Rational Faith and Justified Belief

Lara Buchak, UC Berkeley[1]

1. Introduction

In "Can It be Rational to Have Faith?", I argued that to have faith in some proposition consists, roughly speaking, in stopping one's search for evidence and committing to act on that proposition without further evidence. In that paper, I also outlined when and why stopping the search for evidence and acting is rationally required. Because the framework of that paper was that of formal decision theory, it primarily considered the relationship between faith and *degrees* of belief, rather than between faith and belief full stop (hereafter, "belief"). The purpose of this paper is to explore the relationship between rational faith and justified belief.

Before rehearsing my account of faith, let me briefly say something about the overall epistemological picture that I am working with, which rests on two assumptions. The first is that beliefs come in degrees. This idea can be motivated by noticing that among the propositions I believe, the strength of my belief is not uniform. For example, while I believe that the sun will rise tomorrow and I believe that it will be a cloudless day, I believe the former much more strongly. I would be willing to bet more on the former: I would be willing, for example, to pay 99 cents for a bet that pays me a dollar if the sun rises tomorrow, but I would not pay 99 cents for a bet that pays me a dollar if it is a cloudless day. This distinction holds not only among propositions I believe, but also among propositions I fail to believe: I can rank these according to how likely I think they are to be true. While I don't believe that the Giants will win the World Series and I don't believe that the Cubs will win the World Series, I would be willing to pay 10 cents for a bet that pays a dollar if the Giants win, but I would be unwilling to accept such odds on a bet that pays off if the Cubs win. According to the epistemological picture here, my degree of belief that the sun will rise tomorrow is higher than my degree of belief that today will be cloudless, and my degree of belief that the Giants will win is higher than my degree of belief that the Cubs will win. On this picture, rational degrees of belief (or *credences*) are best thought of as subjective probabilities: subjective in the sense that they are the individual believer's response to her own evidence, and probabilities in the sense that they obey the axioms of the probability calculus. Rational degrees of belief are also *updated* in response to new evidence: so, for example, on my current evidence I might have degree of belief 0.1 that the Giants will win the World Series, but if I learn that their star shortstop is injured, my degree of belief drops to 0.05. We can represent this by writing p(Giants win) = 0.1 and p(Giants win | star shortstop is injured) = 0.05, where p(X | Y) is read as "the probability of *X* conditional on *Y*" or "the probability of *X*

[1] Penultimate draft. Forthcoming in Religious Faith and Intellectual Virtue, eds. Tim O'Connor and Laura Goins (Oxford University Press): please cite that version.

given *Y*. Finally, we can think of an individual's degree of belief in a proposition as an estimate of the proposition's truth-value, given her evidence.[2]

The second assumption is that rational believers ought to proportion their degrees of belief to the evidence. This thesis is known as *evidentialism*. Evidentialism rules out taking into account non-truth-conducive reasons in deciding what degrees of belief to hold. For example, one cannot adopt a high degree of belief in a proposition simply because one wants it to be true. Nor can one adopt a high degree of belief in a proposition for moral reasons, for example, because it is a proposition about a friend, and one has a moral duty to think well of a friend. I won't take a stand on whether a given body of evidence always recommends a unique set of credences or whether, on the contrary, two different individuals could have the same evidence and each rationally adopt different credences – but if this is possible, the difference must be explained by something epistemic rather than their non-epistemic values. Furthermore, I am only assuming evidentialism about *degrees of* belief: as we will see, one can be an evidentialist about degrees of belief while thinking that "on-off" belief ought to be sensitive to non-truth-conducive factors.

Epistemic rationality concerns which credences one ought to have and what one ought to believe. *Instrumental rationality* concerns what one ought to do. We saw above that there is a link between credence and betting behavior: of two propositions, a rational individual would rather take a bet on the proposition to which she assigns a higher credence. More generally, credences figure into a precise theory of rational action. A maxim that guides action is that one ought to take the means to one's ends: one ought to choose the act that brings about the outcome one values most. But this maxim cannot always be followed as stated. In the typical case, one has many competing ends which one values to different degrees; furthermore, one is typically not certain of what the state of the world is, but instead assigns credence to several possible states. For example, consider the choice about whether to bring one's umbrella to work: one values staying dry while carrying an umbrella more than getting wet, but one also values staying dry while not carrying an umbrella more than staying dry while carrying one; and one assigns some credence to the hypothesis that it is raining and some credence to the hypothesis that it is not raining.

Decision theory makes precise the components of this decision, and how they interact to produce a recommendation about what one should do. An act is thought of as a gamble whose payoffs depend on the state of the world. For example, not carrying an umbrella is the gamble that results in getting wet if it rains and staying dry while not carrying an umbrella if it doesn't: and if one thinks there is a 70% chance of rain, then this is the lottery {70% chance of getting wet, 30% chance of staying dry while not carrying

[2] See Joyce (2005). It is controversial how exactly to give content to what degree of belief represents, but the differences between the various proposals won't matter here.

an umbrella}. Also introduced is a *utility function*, a function which measures how much one values particular outcomes. The standard view is that an instrumentally rational individual ought to choose the act with the highest average utility value (the highest *expected utility*), given the probabilities she assigns to the various possible states. We write u(O) to stand for the utility value of some outcome *O*, and p(X) to stand for the probability of some possible state of the world *X*. If u(getting wet) = -3, u(staying dry while not carrying an umbrella) = 3, u(staying dry and carrying an umbrella) = 1, p(rain) = 0.3, and p(not-rain) = 0.7, then EU(don't bring umbrella) = (0.3)(-3) + (0.7)(3) = 1.2 and EU(bring umbrella) = (0.7)(1) + (0.3)(1) = 1; therefore, one ought not to bring one's umbrella.

On this picture, one's values and beliefs are subjective, and from them we can arrive at a recommendation about what to do. I hold – although this is controversial – that we ought additionally to allow an individual to determine for herself how to take risk into account. For example, some people care proportionally more about what happens in the worst-case scenario than what happens in the best-case scenario. Although the chance of rain is 0.3, when contemplating an act where "rain" leads to the worst outcome, the possibility of rain may weigh more heavily in deliberation. To determine which act is instrumentally rational for these *risk-avoidant* individuals, we can still calculate values according to a mathematical formula, but one which weights the minimum more heavily than its "probability share" of the state space.[3] For example, one might shift the "decision weights" – the weights the various possibilities get for decision-making purposes – so that the risk-weighted expected utility (REU) of not bringing an umbrella is (0.6)(-3) + (0.4)(3) = -0.6, and as a result forgoing one's umbrella is not recommended. (A similar point holds of "risk-inclined" individuals, who weight the maximum more heavily than the minimum.) Since this thesis is still controversial, I will make sure to say how things go both on standard decision theory (expected utility theory) and on the alternative I favor (risk-weighted expected utility theory).

This is the basic framework of "formal" epistemology: credences are the epistemological entity; the norm of epistemic rationality is to have credences that obey the probability calculus and (according to evidentialism) to proportion them to the evidence; and the norm of instrumental rationality is to maximize expected utility (or, on my view, to maximize risk-weighted expected utility). You'll notice that I haven't yet said anything about the relationship between credence and belief. That is because there is currently no agreed upon view about how the two frameworks fit together. The bulk of this paper will consider what can be said about faith and belief according to several plausible but competing answers to this question. First, however, let me briefly outline the account of faith I offered in Buchak (2012), and how faith fits into the formal epistemological picture.

[3] For details, see Buchak (2013).

2. The Nature of Faith

An account of the nature of faith should satisfy several criteria. First, it should capture what we take to be paradigm cases of faith, both intuitively and within the context of interpersonal relationships and religious practices. I assume throughout that religious faith is a special case of a general, unified attitude that encompasses "secular" cases of faith as well, such as faith in a friend. Thus, we are interested in a minimal core notion of faith that is consistent with distinct, thicker notions of faith such as Christian faith, faith in one's spouse, and so forth. Second, an account of the nature of faith should be able to distinguish between what we take to be good cases of faith and what we take to be bad cases of faith. If we don't agree about all cases, it should capture the cases we do agree about and yield a verdict on those we don't. Finally, those who endorse faith as a virtue think of it as a central component within some sphere of activity, for example, religious practice or interpersonal relationships. Thus, the final desideratum for an account of faith is that it explains why faith can be a virtue (intellectual or otherwise), and what the attitude of faith can add to human life. To set these criteria out isn't to prejudge the question of whether they can be satisfied (of whether there really is a notion of faith that is common to both religious and mundane contexts and according to which faith serves a positive role) but rather to begin a search for whether there is an account of faith that can meet them. Additionally, there may be no single sense of faith that explains all usages of the term: what I am after is the concept in the neighborhood that is normative – the concept according to which it is true of some people that they ought to have faith that a friend will come through for them or faith that God exists – *if* there is such a concept.

Even given these caveats, there are at least two important senses of the term "faith": *propositional faith* (faith that *X*, where *X* is some proposition) and *interpersonal faith* (faith in *I*, where *I* is some individual). My account is an account of propositional faith. My hope is that the correct account of interpersonal faith will ultimately rest on an account of propositional faith: for example, to have faith in a person is to have faith that some facts about her obtain. But even if these two senses of faith are largely independent, I take it there is still an important question about what propositional faith consists in. (The only substantive thesis I am ruling out is that an account of propositional faith rests on an account of interpersonal faith.) An account of propositional faith has two parts. The first delineates the set of propositions that are potential candidates for faith, and the second what it is to have faith in one of these propositions. While all propositions are potentially the objects of credence and of belief, not all propositions are even candidates for faith. Thus, I introduce three criteria that a proposition must meet in order to be a potential object of faith for a particular individual.

First, in order for a proposition to be a potential object of faith, the individual must care whether or not the proposition is true. Faith that *X* is incompatible with indifference about whether *X*.

Second, the individual must have a positive attitude towards the truth of the proposition. This can be seen by noting that while I can be said to have or lack faith that you will quit smoking, I can't appropriately be said to have or lack faith that you will continue smoking. The exact attitude one must have towards the proposition needs some spelling out, though. In my earlier account I said that the sense in which one must have a positive attitude is that one is basing some act on the proposition. However, this threatens to let in too many propositions as potential objects of faith. For example, it implies that my betting on your continuing to smoke is enough to make "you continue to smoke" an appropriate object of faith. One might instead claim that in order for *X* to be an appropriate object of faith, one must prefer that *X*, aside from any act one takes. However, this threatens to rule out too many propositions as potential objects of faith. For example, consider an individual whose friend brings her news that the individual's child has been kidnapped and that the individual must pay a ransom to rescue him: it is felicitous to say that the individual pays in part because she has faith that her friend is telling the truth, but she of course prefers that the friend be lying.[4] I don't think this example reveals that a positive attitude is not a necessary condition of faith – there is something that one has a positive attitude towards in this example, namely the friend – but that we need a more nuanced account of what sort of attitude is required. I leave this aside for future work.

Third, in order for a proposition *X* to be an appropriate object of faith for a particular individual, she must not take her evidence on its own to support her being certain that *X*: her evidence must leave open the possibility that *not-X*.[5] For example, while it is felicitous to say, before you know the results of a friend's exam, that you have faith that your friend passed the exam, it is infelicitous to say this once she shows you her passing grade. There are certain kinds of propositions for which evidence cannot generally produce certainty, because there is yet no fact of the matter: for those who hold that free actions must not be determined in advance, an example of such a proposition is a proposition concerning the future free act of another individual. Thus, these kinds of propositions will often be candidates for faith.

Now that we have delineated the set propositions that are candidates for faith, we can ask what having faith involves. A key component is that faith that *X* involves a willingness to act on *X* in situations in which doing so constitutes taking a risk. When we have faith that a particular individual will act in a certain way – keep our secret, pick us up from the airport, do what is in our best interests – we take a risk that the individual will let us down. We are vulnerable to the individual not acting as we have faith that she will act.

However, not every case of risk-taking will count as an act of faith. Faith involves a willingness to commit to acting on the proposition one has faith in without first looking for further evidence for or

[4] The general form this example takes is due to Alex Pruss.
[5] "*Not-X*" is hereafter represented as $\bar{X}$.

against that proposition. An individual with faith in her friend's ability to keep a secret must be willing to confide in her friend without first verifying with a third party that the friend isn't a gossip. A man who has faith that his wife is constant must commit to his marriage without first hiring a private detective to observe how his wife behaves when he is not there. An individual who has faith that a particular bridge will hold his weight doesn't test the bridge before stepping onto it. Not only do individuals with faith not need further evidence, they will choose not to obtain it if it is offered to them, when their only interest in obtaining it is in how it bears on the decision about the act. For example, I must decline if a third party offers to tell me about her experiences with my friend's secret-keeping abilities. More specifically, individuals with faith will commit to the risky act without looking for further evidence.

I want to be clear that having faith doesn't mean generally avoiding all evidence for or against the proposition in question; rather, it means *not looking for further evidence for the specific purpose of deciding whether to act on the proposition*, or *not basing one's decision on how the evidence turns out*. I also want to make clear that in many cases, a decision to eschew further evidence will be based on evidence one already has: faith need not be "blind" faith. (And, as we will see, faith tends to be rational to the extent that one has already amassed evidence: one must base one's faith on evidence, even though faith involves eschewing further evidence – one must first climb the ladder before kicking it away, so to speak.) So faith involves two key components: taking a risk and doing so without the need for further evidence.

Let us make this analysis of faith explicit. First, we will say that an act *A* constitutes an individual's taking a risk on *X* just in case there is some alternative available act *B* such that *A* is preferred to *B* under the supposition that *X*, and *B* is preferred to *A* under the supposition that $\overline{X}$. For example, my revealing a secret is a risk on my friend's keeping the secret because I prefer to tell her on the supposition that she will keep it and not tell her on the supposition that she won't. Whether an act is a risk on *X* will be relative to the individual performing the act. So we have:

> For an individual *I*, *A* is an act of faith that *X* if and only if *X* is a candidate proposition for faith and:
>
> (1) *A* constitutes *I* taking a risk on *X*.
>
> (2) *I* chooses to commit to *A* before examining additional evidence rather than to postpone her decision about *A* until she examines additional evidence.[6]

So to perform an act of faith in a proposition is to take a risk on that proposition, and to refrain from gathering further evidence before committing to taking that risk.

---

[6] In Buchak (2012), these conditions were formulated in terms of preference rather than choice because decision theory is primarily a theory about preference. I think the view is more intuitive when formulated in terms of choice, and since choice and preference are linked, I see no harm in doing so, although there may still be questions about whether the requirement is ultimately about choice or preferences, if the two do come apart.

Several points of clarification are in order. The first was already mentioned: to have faith that *X* does not require that one in general ignore future evidence in the matter of *X*. What it requires is that *one choose to commit to the relevant act without first gathering additional evidence*. For example, that the theologian has faith that God exists is compatible with her continuing to study theology, because she does not base her commitment to the Christian life on the results of her study. Indeed, gathering evidence can itself sometimes be a faithful act if the evidence is gathered for purposes other than committing to a further act (though of course the question of how such evidence will then bear on one's faith arises). For the theologian, devoting herself to theological study is itself an act of faith because doing so constitutes taking a risk that God exists (if God does not exist, to study theology is a waste of her time, but if God does exist, then theological study will lead to a deeper understanding of God), and because she is willing to devote herself to study without first verifying through other means that God exists (praying for a sign, for example). Faith requires not engaging in an inquiry *for the purpose of* figuring out whether to take the risk on the claim in question.

The second point of clarification concerns the move from acts of faith to faith itself. Whether one has faith that *X* is a matter of which acts of faith that *X* one is willing to perform. Just as belief comes in degrees, so too does faith. And one's degree of faith will be a matter of which risks one is willing to take on *X* without looking for further evidence. I might have enough faith that God exists to attend a house of worship (a low-stakes risk) without gathering further evidence, but I might not have enough faith that God exists to donate all my money to charity and take up a life of poverty (a high-stakes risk) without gathering further evidence.[7] Faith simpliciter, then, is a matter of one's dispositional profile.

Given that faith simpliciter is determined from the acts of faith one is willing to perform, one point to note about the requirements for a proposition's being a candidate for faith in combination with the second condition in the above account of faith is that the account can distinguish between propositions that a risky act reveals faith in and propositions that are presuppositions of the act. When we say that an act constitutes a risk on *X*, in the sense that it is the preferred act if *X* holds but not the preferred act if $\overline{X}$ holds, this preference is determined against the individual's background credences. Donating all of one's money to Oxfam only constitutes taking a risk on {God exists} if one assumes that if God exists, God commands extreme charitable giving and that Oxfam is the most efficient charity. Therefore, this act also constitutes taking a risk on {if God exists, then God commands poverty} and {Oxfam is the most efficient charity}. Refraining from a prenuptial agreement only constitutes taking a risk on {my spouse will continue to be committed to me} if one assumes that it will be financially beneficial to separate one's

[7] I don't have a formal definition of how precisely one's degree of faith is determined from which risks one is willing to take without gathering further evidence, though at the very least being willing to take higher-stakes risks indicates having more faith. I also don't mean to suggest anything about the structure of degrees of faith, e.g., that they can be represented cardinally rather than ordinally.

assets from his in the event of a divorce. Therefore, refraining from a prenuptial agreement also constitutes taking a risk on {a prenuptial agreement will be financially beneficial in the event of divorce}. But we wouldn't want to say that one has faith in all of these propositions. And, on my account, we don't: the proposition that a prenuptial agreement will be financially beneficial is not an object of faith, because one is indifferent to the truth of that claim. One does not have faith in {Oxfam is the most efficient charity} even though one is not indifferent to the truth of this claim, if one is willing to research charities further before donating. Similarly, whether one has faith that {if God exists, then God commands poverty} is a matter of whether one is willing to read the relevant religious texts to get a better idea of what God commands.

3. When and Why Faith is Rational

What, then, is the relationship between rational faith and degrees of belief? (Another way of putting this question is: under what evidential conditions is it rational to have faith?) Recall that for a proposition to be an object of faith, one must not have $p(X) = 1$: one must not be certain, on the basis of the evidence alone, that X holds. Assuming that X meets the other conditions for being an appropriate object of faith (one cares whether X and has a positive attitude towards the truth of X), we can characterize when it is rational to have faith that X by considering when it is rational to perform risky acts on X without more evidence.

Again, every act can be thought of as a lottery which yields various results in various states. To commit to act A is just to take the gamble that yields $A\&X$ if X obtains and $A\&\overline{X}$ if $\overline{X}$ obtains. This is to say: committing to A can be thought of as holding a lottery ticket that yields $A\&X$ with probability $p(X)$ and $A\&\overline{X}$ with probability $p(\overline{X})$. Committing to tell one's secret to the friend, without looking for further evidence, is a lottery which results in telling one's secret and having it kept, with whatever probability one now assigns to the proposition that the friend will keep the secret; and which results in telling one's secret and having it spilled, with whatever probability one now assigns to the proposition that the friend won't keep the secret.

An act of evidence-gathering can be thought of, in the simplest case, as an act which will result in one of two "evidence-receiving" events, E or $\overline{E}$, one which raises the individual's degree of belief in X and the other which lowers it.[8] Let us assume that E is the "confirmatory" evidence: confirmatory in the sense that it increases one's degree of belief in X, so that the probability of X given E is higher than the probability of X without this information ($p(X \mid E) > p(X)$). And let us assume that $\overline{E}$ is the "disconfirmatory" evidence: disconfirmatory in the sense that it decreases one's degree of belief in X, so that $p(X \mid \overline{E}) < p(X)$. For example, the act of gathering evidence about one's friend's secret-keeping

[8] See Good (1967) for the general framework for this discussion.

ability by asking a third party either results in the third party saying that she trusts the friend, which raises one's credence in the proposition that the friend will keep the secret, or results in her saying that she does not trust the friend, which lowers one's credence.

Gathering the evidence and then choosing an act can also be thought of as a particular kind of gamble. One predicts whether one will do *A* or *B* if *E* obtains by determining what one will prefer when one assigns credence p(*X* | *E*) to *X*. For example, one predicts whether one will reveal one's secret if the third party says the friend is trustworthy by considering one's credence that the friend will keep the secret conditional on the third party saying the friend is not trustworthy. Doing *A* when one gets the confirmatory evidence *E* results in *A&X* if *E&X* and results in *A&*$\bar{X}$ if *E&*$\bar{X}$: it results in telling one's secret and having it kept if the third party says the friend is trustworthy and she is in fact trustworthy, and it results in telling one's secret and having it spilled if the third party says the friend is trustworthy but the friend is not. One likewise predicts whether one will do *A* or *B* if $\bar{\text{E}}$ obtains. Thus, gathering the evidence and then acting is a gamble with four relevant states, *E&X*, *E&*$\bar{X}$, $\bar{E}$*&X*, and $\bar{E}$*&*$\bar{X}$. (In our example: the third party says the friend is trustworthy and she is; the third party says she is trustworthy and she is not; the third party says she is not trustworthy but she is trustworthy; and the third party says she is not trustworthy and she is not.) The case in which the individual will do *A* either way is just equivalent to the act of committing to *A*, and so will have the same results as committing to *A* does, or worse results if the experiment has some cost. The more interesting case is that in which the individual will do *A* if *E* but *B* if $\bar{E}$: reveal the secret if the third party says the friend is trustworthy, and don't reveal it if she does not. In this case, gathering the evidence and then choosing an act is the lottery that yields *A&X* with probability p(E&X); *A&*$\bar{X}$ with probability p(E&$\bar{\text{X}}$); *B&X* with probability p($\bar{\text{E}}$&X); and *B&*$\bar{X}$ with probability p($\bar{\text{E}}$&$\bar{\text{X}}$). (Again, in our example: the secret is told and kept if the third party says the friend is trustworthy and she is; it is told and not kept if the third party says the friend is trustworthy and she is not; it is not told but would have been kept if it had been told, if the third party says the friend is not trustworthy but she is; and it is not told and would not have been kept if it had been told, if the third party says the friend is not trustworthy and she is not – and each of these events can be assigned some probability.) Additionally, we might consider that performing the experiment rather than committing to the act first is *costly*, in which case we can adjust the outcomes of this lottery to include the relevant costs.[9]

Whether committing to *A* or basing one's decision on further evidence is rational is determined by which of these acts has the highest expected utility, or on the more general theory I mentioned above in which individuals' risk-attitudes are allowed to vary, which has the highest risk-weighted expected

[9] Costs might be different for different states or outcomes, e.g., if it is costly to postpone doing *A* but not doing *B*, or if it is costly to postpone doing *A* just in case *X* obtains, but not otherwise.

utility. To see which general conditions will make committing rather than gathering evidence rational, note that the evidence-gathering gamble has three important features. First, it may have costs in one or more states, if there are costs associated with postponing the commitment. Second, if the state of the world is $\bar{E}$, then because gathering the evidence prompts the individual to do *B*, gathering the evidence produces a benefit if $\bar{X}$ holds (it prompts the individual to do *B* when *B* is better than *A* in the true state of the world), and a harm if *X* holds (it prompts the individual to do *B* when *B* is worse than *A* in the true state of the world). Call this latter possibility ($\bar{E}\&X$) the possibility of *misleading* evidence: misleading in the sense that it leads the individual to (rationally) lower her credence in the true state of the world and thereby perform an action that in fact turns out worse for her. The third thing to notice is that if the state of the world is *E*, then since one will do *A* regardless of whether one gathers the evidence or first commits, the outcome one gets is the same either way (aside from the possible cost).

So, under what conditions will committing to *A* rather than gathering additional evidence before deciding be rational? While I leave the mathematical details aside,[10] we can say that doing so will be rational roughly to the extent that three conditions obtain, where each condition is a necessary condition. First, either postponing the decision has significant costs or the individual is risk-avoidant in the sense mentioned above (that she gives more weight to worse possible states in decision-making), or both.[11] (Keep in mind that postponing the decision can mean tentatively beginning a course that will result in the act but backing out if one gets counterevidence, so for the decision to have costs there must be costs to doing this rather than beginning a course that one intends to maintain, or this possibility must not be available.) Second, the individual already has a high credence in *X*.[12] And third, counterevidence $\bar{E}$ is not highly correlated with the true state of the world: it is correlated enough that $\bar{E}$ will prompt one to do *B* rather than *A*, but the possibility of misleading evidence is still significant. Another way to describe the situation in which the possibility of misleading evidence is significant is to say that evidence against *X* would be inconclusive: evidence against *X* still leaves a significant chance of *X* in fact being the true state of the world. What exactly the risk-attitude, costs, and credences need to be depends on the utilities of the outcomes involved.

I have now outlined the circumstances under which it is rational for an individual to have faith that *X*, expressed by some act *A*. Whether such an act of faith is rational depends on her utilities and credences. And since we are assuming that credences must be proportioned to the evidence, whether an

---

[10] See Buchak (2010) and Buchak (2012).

[11] While the reasons that the risk-avoidant individual will eschew further evidence are too complex to go into here, the rough idea is that for her, the risk of getting misleading evidence – evidence that makes her refrain from doing *A* even though *X* is the true state of the world – outweighs the benefits of evidence that is correctly correlated with the true state of the world.

[12] However, there is no necessary credence threshold, because the required credence depends on the utility of the outcomes involved.

act of faith is rational depends on an individual's utilities and *evidential situation*. Additionally, given that faith simpliciter is a matter of which risks one is willing to take on a proposition, whether faith simpliciter is rational depends on an individual's evidential situation. In particular, *whether faith is rational depends on the evidence the individual currently possesses in the matter of X, as well as the character of the possible evidence in the matter of X*.

I haven't yet said exactly which evidence-gathering acts an individual needs to eschew in order to count as having faith relative to a particular risk. There are a few options here. The first is that the individual needs to eschew any possible act of evidence-gathering whatsoever. Since it follows from the third requirement for the rationality of faith that faith is rational only when the potential evidence takes a particular form, this would have the upshot that faith is almost always irrational. The only exception would be in circumstances in which there is no *possible* evidence that is conclusive enough against *X* to make evidence-gathering rational. This might hold, for example, in the case of free actions yet to be performed. Another option is to hold that in order to have faith, an individual must only eschew examining additional evidence that is currently available. This would have the upshot that an act could count as an act of faith simply because additional evidence isn't available at this time: one has faith that one's spouse isn't cheating simply because there are no private detectives to hire at the moment. This upshot initially sounds jarring, but is perhaps more palatable when we notice that having faith simpliciter (and faith to a particular degree) is itself a modal notion, fixed by the contexts in which one is willing to perform acts of faith. An intermediate proposal is that the evidence-gathering acts the individual must eschew are acts of a type normally available in our world, regardless of whether they are currently available. This has the advantage of preserving the idea that faith is a virtue precisely because it is sometimes necessary in worlds like ours, while making what counts as an act of faith less subject to the particular circumstances of the individual. In any case, I won't take a firm stand on this question here.

The key question of this paper is what follows from my account about the relationship between the beliefs of a rational individual and the propositions she has faith in. We will approach this question by considering the way in which what an individual ought to believe depends on her credence, utility, and evidential situation, since we know how the rationality of faith depends on these factors.

4. Belief and Credence

There is not as of yet philosophical consensus about the relationship between on-off belief and degree of belief in a rational individual. There are really two questions here: how the mental states "belief" and "degree of belief" relate to each other, and how the normative states of justified belief and rational degrees of belief relate to each other. I am concerned with the latter question, though of course their answers have mutual implications. On some views, the question of how belief depends on credence

cannot even be given an answer. This holds of views that seek to eliminate either the notion of belief or the notion of credence from our taxonomy altogether;[13] views that reduce rational credence to justified belief;[14] and views that hold that there is no formal relationship between belief and credence.[15] However, many philosophers hold that justified belief can be reduced to credence, utility, and evidential situation. There are four views which all hold that whether a belief is justified reduces to the individual's credences, are all serious contenders according to the philosophical literature, and all have different implications for faith. These views will be the focus of the discussion here.[16]

Without taking a stand on which view is correct, I will consider, on each view, what the relationship between rational faith and justified belief is: if an individual is doing what she ought both in matters of faith and matters of belief, how do her beliefs relate to what she has faith in? One might think of these views as views about the relationship between two different kinds of entities epistemologists discuss, rational credence and justified belief. Alternatively, one might think of these views as views about the relationship between the degree of evidential support one (rationally) takes a proposition to enjoin, in the sense of how likely it is to be true given the evidence, and whether one ought to believe it.

The reader should note that I am approaching the question of the relationship between faith and belief as the question of how they relate for a rational individual: again, I am concerned with the relationship between the normative states rather than the mental states. However, as with the relationship between belief and credence, the normative question will have some bearing on the descriptive question and vice versa. It is also important to note that in order to canvas a broad range of views about the relationship between credence and belief, I will undoubtedly run roughshod over distinctions between different versions of these views: this paper is the beginning of a discussion about how faith shakes out on these views, not as the final word.

---

[13] Eliminativists about belief include Jeffrey (1970) and Christensen (2004), although the latter holds that the notion of binary belief is useful. Eliminativists about credence include Holton (2008), who accepts a notion of partial beliefs but thinks that credence as traditionally understood is problematic.

[14] See Harman (1986).

[15] I myself am drawn to this view, by consideration of cases that seem to present a problem for any formal reduction of belief to credence, namely cases of bare statistical evidence. Given this, one might wonder why I am concerned with the question how rational faith relates to justified belief according to a formal reduction of belief to credence. I have three reasons for this. The first is that I am not yet certain the problem presented by these cases is insurmountable. The second is that many have found the views I consider here plausible, and so it is worth considering what these views imply about faith, on my view of faith. The third, and most important, reason is that on the view I am drawn to, even if there is no formal connection between belief and credence, there will be a *defeasible* connection, because both belief and credence are based on the same body of evidence. So we might see the four views I will talk about here as illuminating something that is generally or approximately true but not always or exactly true of justified belief and rational credence, because each entity is a rational response to a single body of evidence. Each view can be seen, roughly, as a view about how the circumstances one is in and the values one has determines the character of the evidence required for belief.

[16] This list of views isn't exhaustive. Other very different views include those of Sturgeon (2008) and Frankish (2009).

5. The Certainty View

The first view of the relationship between justified belief and rational credence is known as the *Certainty View*: an individual is justified in believing X only if she (rationally) assigns $p(X) = 1$, that is, only if she is certain of X on the basis of her evidence alone. A related view is the Certainty View about *knowledge*: an individual only knows X if she (rationally) assigns $p(X) = 1$.[17] This view is thought to be motivated in several ways. First, it makes sense of the idea that to hold something to be probable is "less" than to believe it outright (or to know it): degrees of belief are "partial" beliefs, and outright belief is "full" belief. Second, the view isn't susceptible to certain worries (that I will discuss shortly) concerning lottery cases. Third, the view captures the following thought, which some have found plausible: if you are justified in believing something, or knowing something, then there's no evidence you think you might get that will undermine your justified belief or your knowledge. If E is some piece of evidence that would lead you to lower your degree of belief in X, then you assign $p(E) = 0$.

However, this view also faces two serious problems. The first is that many think one *can* be justified in believing (or knowing) even if one thinks that there is possible evidence that could make one change one's mind. (In the credence framework: one can be justified in believing X or one can know X even if there is some piece of evidence E, to which one assigns positive probability, that will make one lower one's credence.) Put succinctly, to justifiably believe or know something is not to be certain of it. The second problem is that those working in the framework of formal epistemology typically hold that there is a strong connection between degrees of belief and betting behavior: in particular, if one holds $p(X) = 1$, then one must be willing to take any bet whose payoff on X is better than not taking the bet: what happens on $\bar{X}$ is irrelevant to the value of the bet. So if an individual's degree of belief in X is 1, she must be willing to take the bet {\$1 if X, -\$1 million if $\bar{X}$}.[18] Most individuals believe things that they are not willing to bet on at such high stakes, and do not seem irrational for doing so. This isn't to say these problems are insurmountable: one might, for example, employ a different theory of action in conjunction with degrees of belief.

So, according to the Certainty View, one is not justified in believing X unless one assigns $p(X) = 1$ on the basis of one's evidence. But I pointed out above that a proposition cannot be an object of faith if one's evidence alone yields $p(X) = 1$. On this view, then, faith is only possible when one *lacks* a justified belief. A similar point holds for the Certainty View about knowledge: if knowing X entails assigning $p(X) = 1$, then one can only have faith in propositions that one does not know. We might couple the

[17] This needn't imply that credence is prior to justified belief or knowledge: this view could instead be formulated as: an individual assigns $p(X) = 1$ if she justifiably believes that X or knows that X.
[18] See Kaplan (1996: 91-93).

Certainty View with some plausible norms for action in order to explain the role of faith. In particular, many have found the following norm plausible: one ought to act on what one (justifiably) believes, or act on what one knows.[19] Furthermore, it follows from the Certainty View that if one is justified in believing some proposition, or if one knows something, then one can act on it without gathering further evidence: from one's own point of view, further evidence can't possibly tell against the proposition. We then get the following conclusion: *faith plays the same role in action in situations in which conclusive evidence is not available, as justified belief or knowledge plays in action in situations in which conclusive evidence is available.* When one has conclusive evidence that *X*, one acts on one's knowledge or justified belief that *X*. When one lacks conclusive evidence for or against *X* (and therefore cannot act on knowledge or justified belief), there are two possibilities. If one's evidential situation is such that one ought to look for more evidence, then one ought to gather more evidence rather than acting on *X*, and so one ought not to act on *X* on faith. But (assuming that *X* meets the conditions for being a candidate for faith) if one's evidential situation is such that one ought not to look for more evidence – if current and potential evidence have the required character – then one rationally acts on faith that *X*.

6. The Threshold View

The second view of the relationship between justified belief and credence that I will examine is the *Threshold View*: one is justified in believing a proposition if and only if one's (rational) credence in a proposition is above a particular threshold. The threshold is typically taken to be very high, e.g. 0.95 or 0.99, though sometimes individuals talk as if 0.5 is the right threshold. But the threshold can be fuzzy or vague. This view is not susceptible to the two problems mentioned above for the Certainty View: the Threshold View allows that one can be justified in believing a proposition for which one thinks counterevidence is possible, and the Threshold View allows that one can be justified in believing *X* while refusing high-stakes bets on *X*: indeed, the threshold will determine the exact stakes of the bets one must accept to count as having a justified belief.[20]

A worry for this view arises in the form of the well-known "lottery paradox."[21] For any candidate threshold, consider a lottery with enough tickets that the probability of each ticket losing is above the threshold. (For example, if the threshold is 0.99, consider a 1000 ticket lottery.) Since a rational individual will have credences that accord with the objective probabilities associated with the lottery, her credence in each proposition of the form "ticket *x* is a losing ticket" will be above the threshold.

---

[19] Hawthorne and Stanley (2011) argue that one can only act on what one knows.

[20] As above, one could adhere to a threshold view of when credence gives rise to knowledge. Indeed, all of the remaining views could be formulated as views about knowledge instead of justified belief, so I will omit this proviso hereafter.

[21] Kyburg (1961).

Therefore, she will be justified in believing of each ticket that it is a losing ticket. Nonetheless, we may assume that she is justified in believing that some ticket must be the winning ticket: so she has a set of justified beliefs that are inconsistent with one another.

Again, this problem may not be insurmountable: one view (known as the "*Lockean*" view) holds that the Threshold View is correct and that the individual is justified in believing of each ticket that is the losing ticket, but denies the principle that one is always justified in believing the conjunction of what one justifiably believes. Thus, the Lockean accepts that one can have justified beliefs that jointly contradict, as long as one does not believe their conjunction.

What of belief and faith, on the Threshold View? The discussion in section 3 pointed out that faith is justified only insofar as one's credence in X is high enough. The minimum credence threshold for a particular act of faith will depend on the characteristics of A and the alternative available acts: in particular, the worse that A is relative to the alternatives in case $\bar{X}$ obtains, the higher the required threshold, and the better that A is relative to the alternatives in case X obtains, the lower the required threshold. Therefore, the higher credence one has, the more acts of faith will potentially be justified, although in order to see whether an act of faith is ultimately justified, we need to check whether one's credences concerning the potential evidence imply that it is rational not to gather further evidence. Still, for a given evidential situation, faith accompanied by justified belief (that is, by credences that support justified belief) will be more robust than faith unaccompanied by justified belief.[22] Importantly, faith in X can be rational – or faith to a certain degree in X can be rational – even when one does not have a justified belief that X. Moreover, if the potential evidence is sparse enough, then faith to a certain degree in X may be rational, even if one is justified in disbelieving X. This is because even if one is justified in disbelieving X, it can be rational to perform acts that are much better than the alternatives in case X holds but not much worse than the alternatives in case $\bar{X}$ holds, and there may be certain evidential situations according to which additional evidence may hurt more than it may help (though these will be very limited). As for the converse, belief might be justified without faith being rational, if one's credence in the candidate proposition meets the threshold but the evidential situation doesn't justify eschewing further evidence for purposes of committing to an act.

7. Pragmatic Reduction

A third type of view of the relationship between credence and belief for a rational individual holds that whether a particular credence profile justifies a belief is not independent of the circumstances

[22] Audi (2011) notes that there are two kinds of propositional faith, doxastic and fiducial, where the former is accompanied by belief but the latter is not. My view of faith, in combination with the Threshold View, can explain Audi's distinction.

the individual finds herself in. What credence is required for a justified belief can depend on how much is at stake in being wrong. For example, one can have a credence that will justify the belief that the sandwich in the refrigerator is almond butter rather than peanut butter in circumstances in which one is selecting a sandwich for a guest who slightly dislikes peanut butter, but the same credence might fail to justify this belief when selecting a sandwich for a guest who is allergic to peanut butter.[23] One justifiably believes that *X* when one's credence is enough for practical purposes.

There are two versions of this view: one holds that pragmatic factors affect when a belief is *justified* on the basis of a credence profile and the other that pragmatic factors affect when a credence profile gives rise to a *belief* at all.[24] So, that pragmatic factors encroach on justified belief can be a thesis about how these factors encroach on justification, or about how they encroach on belief – but since our topic is the relationship of rational faith to justified belief, it will not matter which.

The basic idea is that an individual can have a justified belief that *X* if and only if she is rational in assuming that *X* for purposes of action, in a relevant range of circumstances. Ross and Schroeder provide a helpful summary of this type of view, which they call Pragmatic Credal Reductivism (PCR):

> "[Proponents of PCR] maintain that what it is to believe that [*X*] is to have a sufficiently high credence in [*X*] to rationalize acting as if [*X*] 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."[25]

Notice that to act as if *X* can mean two different things. The first is to explicitly adopt *X* in decision-making; the second is to perform the same actions one would perform if one was certain that *X* held. Harsanyi's (1985) view takes the first line. According to Harsanyi, a rational individual believes that *X* if it is permissible for her to use her conditional credence function p($\mid X$) – the result of her updating her credences on *X* – rather than to use her "unconditional" credence function p( ), in decision-making. When one uses one's conditional credences in decision-making, one "accepts" or assigns "direct practical certainty" to the statement conditionalized on. One assigns persistent practical certainty to a statement *X* if one chooses to assign practical certainty to *X* in all decisions one will face until explicitly deciding to reconsider. (Harsanyi doesn't explicitly say that his theory is a theory of *belief*, but because of the

---

[23] Example suggested by Ross and Schroeder (2012).

[24] For an example of the former, see Fantl and McGrath (2002). See also Pace (2011), who argues that *moral* factors determine when a belief is justified. For an example of the latter view, see Weatherson (2005). Since Harsanyi's (1985) discussion is stated in terms of an ideally rational individual, he could be interpreted in either way; however, given that his discussion is framed in terms of whether one ought to accept a statement (rather than whether one *does*), it seems to me more natural to interpret his as a theory of when acceptance or belief is justified. For a different kind of view on which pragmatic factors determine when an individual is justified in believing, see Ross and Schroeder (2012). Unfortunately I do not have space to consider the relationship between justified belief and rational faith on Ross and Schroeder's account.

[25] Ross and Schroeder (2012: 5).

problem he was responding to, and the similarity to later theories that explicitly mention belief (such as Weatherson (2005)), we can proceed as if belief corresponds to persistent practical certainty.)

Weatherson's (2005) view takes the second line. According to Weatherson, one believes that *X* if and only if one's preference ordering of the relevant actions according to one's conditional credence function p(| *X*), is the same as one's preference ordering according to one's unconditional credence function p( ). The difference between these formulations is that Harsanyi's allows that one's conditional preference ordering can be different from one's unconditional preference ordering, as long as the differences are small enough to be outweighed by the cognitive costs of using one's conditional credences.[26]

Pragmatic Credal Reductivism has the advantage of explaining what is going on in cases like the almond butter case above. Nonetheless, this view has its downsides as well. If belief, justification, and knowledge are supposed to be purely epistemic phenomena, phenomena that aim at the truth, it is unclear why non-truth-conducive factors like stakes should encroach on the standards they must meet. The previous two views accepted evidentialism about belief; this one does not.

With these expository details out of the way, let us consider the relationship between faith and belief on Pragmatic Credal Reductivism. Recall first that the set of propositions that are candidates for faith is a subset of the propositions that are candidates for belief. So there will be some things that a rational individual believes but fails to have faith in because they are not candidates for faith in the first place. Setting these aside, however, we can ask, for propositions that are candidates for both faith and belief, how the conditions for rational faith relate to the conditions for rational belief. Since one's degree of faith is a matter of whether one has faith within a particular context or relative to some particular set of acts, the question will be whether, if one has the credence required for faith within some context, one has a justified belief within this context, and vice versa.

To remind the reader, in order for an individual to have rational faith that *X*, within a particular context, *X* must belong to the set of propositions that are candidates for faith, and the individual's credences must satisfy two conditions, each for a particular set of acts:

(1) It must be practically rational to act as if *X* for some relevant set *F* of pairs of actions, the choice pairs that involve taking a risk on *X*.

(2) It must be practically rational to commit to acting as if *X* for choices in set *F* rather than to first look for new evidence: that is, it must be practically rational to commit to acting as if *X* for choices in set *F* rather than to first perform some act in *G*, the set of relevant acts that constitute looking for new evidence. (The risky act in each pair in *F* must rank above each act in *G*.)

[26] Weatherson adds that a threshold of 0.5 is necessary for belief, but we will ignore that detail in our general discussion.

And in order to justifiably believe that *X*, within this context, the individual's credence must satisfy one condition:

(3) It must be practically rational to act as if *X* for some relevant set *H* of actions, the set of acts relevant to the question of whether one believes.

Now we can consider the relationship between credences that satisfy (1) and (2), and credences that satisfy (3), for a given utility function, by considering the relationship between the sets *F* and *G* on the one hand, and *H* on the other hand. *F* is a set of pairs of acts, but we can consider *F**, the set of all acts that are in some pair in *F*. If it is rational to act as if *X* for acts in *F**, then it is rational to acts as if *X* for all pairs in *F*, and vice versa. This is just to say: we don't need to separately consider what it is rational to do when considering different risks on *X*: if it is rational to act as if *X* for all acts in a set, then it is rational to take all of the risks on *X* that involve a choice between two acts in this set, and vice versa.

We must also consider how evidence-gathering relates to acting as if *X*. Consider some pair of acts in *F*: *A* (the risk on *X*) and *B* (the alternative). How does performing some evidence-gathering act in *G* rank compare to committing to *A*, if one uses one's credences conditional on *X* rather than one's unconditional credences? Recall from section 3 the idea of evidence-gathering acts. If we assume that the confirmatory evidence *E* will support doing *A* and the disconfirmatory evidence $\bar{E}$ will support doing *B*, then to perform a particular two-outcome evidence-gathering act is to take the gamble {A&X if E&X, A&$\bar{\text{X}}$ if E&$\bar{\text{X}}$, B&X if $\bar{\text{E}}$&X, B&$\bar{\text{X}}$ if $\bar{\text{E}}$&$\bar{\text{X}}$}. If one evaluates this gamble using one's conditional credences p(| X), then the only possible states will be *E&X* and $\bar{E}$*&X* – because $\bar{X}$ states are ruled out – and so the only possible outcomes will be *A&X* and *B&X*. But since *A&X* is strictly preferred to *B&X*, this gamble will be strictly dispreferred to committing to act *A*. If we assume that both *E* and $\bar{E}$ will support doing *A*, then the individual will be indifferent between committing to *A* and gathering the evidence – or will strictly prefer committing to *A* if not doing so has costs. Therefore, acting as if *X* with respect to the choice between *A* and *B*, and with respect to gathering evidence before making this choice, requires (strictly or weakly) preferring to commit to *A* rather than to first gather additional evidence.

What this discussion implies is that if $\text{F}^* \cup \text{G} \subseteq \text{H}$, then if one's credences support justified belief that *X*, they support rational faith that *X*. Less formally: if the set of acts relevant to whether an individual has a justified belief that *X* includes all the relevant acts that constitute taking a risk on *X* (or failing to do so) and includes all the relevant acts that constitute gathering evidence about whether *X*, then if one has a justified belief that *X*, one's faith in *X* is rational; but it may be possible to have rational faith without justified belief. Conversely, if $\text{H} \subseteq \text{F}^* \cup \text{G}$, then one's credences support rational faith that *X* only if they support justified belief that *X*. Less formally: if all of the acts relevant to whether an individual has a justified belief that *X* are either relevant acts that constitute taking a risk on *X* (or failing to do so) or

relevant acts of evidence-gathering, then one must have a justified belief that *X* in order to have rational faith that *X*; but it may be possible to have a justified belief without faith being rational.

Which, if either, of these conditions holds will depend on our views about the relevant sets. It is natural to hold that *H*, the actions relevant to belief, are the actions available to the individual or the actions she is seriously considering, in the context.[27] Thus, aside from the question of the relevant evidence-gathering acts, *F** and *H* coincide.[28] What is the relationship between *G* and *H*? Recall the discussion in section 3 about which set of evidence-gathering acts the faithful individual is supposed to eschew. I mentioned three possibilities: all of the possible acts, all of the acts normally available in the context in question to creatures like us in our world, and all of the acts actually available in the context. For the first and second possibility, the acts relevant to whether belief is justified are a proper subset of the relevant acts for rational faith, and so a rational individual can only have faith if she has a justified belief, but justified belief is not enough.[29] For the third possibility, the acts relevant to whether a belief is justified are the same as the acts relevant to whether faith is rational, and so a rational individual has faith exactly when she has a justified belief. This third possibility, then, suggests the view that faith is belief with a particular kind of content.

To summarize: on PCR, there is a close relationship between when a rational individual has faith that *X* and when she believes that *X* within a context: both facts are determined by whether her credence in *X* is sufficient to justify acting in a particular way in a particular set of circumstances determined by the context. Depending on how one delineates the set of circumstances – and assuming that *X* is a candidate for faith in the first place – the two facts may have identical truth conditions, or their truth conditions might come apart: in particular, on some plausible delineations, rational faith that *X* is stronger than justified belief that *X*, because it requires acting as if *X* with respect to a wider range of acts, namely particular acts of evidence-gathering.

8. Epistemic Reduction

The fourth kind of view we might call *Epistemic Credal Reductivism*. This view agrees with Pragmatic Credal Reductivism that whether a set of credences justifies a belief depends on the utilities

[27] For a more detailed proposal, see Weatherson (2005).

[28] There might be acts in *H* that are not part of pairs that constitute risks on *X*, because whether *X* holds does not matter to the value of these acts. However, in this case it is trivially true that one can act as if *X* with respect to these acts: conditionalizing on *X* does not change their value.

[29] Another possibility is that acts of evidence-gathering outrun acts relevant to belief because the former do not count as ordinary acts. This might be plausible to PCR theorists who want justified belief that *X* to be determined solely by one's credence in *X*, not by one's credences in *X* conditional on various evidence-gathering results. For those who hold to this view, rational faith that *X* will again require more than justified belief that *X*: justified belief will be a matter of one's current credence in *X* and one's utilities, and rational faith will be a matter of these two things plus a view about the evidence.

involved. However, on this view, the utilities must track purely epistemic values. Consider William James's (1896) point that we must balance two epistemic goals: knowing truth and avoiding error. This can be cashed out in our framework as follows. First, we can treat the act of believing some proposition and the act of not believing it as acts evaluable in the standard way by the apparatus of decision theory. The relevant states are that the proposition is true and that the proposition is false; and the outcomes, therefore, are believing the proposition when it is true, believing it when it is false, failing to believe it when it is true, and failing to believe it when it is false. These outcomes are each assigned an "epistemic utility."[30] Naturally, one important feature of a candidate proposition for belief is whether it contradicts another proposition one believes. To take account of this consideration, the apparatus of decision theory can be applied to the act of believing an entire set of propositions, rather than separately to each proposition under consideration.

Epistemic Credal Reductivism (ECR) is attractive to those who hold that the only factors relevant to whether one has a justified belief are epistemic factors, but that some of those factors can be up to the individual to trade off. It is also attractive to those who hold that belief involves the will in some way: on Epistemic Credal Reductivism, since belief is evaluated in the way any ordinary act is evaluated – does a belief maximize expected utility relative to the other options, in this case not believing? – it is easy to fit this view together with a picture on which belief *is* an act of some kind.

The view sketched so far is rough, and there may be several ways to give formal content to it. One way is found in Levi (1967).[31] Consider an individual who partitions the world into three possibilities, *X*, *Y*, and *Z* (in an example Levi considers, the possibilities are that each one of three candidates wins an election). The individual has eight possible "cognitive options," distinguished by which proposition she accepts as the logically strongest she is committed to: accept the disjunction of all three, accept the disjunction of any two, accept one of the possibilities, or accept all of the possibilities (i.e., accept a contradiction). There are two desiderata, which correspond to knowing truth and avoiding error: relief from agnosticism, i.e., accepting some possibility as strongest when that possibility in fact obtains, and avoiding error, i.e., refraining from accepting some possibility as strongest when that possibility in fact does not obtain.

To illustrate, here is a toy example, involving a partition into only two states of the world. Let us consider an individual deciding whether to believe that God exists (G), believe that God doesn't exist ($\overline{G}$), or remain agnostic. And let us assume that u(G & believe(G)) = u($\overline{G}$ & believe($\overline{G}$)) = 0.9; that u($\overline{G}$ & believe(G)) = u(G & believe($\overline{G}$)) = -0.1; and that u(G & believe(G or $\overline{G}$)) = u($\overline{G}$ & believe(G or $\overline{G}$)) = 0.7.

[30] Here we are only assigning utility values to on-off beliefs in the possible states. But one might also assign values to having various degrees of belief in possible states: see the literature on "scoring rule" arguments for probabilism.
[31] See also Maher (1993), and the somewhat related view in Kaplan (1996).

Then, before the question of evidence-gathering comes into the picture, the expected utility maximizer will need a credence of at least 0.8 in one or the other of the hypotheses in order to believe it rather than to remain agnostic. (The risk-avoidant individual will need a slightly higher credence than this in one or the other of the hypotheses, since the possibility of having a false belief plays a larger role in her evaluation of the act of believing.) How does having a justified belief bear on whether to acquire further evidence, on this type of view? According to Levi, justified belief in a proposition is a matter of one's credence in that proposition at a particular time, and both justified belief and agnosticism are compatible with continuing to look for additional evidence. In particular, Levi cautions that we shouldn't assume that a decision to collect evidence should be seen as a decision to remain agnostic.[32]

According to ECR, the act of believing that *X* is to be evaluated in the same way we evaluate any ordinary act. Recall that according to the above analysis of faith, propositional faith involves a disposition to choose certain acts, and its rationality is evaluated with respect to one's utilities for the possible outcomes of these acts as well as one's credence concerning the target proposition and the potentially available evidence. So we've been asking whether, for example, faith that my friend will keep my secret, expressed by the act of telling her my secret, is rational. But since belief is itself an act, we can also ask whether faith that my friend will keep my secret, expressed by the act of believing that my friend will keep my secret, is rational. That is, we can ask about the status of *belief as an act of faith*. To have a justified belief that *X*, according to ECR, is to have a high enough credence in *X* that believing *X* maximizes (risk-weighted) expected utility. To be rational in believing as an act of faith is to have a high enough credence in *X* to meet this condition *and* for the character of potential new evidence to be such that it is rational not to gather more evidence.

Therefore, restricting ourselves to the set of propositions whose contents make them candidates for faith, there will be three different cases of justified belief. The first case is when one's credence in *X* is 1. In this case, belief in *X* is justified, but the evidence itself rules out the proposition as a candidate for faith. The second case is when one's credence in *X* justifies belief in *X* despite being less than 1, but one's credences also imply that one ought to look for more evidence. In this case, belief in *X* is justified, and *X* is a candidate for faith, but belief in *X on faith* is not justified. For example, if the relevant experiment concerning *G* is cost-free and potentially conclusive, then the individual with $p(X) > 0.8$ will seek out more evidence even though she has a justified belief. Finally, one's credence in *X* might justify belief in *X* despite being less than 1, and one's credences might also imply that one ought to eschew further evidence. In this case, belief is justified and belief by faith is also justified, because the potential additional evidence is such that a rational individual will not seek it out even though she is not certain of

[32] Levi (1967: 144, ff. 7). One might dispute this claim, in which case the second of the three cases mentioned below will be ruled out.

the proposition. There will not be any cases in which belief by faith is rational but belief is not, because the conditions for rational faith are strictly stronger than the conditions for rational belief. All this is to say: believing that *X*, as an act of faith, is rational if the belief itself is justified according to one's current evidence and if, additionally, it is rational not to look for further evidence. Thus, on this view, there are some things the rational individual is certain about; there are some things she believes but will continue to explore; and there are other things she believes on faith: she believes them, and even though she is not certain, she does not seek additional evidence.[33] These acts of belief without certainty will be acts of faith: the individual not only believes, but commits to believing without looking for further evidence.[34]

One final point about the relationship between rational faith and justified belief bears mentioning. All of the reductions of justified belief to rational credence mentioned here are synchronic. Justified belief is a matter of meeting certain requirements at a time: credence 1, credence above a threshold, or credence that justifies a particular preference ordering of the possible actions or cognitive options. Faith involves a commitment to the act in question. Whether this commitment is rational is determined at the time of making it. However, commitments impose diachronic constraints. An individual might have reason to act in accordance with a prior commitment, even if she wouldn't have reason to perform the same action in the absence of a commitment. Thus, one additional way in which faith and belief can come apart is that an individual might carry through with an act of faith because she is committed to it – where committing to it was rational and she had a justified belief in the relevant proposition – even though she no longer has a justified belief. For example, if we adopt PCR, then while committing to an act before looking at evidence (acting as if *X*) might be rational before looking at the evidence, performing the act after looking at the evidence might not be. If we adopt ECR, then holding a belief might maximize cognitive utility before looking at evidence, and it might be rational to commit to holding the belief before looking at the evidence, but once the evidence is in, the belief might no longer maximize cognitive utility. Thus, one role that faith can play, particularly on the last two views, in which the requirements for rational faith at a time may be stronger than the requirements for justified belief at that

[33] One additional feature of Levi's picture worth mentioning is that he adopts a twofold classification: acceptance, and acceptance *as evidence*. Propositions accepted as evidence are propositions that one will not continue to look for evidence for. If we accept Levi's classification, then whenever the rational individual accepts a proposition as evidence without being certain of it, she accepts it on faith.

[34] We might also ask about the relationship between justified belief that *X* and rational faith that *X* expressed by acts other than believing. If it is rational to have faith in *X*, expressed by an act riskier (in utility terms) than belief, then it follows that it one is justified in believing that *X*. But if one is justified in believing that *X*, then we cannot even conclude that less risky acts of faith are justified, because whether they are justified depends on the character of the available evidence.

time, is that faith can provide a reason to perform an act or maintain a belief even when the act or belief ceases to be justified on its own.[35]

9. Conclusion

I have considered four views about the relationship between justified belief and rational credence, and what each view says about the relationship between justified belief and rational faith. Since these views can be seen as views about what degree and type of evidential support is required for belief, each view has implications for how the evidential situations that justify belief relate to the evidential situations that make faith rational. On the Certainty View, faith and belief do not overlap at all: a rational individual can only have faith in those propositions she does not have sufficient evidence to believe. Furthermore, faith can be seen as an *alternative* to belief as a basis for action. On the Threshhold View, faith can be accompanied by belief or unaccompanied by belief – since one's credence does not need to meet a threshold in order to make faith rational – but faith accompanied by belief is more robust than faith unaccompanied. Furthermore, there can be cases of rational agents who believe but lack faith, if their evidence supports a credence above the threshold but the character of the potential evidence implies that they should continue to look for more evidence concerning the proposition in question. On Pragmatic Credal Reductivism, the requirements for justified belief in a proposition, within a context, are similar to the requirements for rational faith in that proposition within that context: both require credences according to which it is rational to act as if the proposition in question is true. Rational faith in a proposition requires justified belief in the proposition. However, depending on which acts are relevant to the question of whether one counts as acting as if the proposition is true, rational faith might require more, namely that one must eschew not only the evidence-gathering acts *available* to one, but others as well: all of the possible evidence-gathering acts, or all of the normally available ones. Finally, on Epistemic Credal Reductivism, believing a proposition can itself be an act of faith. Such an act will be rational only when one's current evidence justifies the belief, and the character of the potential future evidence makes it rational to look no further into the matter.

One debate in religious epistemology is whether rational religious faith requires justified religious belief. Other questions about the nature and ethics of belief raise the question of whether having justified religious beliefs necessarily entitles one to make the practical commitments of religious faith. One upshot of this paper is that we cannot reach consensus on these questions at present because of the more general epistemological lack of consensus over how belief relates to rationality: in particular, over how belief

[35] One might wonder why this kind of case does not illustrate that it can be rational to renege on one's earlier commitments. My view is that in these cases, one's over-time reasons and one's at-a-time reasons conflict in ways they don't in ordinary cases. One might interpret what is going on in these cases as the coexistence of faith and doubt.

relates to the degrees of belief it is rational to have given one's evidence. Having rational religious faith might or might not require one to be justified in one's religious beliefs. Justified religious belief might or might not entail that it would be rational for one to have faith. These questions turn on broader epistemological debates, and this paper has begun the discussion of how resolving these debates will answer these questions.

BIBLIOGRAPHY

Audi, Robert (2011). Rationality and Religious Commitment. Oxford, UK: Oxford University Press.

Buchak, Lara (2010). "Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering," *Philosophical Perspectives* 24: 85–120.

Buchak, Lara (2012). "Can it be rational to have faith?" Chapter 12 of Probability in the Philosophy of Religion, eds. Jake Chandler and Victoria S. Harrison, Oxford University Press.

Buchak, Lara (2013). Risk and Rationality. Oxford, UK: Oxford University Press.

Christensen, David (2004). Putting Logic in its Place: Formal Constraints on Rational Belief. Oxford: Oxford University Press.

Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics, and Justification." *The Philosophical Review* 111(1): 67-94.Frankish, Keith (2009). "Partial Belief and Flat-Out Belief," Degrees of Belief, eds. Franz Huber and Christoph Schmidt-Petri. Springer.

Good, I. J. (1967) 'On the Principle of Total Evidence', *British Journal for the Philosophy of Science*, 17, 4: 319–321.

Harman, Gilbert (1986). Change in View. MIT Press.

Harsanyi, John C. (1985). "Acceptance of Empirical Statements: A Bayesian Theory Without Cognitive Utilities." *Theory and Decision* 18 (1):1-30.

Hawthorne, John and Jason Stanley (2011). "Knowledge and Action." *Journal of Philosophy*.

Holton, Richard (2008). "Partial Belief, Partial Intention," *Mind* 117: 27-58.

James, William (1896). "The Will to Believe." Reprinted (excerpt) in Reason and Responsibility, eds. Joel Feinberg and Russ Shafer-Landau. Wadsworth/Thomson 2005. 101-109.

Jeffrey, Richard C. (1970). "Dracula Meets Wolfman: Acceptance vs. Partial Belief," Induction, Acceptance, and Rational Belief, ed. Marshall Swain, Reidel.

Joyce, James M. (2005). "How Probabilities Reflect Evidence," *Philosophical Perspectives*, 19: 153–178.

Kaplan, Mark (1996). Decision Theory as Philosophy. Cambridge, UK: Cambridge University Press.

Kyburg (1961). *Probability and the Logic of Rational Belief*. Wesleyan University Press.

Levi, Isaac (1967). Gambling with Truth. Alfred A. Knopf, New York and Routledge & Kegan Paul, Ltd., London.

Maher, Patrick (1993). "The Concept of Acceptance." Chapter 6 of Betting on Theories. *Cambridge Studies in Probability, Induction, and Decision Theory*. Cambridge, UK: Cambridge University Press.

Pace, Michael (2011). "The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe'" *Noûs* 45 (2):239-268.

Ross, Jacob and Mark Schroeder (2012). "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research*.

Sturgeon, Scott (2008). "Reason and the Grain of Belief." *Nous* 42(1): 139-165.Weatherson, Brian (2005). "Can We Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19: 417-43.