it is rationally permissible for the agent to *choose* X when faced with such a choice. Note that it is compatible with this principle that the agent may rationally continue to *dispositionally* prefer X to Y even when faced with a situation in which there is a choice between X and Y and it is not rationally permissible to choose X . The principle requires only that it is not rational to *occurrently* prefer X to Y in such a situation.

We can now derive the Knowledge Action Principle as follows (in each line, S , ϕ , p and C range over all agents, actions, propositions and choice situations, respectively).

- (1) [From the reasoning disposition account] If, in C , it is relevant whether p , then believing that p essentially involves being disposed to treat p as true in one's reasoning in C .
- (2) [From 1 and the Justification Condition on Occurrent Attitudes] In C , if it is relevant whether p , and S is justified to occurrently believe that p , then it is rationally permissible for S to treat p as true in her reasoning.
- (3) In C , if S knows that p , then S is justified to occurrently believe that p .
- (4) In C , if it is rationally permissible for S to treat p as true in her reasoning, then it is rationally permissible for S to act as if p .
- (5) [from 2, 3 and 4] In C , if it is relevant whether p , and S knows that p , then it is rationally permissible for S to act as if p .
- (6) In C , if it is rationally *impermissible* for S to act as if p , then what is unconditionally optimal must differ from what is optimal conditional on p , and so it is relevant whether p .
- (7) [From 5 and 6] In C , if it is rationally impermissible for S to act as if p , then S does not know that p .

Thus, from the reasoning disposition account of belief, together with some other plausible assumptions, we have derived the Knowledge Action Principle. And so this principle, and with it pragmatic encroachment on knowledge, can be explained using either of two accounts of

belief: the pragmatic credal reductionist account or the reasoning disposition account. Moreover, the defender of either of these accounts can provide such an explanation without appealing to any pragmatic encroachment on justified credence. The pragmatic credal reductionist needn't appeal to pragmatic encroachment on justified credence, for she maintains that pragmatic differences can lead to differences between what is known by people with identical justified credences, by affecting whether these justified credences amount to justified beliefs. And the proponent of the reasoning disposition account needn't appeal to pragmatic encroachment on justified credence, because she denies that belief is to be understood in terms of credence, and so she can deny that knowledge and justified belief must be understood in terms of justified credence.

2. Why Our Account is Preferable to Pragmatic Credal Reductionism

So far, we have seen that there are at least two alternative accounts of belief that could explain the Knowledge Action Principle, and that could thus explain pragmatic encroachment on knowledge, without committing to pragmatic encroachment on justified credence. The question therefore arises as to which of these two accounts of belief is preferable. We will now argue that there is strong reason to prefer the reasoning disposition account. We will do so by considering four claims about belief that we take to characterize central features of belief as this attitude is ordinarily conceived. In the case of each claim, we will argue that the same pattern can be found: the reasoning disposition account can be used to explain the truth of the claim in question, whereas the pragmatic credal reductionist account, far from explaining the truth of the claim in question, cannot plausibly be reconciled with it. Hence, insofar as these four claims about belief are plausible, they make a very strong case for favoring the reasoning disposition account of belief over the reductionist alternative.

But first, a brief digression. Since three of the four claims will concern the *rationality* of outright belief, before we can evaluate the two accounts of belief in relation to these claims, we'll first need some understanding of what these accounts imply concerning such rationality.

2.1. *Rationality Conditions on Outright Belief*

On the pragmatic credal reductionist picture, the conditions under which it is rational to believe a proposition, in a given circumstance, are fairly straightforward. Under circumstance *C*, it's rational for *S* to believe that *p* just in case it's rational for *S*'s credence in *p* to be high enough to constitute believing that *p* in *C*. And this in turn will be true just in

case it's rational for *S*'s credence in *p* to be high enough to rationalize acting as if *p* in every relevant choice situation.

On the reasoning disposition account, determining the rationality conditions for outright belief is less straightforward, but one thing is clear. On this account, belief is a kind of heuristic. And heuristics are justified pragmatically, in terms of the ends they serve. But on the reasoning disposition account, belief serves two competing ends. On the one hand, insofar as beliefs dispose one to treat what is believed as true in reasoning, belief serves the end of reasoning, which is to arrive at a good conclusion. Thus, one of the ends of belief is to allow the believer to arrive at good deliberative conclusions—which, in the case of practical reasoning, amounts to choosing courses of action with high expected utilities. On the other hand, belief serves the end of preventing the cognitive overload that would result from reasoning in an ideal Bayesian manner on the basis of one's credences. Thus, the second end of belief is to minimize the cognitive costs of reasoning. And these two ends compete: reasoning in an ideal Bayesian manner would be optimal in relation to the first end, but far from optimal in relation to the second. And the more we depart from ideal Bayesian reasoning in order to promote the second end, the more we risk sacrificing the first end. What is required, therefore, is a balance between these ends.

But if belief is a heuristic whose function is to balance these two ends, then we would expect the rationality of belief to be intelligible in terms of these ends. There are several ways in which this might work. Here's one.

Procedural Rationality Condition: A set of beliefs is rational only to the extent that it is licensed by rules or procedures that strike an optimal balance between minimizing expected cognitive costs and maximizing the expected value of the agent's deliberative conclusions.¹⁶

This principle connects the rationality of belief to the two ends indirectly, via rules or procedures, and so it can be thought of as the doxastic analogue of rule consequentialism in moral theory. And just as rule consequentialism is one among a plethora of ways in which the moral rightness of actions may be based on the moral values of outcomes,¹⁷ so the Procedural Rationality Condition is one among a plethora of ways in which the rationality of belief may be based on the two

¹⁶ We leave it open how these 'expected values' are to be understood. They might be defined in terms of the agent's actual or idealized credences, in which case they will vary as the agent's credences or evidence changes. Alternatively, they might be understood in terms of hypothetical priors, in which case they will not so vary.

¹⁷ See, in particular, Sinnott-Armstrong [2006] and Driver [forthcoming].

ends we have indicated. We expect that many of these ways of connecting the rationality of belief to these ends would serve our explanatory purposes equally well, and we do not mean to commit ourselves to any one of them in particular. But for the sake of concreteness, in illustrating the explanatory potential of the reasoning disposition account, we will assume the Procedural Rationality Condition.¹⁸

We now turn to the four claims. While our discussion of the first claim won't require appealing to any principles of rationality, our discussions of the remaining three will.

2.2. *Correctness*

The first claim about belief may be stated as follows.

Correctness: Believing that p when p is true constitutes being right about whether p , whereas believing that p when p is false constitutes being wrong about whether p .

This claim about belief seems undeniable, and so it would seem to be a criterion of adequacy for an account of belief that it be compatible with this claim. Fantl and McGrath [2010] concur, and they criticize another version of credal reductivism—a version according to which the level of credence required for believing a proposition does not vary from context to context—precisely on the ground that it conflicts with Correctness:

Consider a standard Lockean View under which belief is a matter of having a credence greater than some $d < 1$. Suppose d is .98. If you have a .99 credence for p , and p turns out to be false, it does not follow that you were wrong about whether p . If you were told 'Ha, so you were wrong about whether p , weren't you?' you could reasonably say in your defense: 'Look, I took no stand about whether p is true or false; I just assigned it a high probability; I assigned its negation a [positive] probability, too' (p. 141).

Here we agree entirely with Fantl and McGrath. The considerations they adduce do seem to show that Correctness conflicts with what they

¹⁸ One might object that *any* principle that ties the rationality of belief to the aforementioned ends must concern the practical rather than the epistemic rationality of belief. For the epistemic rationality of belief, one might argue, cannot depend on pragmatic factors. We do not take up this objection here, for here our principal opponents (namely pragmatic credal reductivists) are committed to pragmatic encroachment on both knowledge and the epistemic rationality of belief. For a detailed discussion of this issue, see Schroeder (forthcoming). And for a general discussion of the bearing of pragmatic factors on the rationality of attitudes, see Ross (2012).

call the Lockean View of belief. But these considerations seem equally to show that Correctness conflicts with PCR. For suppose that, given your practical circumstances, the minimum level of credence you'd need to have in p for it to be rational for you to act as if p in all there relevant choice situations is .98. And suppose your credence in p is .99, so that you count, on Fantl and McGrath's view, as believing that p . If p turns out to be false, and you were told 'Ha, so you were wrong about whether p , weren't you?' then once again it seems you could reasonably say in your defense: 'Look, I took no stand about whether p is true or false; I just assigned it a high probability (a probability that happened to suffice to rationalize acting as if p in the relevant choice situations); but I assigned its negation a positive probability, too.' This reply would seem to have just as much force as the original reply. Thus, the move from the Lockean view to PCR doesn't solve the serious problem with the former view that Fantl and McGrath identify.

Here's another way to put the problem. Whatever it is that constitutes, or makes it the case, that an agent is wrong about whether p when p is false, it can't be an attitude that involves, or commits the agent to, *acknowledging the possibility that p is false*. But having a credence in p of less than 1 involves, or commits one to, having a positive credence in $\sim p$, and so it involves, or commits one to, acknowledging the possibility that p is false. Hence, having a credence of less than one in a false proposition p can't constitute being wrong about whether p . But if having a credence of less than one in a false proposition p could constitute falsely believing that p , and if falsely believing that p always constitutes being wrong about whether p , then having a credence of less than one in a false proposition p could constitute being wrong about whether p . It follows that if having a credence of less than one in a false proposition p could constitute falsely believing that p , then it would not be the case that falsely believing that p always constitutes being wrong about whether p . Consequently, any view on which having less than full credence in a false proposition can constitute believing this proposition—including PCR—is incompatible with Consistency.

By contrast, the reasoning disposition account is not only consistent with Correctness, but can help to explain why it's true. For according to the reasoning disposition account, the belief that p is a mental state that disposes one to treat p as true in one's reasoning. But if p is true, then someone who is treating p as true in her reasoning is operating with a conception that corresponds to the actual situation with respect to p , and in this manner she is getting things right with respect to p ; whereas if p is false then she is operating with a conception that corresponds to a non-actual situation with respect to p , and in this manner she is getting things wrong with respect to p . Hence, the reasoning

disposition account implies that if p is true, then someone who believes that p is thereby right about p in the sense that she is disposed to reason in a way that involves getting things right with respect to p , whereas if p is false then someone who believes that p is thereby wrong about whether p in the sense that she is disposed to reason in a way that involves getting things wrong with respect to p .

2.3. Stability

While Correctness raises a general problem for just about *any* version of credal reductivism, the remaining three claims we will consider raise special problems for *pragmatic* credal reductivism. Here is one such claim.

Stability: A fully rational agent does not change her beliefs purely in virtue of an evidentially irrelevant change in her credences or preferences.

In relation to changes in one's belief in a proposition p , when we say that a given preference is evidentially irrelevant, we mean that one's having this preference provides no evidence for or against p . And when we say that one's credence in q is evidentially irrelevant, we mean both that q provides no evidence for or against p , and that one's credence in q provides no evidence for or against p . Understood in this way, Stability seems very plausible. If a fully rational agent's ice cream preferences provide no evidence for or against the cat's being on the mat, then the agent shouldn't gain or lose the belief that the cat is on the mat purely in virtue of changing her ice cream preferences. Similarly, if neither facts about the weather in Kalamazoo nor facts about an agent's credences concerning the weather in Kalamazoo provide any evidence for or against the cat's being on the mat, then the agent shouldn't gain or lose the belief that the cat is on the mat purely in virtue of changing her credences concerning the weather in Kalamazoo. However, PCR conflicts with Stability. For according to PCR, to believe that p is to have sufficient credence in p to rationalize acting as if p in the relevant choice situations. But how much credence in p this will require will generally depend on evidentially irrelevant credences and preferences.

To see how this can depend on evidentially irrelevant *credences*, consider High and Low, where it appears that in Low, but not in High, Sarah knows that r (that is, that the sandwich on the right is an almond butter sandwich). The pragmatic credal reductivist explains this difference by maintaining that, while the level of credence that it is rational for Sarah to have in r is the same in High and in Low, it is true in Low, but not in High, that having this level of credence suffices

for believing that r . Suppose, however, that in Low, while Sarah's credence in r remains constant, she comes to have high credence in the proposition that her guest has a peanut allergy. In this case, her situation would become relevantly similar to her situation in High. And so it seems the pragmatic credal reductivist is committed to saying that Sarah's credence in r would cease to be sufficient to constitute believing that r . Conversely, suppose that in High, Sarah were to become confident that her guest has *no* peanut allergy. In this case, her situation would become relevantly similar to her original situation in Low. And so it seems the pragmatic credal reductivist is committed to saying that Sarah's credence in r would come to be sufficient to constitute believing that r . Thus, it seems the pragmatic credal reductivist is committed to saying that Sarah could gain or lose the belief that r purely in virtue of an evidentially irrelevant change in her credences.

To see how, according to PCR, a rational agent's beliefs can change purely in virtue of a change in her evidentially irrelevant *preferences*, consider a case in which Sarah is dining alone, and she must choose between the three sandwiches in the fridge. Suppose Sarah has no allergies, and she has a preference, of an ordinary magnitude, for the almond butter sandwich over the tuna sandwich, and she has a preference of an equal magnitude for the tuna sandwich over the peanut butter sandwich. This would seem to be an ordinary, low-stakes context where Sarah should count as knowing that r , and so the pragmatic credal reductivist should maintain that, in this case, Sarah's justified credence in r is sufficient for her to count as believing that r . But now suppose that while Sarah's credence in r remains constant, the degree to which she prefers the almond butter sandwich to the tuna sandwich progressively diminishes. There will come a point, as the magnitude of this preference approaches zero, when Sarah's credence in r will cease to suffice to rationalize her acting as if r by taking the sandwich on the right. Hence, PCR will imply that there will come a point where Sarah's credence in r ceases to suffice for her to believe that r . And in the reverse case, where the magnitude of Sarah's preference progressively increases from zero, PCR will imply that there will come a point where Sarah's credence in r comes to suffice for her to believe that r .

Note that PCR doesn't merely imply that changes in a rational agent's beliefs can be *caused* by changes in her evidentially irrelevant credences or preferences. Rather, it implies that the latter changes can *constitute* the former changes. Thus, it implies that, in the case just described, as the strength of Sarah's preference increases or decreases, *by that very fact* she would gain or lose the belief that r . And this implication is very hard to accept.

The instability predicted by PCR is particularly implausible in relation to propositions for which we have overwhelming evidence, but in which we aren't absolutely certain. Suppose Stella is extremely confident that steel is stronger than Styrofoam, but she's not so confident that she'd bet her life on this proposition for the prospect of winning a penny. PCR implies, implausibly, that if Stella were offered such a bet, she'd cease to believe that steel is stronger than Styrofoam, since her credence would cease to rationalize acting as if this proposition is true.

While PCR is inconsistent with Stability, the reasoning disposition account supports a conception of belief that would explain why Stability is true. For, as we argued in section 1.4, the reasoning disposition account assigns to belief a role that credences and preferences simply cannot play. And so it supports a conception of belief as an autonomous mental state distinct from credences and preferences. And from this conception it follows that changes in evidentially irrelevant preferences and credences can't constitute changes in one's beliefs.

Moreover, the reasoning disposition account can explain more than this. Consider cases where an agent moves from a low- to a high-stakes context, or vice versa, in virtue of changes in evidentially irrelevant credences or preferences (where, relative to a proposition p , a *high-stakes* context is one where the costs of acting as if p if p is false are high relative to the costs of failing to act as if p if p is true, and a *low-stakes* context is one where the reverse obtains). In such cases, the reasoning disposition account supports not only the claim that these changes don't *constitute* changes in the agent's beliefs, but also the claim that these changes don't *rationalize* changes in the agent's beliefs.

For assuming the Procedural Rationality Condition from section 2.1, it will be rational for an agent to change her beliefs as she moves between high- and low-stakes contexts only if such changing beliefs are licensed by procedures that strike an optimal balance between minimizing expected cognitive costs and maximizing the expected value of the agent's deliberative conclusions. (Let's call such procedures *optimific*.) But if the reasoning disposition account is true, then we should not expect the optimific procedures to license such changes in one's beliefs, as such changes would require an unnecessary expenditure of cognitive resources. For recall that according to the reasoning disposition account, the belief that p involves a *defeasible* disposition to treat p as true in reasoning, a disposition that is overridden when the costs of mistakenly acting as if p are salient. Hence, on this account, there will be no need to drop the belief that p when one enters a high-stakes context. For one can instead retain this belief while overriding the disposition to treat p as true, by attending to the costs of mistakenly acting as if p . And, having retained the belief that p , there will be no need to

reacquire this belief when one reenters a low-stakes context. We should expect, therefore, that the optimific rules would prescribe stability in one's beliefs as one moves between high- and low- stakes contexts.

2.4. *Sufficient Evidence*

Our third claim about belief can be stated as follows.

Sufficient Evidence: It is rational to believe a proposition p only if one's evidence significantly favors p over its negation.

This claim is difficult to deny, and yet it conflicts with PCR. The conflict is particularly striking in relation to what we may call *practically irrelevant propositions*, that is, propositions such that one's credence in them makes no difference to how it would be rational for one to act. Here's one example:

m : The width of Salvador Dali's mustache when he painted *The Persistence of Memory* was twice the length of Cleopatra's nose when she met Mark Antony.

For most of us, most of the time, our credence in m makes no difference to how we should act. And so, regardless of what our credence in m may be, it will be rational for us to act as if m , in the sense of doing what would be rationally optimal conditional on m . Thus, PCR seems to imply that, even if our credence in m is zero, we will count as believing that m . And this result is clearly unacceptable.

Perhaps the solution to this problem lies in adopting an appropriate conception of the choice situations that count as relevant in evaluating what an agent believes. If the relevant choice situations are defined in a way that ties them to the agent's actual circumstances—if they are defined, for example, as the choice situations that the agent actually faces or that she could reasonably expect to face—then it will be hard to avoid the conclusion that there are some propositions such that our credence in them makes no difference in any relevant choice situation. Hence it will be hard to avoid the conclusion that there are propositions for which *any* level of credence suffices for belief. The pragmatic credal reductivist can avoid this conclusion, however, if she defines the relevant choice situations in a manner that divorces them from the agent's actual circumstances, e.g., by identifying them with the class of choice situations the agent could *conceivably* face. For, given a sufficiently liberal conception of a relevant choice situation, low credence in a proposition will never suffice to rationalize acting as if it's true in every relevant choice situation. Such a move, however, would

undermine the pragmatic credal reductionist explanation of pragmatic encroachment. For recall that, according to this explanation, the reason that Sarah knows that r in Low but not in High is that it is true in Low, but not in High, that Sarah's credence in r suffices to rationalize acting as if r in the relevant choice situations. But this explanation only works if the choice situations that count as relevant are closely tied to Sarah's actual circumstances.

Is there any way for the pragmatic credal reductionist to solve the problem raised by practically irrelevant propositions? Among the proponents of this view, the only one to recognize and address this problem is Brian Weatherson. To understand his proposed solution, we'll need to consider it in some detail. On Weatherson's view, S believes that p just in case, for all A and B in the appropriate class of options, and all q in the appropriate class of propositions, S prefers A to B conditional on q iff S prefers A to B conditional on $(p \& q)$. The key to solving the problem of practically irrelevant propositions is to give the right account of appropriate classes of propositions and options. In evaluating what an agent S believes, the appropriate class of propositions, according to Weatherson, are the 'active' propositions, defined as propositions that are either practically relevant to S , or salient to S , or conjunctions of propositions that are practically relevant or salient to S . And when evaluating whether an agent S believes a proposition p , the appropriate class of options, O_p , is defined as follows. Where O is the class of options that are available and salient to S , O_p is the union of O and the following options: the option of believing that p , the option of not believing that p , as well as the options of believing and of not believing that q , for every proposition q that is relevant or salient to S . Lastly, Weatherson stipulates that for any proposition p , S prefers believing that p to not believing that p just in case S 's credence in p exceeds .5 and conditionalizing on p would not affect S 's conditional preferences (conditional on any active proposition) over the options in O . And this stipulation holds both unconditionally, and conditional on q , for any q . (Weatherson [2005], pp. 421–424).

As Weatherson shows, his view implies that a *necessary* condition for believing a proposition p is that one's credence in p exceed .5. For if one's credence in p does not exceed .5, then conditionalizing on p will affect one's conditional preferences over the options of believing that p and not believing that p . However, Weatherson's theory still has the implausible implication that, very often, a *sufficient* condition for S 's believing that p is that S have credence above .5 in p . This will be true, in particular, if p is *strongly* practically irrelevant for S , in the sense that the following two conditions obtain:

- (i) Apart from the options of believing that p and of not believing that p , conditionalizing on p would not affect S 's conditional preferences (conditional on any active proposition) over the options in O_p .
- (ii) p is probabilistically independent of every active proposition. That is, for every active proposition q , S 's credence in p conditional on q is the same as S 's unconditional credence in p .

For it follows from Weatherson's account of belief, together with condition (i) above, that if p is strongly practically irrelevant for S , then S believes that p so long as conditionalizing on p wouldn't affect S 's conditional preferences over the options of believing that p and of not believing that p . And it follows from Weatherson's account of conditional preferences over beliefs, together with condition (ii), that if p is strongly practically irrelevant for S , then conditionalizing on p wouldn't affect S 's conditional preferences over these two options so long as S 's credence in p exceeds .5. And so it follows that if p is strongly practically irrelevant for S , then S will believe that p so long as S 's credence in p exceeds .5.

But this conclusion has very implausible implications. Suppose, for example, that Flip knows that the objective chance of an American penny coming up heads is .5005 (Yates et al. [1999], p. 314). Suppose, therefore, that Flip has credence .5005 in the following, strongly practically irrelevant proposition (call it h) that the next time an American penny is tossed by a Lithuanian xylophonist, it will come up heads. Surely Flip could withhold judgment concerning h , neither believing that h nor believing that $\sim h$. But Weatherson's theory implies that Flip must believe that h . Consequently, Weatherson's theory conflicts with Sufficient Evidence. For it implies that rationality permits (and indeed requires) Flip to believe that h , even though his evidence doesn't significantly favor h over its negation.

We've seen that it's hard to square PCR with Sufficient Evidence. By contrast, the reasoning disposition account of belief, together with the Procedural Rationality Condition, can explain why Sufficient Evidence is true. To see how, let's say that p is a *low probability proposition* just in case the evidence doesn't significantly favor p over $\sim p$, so that the probability of p on the evidence doesn't significantly exceed .5. And let's say that p is a *problematic proposition* just in case the *benefits* of believing that p (in relation to the end of minimizing expected cognitive costs) are outweighed by the *costs* of believing that p (in relation to the end of maximizing the expected value of one's deliberative conclusions). Now if the reasoning disposition account is true, then we

should expect that, by and large, low probability propositions will be problematic, since, for most low probability propositions, being disposed to treat them as true in reasoning would do more harm than good. There may be the occasional low probability proposition that is unproblematic, but these will be rare, and the cognitive costs of identifying them would be high. And so any procedures that permitted belief in low probability propositions would either result in one's believing numerous problematic propositions (with the result that the procedures in question would fare poorly with respect to the end of maximizing the expected value of one's deliberative conclusions) or else these procedures would have to involve the screening out of problematic proposition (with the result that they would fare poorly with respect to the end of minimizing expected cognitive costs). Consequently, if the reasoning disposition account of belief is true, then we should expect the optimistic rules or procedures to prohibit belief in low probability proposition. And so, if the reasoning disposition account and the Procedural Rationality Condition are both true, then we should expect belief in low probability propositions to be rationally impermissible, and so we should expect Sufficient Evidence to be true.

2.5. Consistency

The last claim about belief we will consider is the following.

Consistency: There is rational pressure to avoid logically inconsistent beliefs. In particular, where A is a small set of related propositions that are jointly inconsistent, it is rationally impermissible to simultaneously believe every proposition in A .

The conflict between Consistency and PCR can be seen by considering practically irrelevant propositions, such as m . As we saw in the last section, any version of PCR on which the relevant choice situations are limited to the agent's actual and foreseeable choice situations will imply *any* level of credence in m suffices for believing that m , since our credence in m makes no difference to how we should act in such situations. Note, however, that if our credence in m makes no difference to how it would be rational for us to act in our actual or foreseeable choice situations, the same will be true of our credence in m 's negation, $\sim m$. And so such version of PCR will likewise imply that any credence in $\sim m$ will suffice for believing $\sim m$. They will imply, therefore, that, regardless of what one's credences in m and in $\sim m$ may be, one will count as believing both propositions. And so they will imply that a rational agent will believe both these propositions, contrary to Consistency.

Weatherson's theory avoids the implication that one can rationally believe both a proposition and its negation. For, as we saw in the last section, his theory implies that one believes a proposition only if one's credence in it exceeds .5. It therefore implies that a rational agent will never believe both a proposition and its negation, since she will never have more than .5 credence in each. However, Weatherson's view still conflicts with Consistency. To see how, consider the following three propositions:

- x: The number of hairs in the left half of Dali's mustache is not evenly divisible by three.
- y: The number of hairs in the right half of Dali's mustache is not evenly divisible by three.
- z: The number of hairs in one or other of the halves of Dali's mustache is evenly divisible by three.

Suppose that Gala rationally has $\frac{2}{3}$ credence in each of these propositions, and that each one is strongly practically irrelevant for her. As we argued in the last section, Weatherson's theory implies that one will count as believing a strongly practically irrelevant proposition so long as one's credence in it exceeds .5. It thus implies that, in virtue of having these rational credences, Gala believes all three propositions. And so Weatherson's theory implies that one can rationally believe each proposition in a small set of closely related, jointly inconsistent propositions, contrary to Consistency.

Once again, while PCR is hard to reconcile with Consistency, the reasoning disposition account can help to explain why Consistency is true. For according to the reasoning disposition account, someone who believes that p is disposed to treat p as true in any reasoning in which it is relevant whether p . And so if this account of belief is true, there would be an inherent danger in believing jointly inconsistent propositions. For if one is ever in a situation where all these propositions are relevant to one's reasoning, then one will be at risk of treating each of them as true in one's reasoning. And this would prevent one from reasoning coherently, and hence from arriving at a good deliberative conclusion. This risk would be particularly great if one had inconsistent beliefs among a small set of closely related propositions, for such propositions are particularly likely to be jointly relevant in one's reasoning.

There are two ways in which cognitive rules or procedures could prevent such incoherent reasoning: they could either prevent the formation of inconsistent beliefs in the first place (or at least prevent

the formation of inconsistent beliefs among small sets of closely related propositions), or else they could allow such beliefs to be formed but prevent them from being jointly operative. But preventing the joint *formation* of such inconsistent beliefs would be much less cognitively costly than preventing their joint *operation*. For preventing the *formation* of jointly inconsistent beliefs would require only that a check for consistency be performed prior to forming any new belief. However, preventing the *operation* of jointly inconsistent beliefs would require that checks for consistency be performed far more frequently, prior to the employment of any given belief. Consequently, if the reasoning disposition account of belief is true, then we should expect the optimific rules to prohibit the formation of jointly inconsistent beliefs, at least in relation to small sets of closely related propositions. Hence, by the Procedural Rationality Principle, we should expect such inconsistent beliefs to be irrational. And so, if the reasoning disposition account and the Procedural Rationality Principle are both true, then we should expect Consistency to be true as well.¹⁹

Historical Postscript

It is said, rightly, that entities should not be multiplied beyond necessity. But it is likewise true that entities should not be reduced to the point of inadequacy. In theorizing about the mind, philosophers have tended to focus on the first of these principles. One result of this tendency has been a longstanding tradition of attempting to understand human action in terms of only two types of attitude, a single type of belief-like or cognitive attitude, and a single type of desire-like or conative attitude. This tradition can be traced back at least as far as Leibniz,²⁰ and it finds its most sophisticated contemporary expression in decision theory, where rational action is understood purely in terms of credences and preferences.

¹⁹ This explanation of Consistency, like the explanations we offer of Stability and of Sufficient Evidence, relies on claims about contingent features of human beings. There may be believers who are very different from us about whom these claims are false. Hence, our arguments are not meant to show that Stability, Sufficient Evidence and Consistency are true unrestrictedly, of every possible believer. And this is as it should be, as it is far from obvious that these principles are true unrestrictedly, of all possible believers. Moreover, even in relation to beings very different from us, the Procedural Rationality Condition will never imply that any of these three claims is *false*. That is, it will never imply that it would be *rational* for such beings to have unstable, improbable, or inconsistent beliefs. For the Procedural Rationality Condition provides only necessary, not sufficient, conditions for rationality, and so it never implies that any beliefs are rational.

²⁰ See Leibniz's *Monadology*, sec. 14ff. and his *Principles of Nature and Grace*, sec. 2, both of which are printed in Leibniz [1898].

One thing that has emerged in the practical reason literature, however, is that for cognitively limited agents like us, credences and preferences aren't enough; we need a separate attitude of intention. This has been argued most forcefully by Michael Bratman.²¹ If we had infinite cognitive resources, Bratman argues, we'd have no need for prior intentions. At each moment of action we could consider every possible alternative, instantly calculate which alternatives have the highest expected utility relative to our credences and preferences, and act accordingly. But such exhaustive and instantaneous consideration of alternatives is not feasible for cognitively limited agents like us. And so we need an attitude of *prior intention* or of *settling on a course of action in advance*, so as to guide our future actions and limit our future deliberations to options consistent with what we have settled on. And Bratman argues that no combination of preferences and credences, nor any similar combination of desire-like and belief-like attitudes, can play this role. In part because of Bratman's arguments, and in part because of the disappointing track record of attempts to reduce intentions to desire-like and belief-like attitudes,²² the project of carrying out such a reduction has been largely abandoned, and the autonomy of intention is now widely acknowledged.²³

In this paper we have been arguing that an analogous situation exists with respect to outright belief. If we had infinite cognitive resources, then we'd have no need for an attitude of outright belief by which to guide our actions, for we could reason in an ideal Bayesian manner on the basis of our credences and preferences alone. But such reasoning isn't feasible for cognitively limited agents like us, and so we need an attitude of *outright belief* or of *settling on the truth of propositions*, so as to limit what we consider in our reasoning to possibilities consistent with what we have settled on. And we have argued that no combination of preferences and credences can play this role. Thus, the reasons for acknowledging the autonomy of outright belief are closely analogous to the reasons for acknowledging the autonomy of intention. What these considerations suggest, in both cases, is that to reduce the bases of rational action to credences and preferences would be to reduce entities to the point of inadequacy.

²¹ See especially Bratman [1987]. See also Bratman [1985].

²² For examples of reductivist accounts of intention, see Audi [1973]; Davidson [1980], chaps 1, 2, and 5; Brand [1984]; and Davis [1984]. For anti-reductivist arguments, in addition to those given by Bratman, see Donagan [1987] and Mele [1988].

²³ There are, of course, some exceptions. For a recent defense of a view on which intention is a kind of belief, see Setiya [2007]. And for a critical discussion, see Ross [2009].

References

- Audi, Robert [1973]. 'Intending.' *Journal of Philosophy* 70(13): 387–403.
- Brand, Myles [1984]. *Intending and Acting: Toward a Naturalized Action Theory*. Cambridge, MA: The MIT Press.
- Bratman, Michael [1985]. 'Davidson's Theory of Intention.' *Actions and Events: Perspectives on the Philosophy of Donald Davidson*. Edited by Ernest Lepore and Brian McGlauchlin. Oxford: Basil Blackwell, 14–28.
- [1987]. *Intention, Plans, and Practical Reason*. Chicago: University of Chicago Press.
- Cohen, Stewart [1999]. 'Contextualism, Skepticism, and the Structure of Reasons.' *Philosophical Perspectives* 13: *Epistemology*: 57–89.
- Davidson, Donald [1980]. *Essays on Actions and Events*. Oxford: Oxford University Press.
- Davis, Wayne [1984]. 'A Causal Theory of Intending.' *American Philosophical Quarterly* 21(1): 43–54.
- DeRose, Keith [1992]. 'Contextualism and Knowledge Attributions.' *Philosophy and Phenomenological Research* 52(4): 513–529.
- Donagan, Alan [1987]. *Choice: The Essential Element in Human Action*. London: Routledge & Kegan Paul.
- Driver, Julia [forthcoming]. *Consequentialism*. Forthcoming with Routledge.
- Fantl, Jeremy and Matthew McGrath [2002]. 'Evidence, Pragmatics, and Justification.' *Philosophical Review* 111(1): 67–94.
- Fantl, Jeremy and Matthew McGrath [2010]. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Frankish, Keith [2004]. *Mind and Supermind*. Cambridge: Cambridge University Press.
- [2009]. 'Partial Belief and Flat Out Belief.' *Degrees of Belief*. Edited by Franz Huber and Christoph Schmidt-Petri. Springer (Synthese Library), pp. 75–93.
- Ganson, Dorit [2008]. 'Evidentialism and Pragmatic Constraints on Outright Belief.' *Philosophical Studies* 139 (3): 441–458.
- Hawthorne, John [2004]. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John and Jason Stanley [2008]. 'Knowledge and Action.' *Journal of Philosophy* 105(10): 571–90.
- Joyce, James [1999]. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Leibniz, Gottfried [1898]. *The Monadology and Other Philosophical Writings*. Translated by Robert Latta. Oxford: Oxford University Press.

- Mele, Alfred [1988]. 'Against a Belief/Desire Analysis of Intention' *Philosophia* 18 (2-3): 239–243.
- Resnik, Michael [1987]. *Choices: An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press.
- Ross, Jacob [2006]. *Acceptance and Practical Reason*. Doctoral Dissertation, Rutgers University.
- [2009]. 'How to be a Cognitivist about Practical Reason.' *Oxford Studies in Metaethics* 4: 243–281.
- [2011]. Comments on 'Ross and Schroeder on Belief.' *Thoughts, Arguments, and Rants*. <http://tar.weatherson.org/2011/03/31/ross-and-schroeder-on-belief> (accessed May 27, 2011).
- [2012]. 'Rationality, Normativity, and Commitment.' Forthcoming in Russ Shafer-Landau (ed.) *Oxford Studies in Metaethics*, vol. 7. New York: Oxford University Press.
- Savage, Leonard [1972]. *The Foundations of Statistics*. Second edition. New York: Dover Publishing.
- Schroeder, Mark [forthcoming]. 'Stakes, Withholding, and Pragmatic Encroachment on Knowledge.' Forthcoming in *Philosophical Studies*.
- Setiya, Kieran [2007]. *Reasons without Rationalism*. Princeton: Princeton University Press.
- Sinnott-Armstrong, Walter [2006]. 'Consequentialism.' *The Stanford Encyclopedia of Philosophy*. Edited by Edward Zalta. <<http://plato.stanford.edu/entries/consequentialism>>
- Stanley, Jason [2005]. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Weatherson, Brian [2005]. 'Can We Do Without Pragmatic Encroachment?' *Philosophical Perspectives* 19: 417–443.
- [2011]. 'Ross and Schroeder on Belief.' *Thoughts, Arguments, and Rants*. <http://tar.weatherson.org/2011/03/31/ross-and-schroeder-on-belief> (accessed May 27, 2011).
- [Unpublished]. 'Knowledge, Bets, and Interests.' Unpublished manuscript.
- Weisberg, Jonathan [Unpublished]. 'Belief: Full and Partial.' Unpublished manuscript.
- Yates, Daniel, David Moore and George McCabe [1999]. *The Practice of Statistics*. New York: W. H. Freeman and Company.