

Can it be Rational to Have Faith?

Lara Buchak

I. Introduction

My concern here is the relationship between faith and practical rationality. I seek to develop a unified account of statements of faith concerning mundane matters and those concerning religious faith. To do so, I consider the sense in which faith requires going beyond the evidence, and argue that faith requires terminating the search for further evidence. Having established this, I turn to the question of whether it can still be rational to have faith; arguing that, contrary to common assumptions, there need be no conflict between faith and rationality. We shall see that whether faith can be practically rational depends both on whether there are extrinsic costs associated with postponing the decision involving faith and the extent to which potential counter evidence would be conclusive.

II. Preliminaries

I begin with the idea that faith statements in religious contexts and in more mundane contexts express the same attitude and so share some typical characteristics. By 'faith statements' I simply mean statements involving the term 'faith'. The following are representative examples:

I have faith in your abilities.

He has faith that his spouse won't cheat on him.

I have faith in you.

He has faith that you won't reveal his secret.

She acted on faith.

She has faith that her car will start.

It was an act of faith.

I have faith that God exists.

I have faith in God's goodness.

I have faith in God.

These statements share three important features: they involve a relationship between the agent and a particular proposition, between the agent and a particular (actual or potential) action, and between the proposition and the evidence the agent currently possesses.

The first thing to notice is that faith statements typically involve a proposition to which the actor involved acquiesces. This is obvious in the case of 'faith that'

statements: when a person has faith that p , he acquiesces to p .¹ It is also clear in the case of those statements that can be easily translated into 'faith that' statements: for example, the statement *I have faith in your abilities* is equivalent to *I have faith that you will be able to do such-and-such*. It is less obvious in the case of those statements claiming that an individual has faith in a person; however, upon further inspection, having faith in a person does typically require acquiescing to particular propositions about that person. For example, having faith in a person might involve acquiescing to the claim that the person will do the right thing or will succeed at a task, and having faith in God might involve acquiescing to (at least) the claims that God exists and that God is good. By the same token, performing an act of faith or acting on faith seems to involve acquiescing to a proposition, and which proposition one acquiesces to will be set by the context. For example, if setting down one's own weapons is an act of faith, then this is because setting down one's own weapons involves acquiescing to the claim that the other person will then set down his.

The second thing to notice about faith statements is that the truth or falsity of the proposition(s) involved is ordinarily a matter of importance to the actor. For example, it does not seem apt to state that I have faith that the Nile is the longest river in Egypt, because I do not care whether or not this is true. We do not attribute faith to a person unless the truth or falsity of the proposition involved makes a difference to that person. I might consider whether to have faith that *my spouse won't cheat on me* or *that my car will start* precisely because it makes a difference to me whether or not these things are true.

Along the same lines, having faith typically involves an action: a person's having faith in something should make a difference to her behaviour. However, this needn't be an actual action. It would be enough for faith that if a person were put in a particular situation, she would then manifest the relevant behaviour (assuming that there are no forces that would stop her). Faith is thus linked to a *disposition* to act.

This brings us to the next point about the relationship between faith and behaviour: it seems that one can have faith in a particular proposition relative to one action but not to another. For example, I might have faith that my car will start when I only need to drive to work but lack that faith when I am relied upon to carry a life-saving organ to the hospital (as evidenced by the fact that I may double-check my engine or arrange for a backup mode of transportation in the latter case but not the former). A person might have faith in God when it comes to giving weekly donations to the poor but lack faith in God when it comes to allowing himself to be martyred.

¹ I speak of acquiescing to a proposition rather than believing it because I am not sure that if I have faith in something, I thereby believe it. While it sounds infelicitous to say 'I believe that $\sim X$ but I have faith that X ', there may not be anything wrong with saying 'I don't know whether X —I have no idea whether I believe that X or not—but I have faith that X '. So as not to prejudge that issue, I make a weaker claim: that having faith involves taking the proposition to be true, that is, 'going along with it', but not necessarily adopting an attitude we might describe as belief.

There are two ways in which we might interpret the fact that one might have faith when it comes to the performance of some actions but not others: we might say that faith is context-dependent, or we might say that faith comes in degrees. There is something to be said in favour of each of these approaches. However, whether one has faith in X expressed by a particular act A will be determinate on either approach, and since this will be our basic unit of analysis in this paper, we needn't choose between them.

The next thing to bring into the picture is the relationship between the agent who has faith in X (expressed by some act A) and the evidence he has for X . We make assertions of faith only when the status of the proposition involved is uncertain or when the evidence we have is inconclusive. For example, when a friend is worried about the outcome of an exam, we might reassure her by saying 'I have faith that you passed'; however, once she shows us that she got an A we would no longer say this. Clearly, this is not because we are less willing to acquiesce to the claim that she passed, but because we now know for certain that she did. For similar reasons, it seems odd to claim to have faith in logical truths.

These considerations suggest that a person cannot have faith in propositions of which he is antecedently certain or for which he has conclusive evidence.² Are there further restrictions on which propositions a person can have faith in? I don't believe so. Indeed, a person may have no evidence at all for the proposition he has faith in, or even may have evidence that tells *against* the proposition. For example, we could imagine someone saying 'Although she's spilled all the secrets I've told her so far, I have faith that this time will be different', or 'I don't think there's any evidence that God exists, but I have faith that he does'. Therefore, that a person has faith that X implies nothing about his evidence for X , aside from its inconclusiveness. Statements in which the actor has faith despite no or contrary evidence do seem correctly described as cases of faith, even though they are not cases in which we are inclined to think that the actor is *wise* to have faith; rather, we think his faith is misplaced. We will later see that we can do justice to the distinction between well-placed faith and misplaced faith.

My final preliminary observation is that having faith seems to involve going *beyond* the evidence in some way. The bulk of my argument will be devoted to spelling out

² Although the following worry arises from the possibility of overdetermination. I might have faith in my friend, and therefore have faith that my friend hasn't transformed me into a brain in a vat for his own merriment (this example is due to an anonymous referee), and yet I might be antecedently certain that I'm not a brain in a vat (on the basis of philosophical arguments, perhaps). Or I might have complete faith in a friend's testimony, and thus have faith in anything he says; however, it might be that he sometimes says things of which I am already certain. What should we say in these cases? One possible response is to say that a friend's testimony simply can't produce faith in propositions of which I am already certain. After all, we may think it sounds strange to say *I am independently convinced that I can't be a brain in a vat, and I also have faith that you haven't envatted me*. Another possibility is to claim that these statements, to the extent that we can imagine circumstances in which they could be uttered felicitously, are really modal in character: the actor is claiming that if she wasn't independently convinced that she wasn't a brain in a vat, she would have faith that you haven't envatted her.

in what way one must go beyond the evidence in order to count as having faith. I postpone discussion of this to the next section.

We can now begin to give a formal analysis of faith. As we've seen, the term 'faith' appears in many different grammatical constructions: you might have faith in a person, you might have faith in a proposition, you might perform an act of faith, or you might act on faith. We require an account that makes sense of all of these uses of the term. I've already pointed out that faith typically involves a proposition as well as an action to which the truth or falsity of the proposition makes a difference. I propose, then, to make faith that X , expressed by A the basic unit of analysis, where X is a proposition and A is an act, and define the other constructions in terms of this one.

It is important that our analysis expresses the relationship between the proposition and the act. I have explained that a person can have faith that X only if he cares whether X is true or false, and presumably this is because the person would like to perform the act if X is true but would like to do some other act if X is false. So, as a first pass, we might say:

A person has faith that X , expressed by A , only if that person performs act A when there is some alternative act B such that he strictly prefers $A \& X$ to $B \& X$ and he strictly prefers $B \& \sim X$ to $A \& \sim X$.

Thus, we might say that the agent bases A on the truth of X , since there is some alternative action that the agent would perform if he knew that X were false.

This is not yet the whole story, but it does allow us to go a step further and identify what it is to have faith in a person, and to perform an act of faith:

A person P has faith in another person Q if and only if there is some act A and some proposition(s) X that express(es) a positive judgement about Q such that P has faith in X , expressed by A .³

So, Bob might have faith in Mary because he has faith that Mary won't reveal a secret he tells her, expressed by the act of telling her his secret. Paul might have faith in God because he has faith that God exists and that God is good, expressed by the act of praying. Again, faith only requires a disposition to choose particular acts, and these acts need not be actually available.

We can now take the next step and identify what it is to perform an act of faith, or to act on faith:

A person performs an act of faith (or acts on faith) if and only if he performs some act A such that there is a proposition X in which he has faith, expressed by A .

³ Notice that the judgement must be positive from the point of view of the agent, in the sense that the agent has a preference for $A \& X$, otherwise the account would be subject to the following counterexample: we think that if a person prefers that his friend refrain from smoking, even though he thinks his friend is inclined to smoke, he can't appropriately be said to have faith that she will smoke.

With these preliminaries in place, the rest of this paper will elaborate what else faith that X , expressed by A , requires and under what circumstances it is rational to have such faith.

III. Going Beyond the Evidence: Three Views

Before outlining my own view I consider three initially promising ways to make sense of the requirement that faith goes beyond the evidence. I conclude that each of these attempts fails to reveal a genuine requirement of faith.⁴

The first analysis claims that faith in X requires believing X to a higher degree than one thinks the evidence warrants.⁵ More precisely, for an agent to have faith in X , he must think that the evidence warrants believing X to some degree, say, r , but nonetheless believe X to degree q , where $q > r$. As a special case of this, one might think that faith requires believing X to degree 1, even though one thinks that the evidence warrants a definite credence less than this.

On this analysis, having faith involves being entrenched in a kind of partial belief version of Moore's paradox: one thinks something like X is likely to degree r , but I believe X to degree q . Ignoring the issue of whether this could ever be rational—since we don't want to prejudice the main issue by assuming there must be cases of rational faith—there are two problems with taking this to be a requirement of faith. First, it seems hard to imagine someone actually having faith in this sense and, especially, recognizing that he has faith in this sense. Yet having faith seems to be a common occurrence, one that does not involve psychological tricks or self-deception. Second, because it is unclear that one can reliably or stably have faith in this sense, or even take steps to set oneself up to have faith in this sense, it does not seem to be the kind of thing that ethics would require. And yet, religious ethics and the ethics of friendship do seem to require faith in certain cases.

The second analysis is more initially plausible. According to this, for the person who thinks that the evidence warrants believing X to degree r , faith requires acting as if he has degree of belief q —that is, performing the actions that he would perform if he had degree of belief q —where $q > r$. Thus, one can maintain one's degree of belief in r —and so avoid epistemic inconsistency—while still behaving, as regards the likelihood of X , in a way that goes beyond the evidence. Indeed, the paradox sounds less paradoxical when degree of belief is cashed out in terms of betting behaviour. One thinks, the evidence warrants my adopting a betting quotient of r , but I adopt a betting quotient of q . Again, as a special case of this, we might think that faith requires adopting a betting quotient of 1, that is, using $p(X) = 1$ when making decisions. This would

⁴ I presuppose a standard 'partial belief' framework, where beliefs come in degrees (sometimes called 'credences') between 0 and 1.

⁵ In some of the analyses under review here the act does not figure into the proposal. Therefore, for readability, I will say 'faith in X ' when I really mean 'faith that X expressed by A , for some particular A '.

involve not considering or caring about states of the world in which $\sim X$ holds when making decisions.

I admit that this analysis has some plausibility. However, I again think that there are problems concerning both the phenomenology and ethics of faith.

The phenomenological worry is that, on this analysis, faith requires simultaneously keeping track of two things: one's actual credences, and the 'faith-adjusted' credences that one employs in decision making. However, the phenomenology of faith doesn't seem to involve a lot of mental accounting. Yet perhaps this is not a serious problem, because the defender of this view could argue that since faith is relative to particular acts, one only needs to consider one's faith-adjusted credences when making the relevant decision.

There are two more serious problems arising from the fact that although religious ethics and the ethics of friendship endorse faith in many situations, they wouldn't endorse certain demands that this analysis suggests. First, consider what this analysis recommends that a faithful person do when asked whether he believes that X . Since this is an action, presumably the faithful person ought, if he ought to have faith, to figure out what to do using his faith-adjusted credences, not his actual ones. So he ought to claim to believe X more strongly than he does; that is, he ought to lie. But those that endorse faith often strongly denounce lying.

The second problem is brought to the fore when we consider the special case view that faith requires acting as if $p(X) = 1$. On this view, the faithful person ought to take any bet that is favourable on the condition that X obtains, regardless of the stakes. So, if asked to bet \$1m on a gamble that pays 1 penny if X obtains, the person with faith in X ought to say yes: after all, he can disregard the possibility of $\sim X$ for the purposes of decision making. I'm extremely doubtful that religious ethics would endorse the claim that the truly faithful ought to risk \$1m for a mere penny if God exists, especially since they recognize that the evidence isn't conclusive.⁶

So we can dismiss the first two analyses which held that in order for a person to count as having faith in X , he must treat his credence in X as higher than it in fact is, either by actually raising it or by acting as if it were higher. Perhaps these analyses have gone astray because they take for granted an inadequate account of when faith enters into one's belief formation process. They both assume that one examines all of the evidence dispassionately, forms a belief, and then decides whether to adjust this belief in light of faith. But perhaps the relationship between faith and belief formation is more complex than this assumption recognizes. Instead, faith might require taking evidence into account in a particular way—a way that favours X or gives the truth of X the benefit of the doubt, so to speak. Following this line of thought, a third

⁶ Perhaps the defender of the second analysis could claim that his view doesn't entail an affirmative answer because betting itself, when the payoffs are so frivolous, has inherent disutility. But it is not clear that he can respond in this way to the case in which the 'payoffs' are goods of real value, for example, in the gamble that results in a million lives lost if God doesn't exist and one life improved mildly, if God does exist.

analysis of faith holds that faith requires setting one's prior probability to $p(X) = 1$ before examining the evidence. On this view, one interprets evidence, not with an eye towards finding out whether or not X holds, but in light of the assumption that X does hold. On this view, we might say that faith goes *before* the evidence, not beyond it.

Note that this third analysis is different from the 'special cases' of the first two analyses. On those analyses, the faithful agent sets $p(X) = 1$ even though he believes that the evidence warrants something less. On the present analysis the agent doesn't have an opinion about what the evidence warrants that is separate from the question of whether he has faith. So let's say that I have faith that my friend won't reveal a secret I told him, and I overhear a third party complaining that my friend is a gossip. On the first analysis, I consider this to be evidence against the claim that my friend will keep my secret, but I nonetheless ignore it and continue to have a high degree of belief in the claim. On the second analysis, I consider this to be evidence against the claim, and I lower my degree of belief in the claim, but I nonetheless continue to act as if I have a high degree of belief (I carry on as if no one knows my secret, and I continue to confide in this friend). On the analysis we are now considering, I don't consider this to be evidence against the claim, precisely because I have faith in the claim. Indeed, there will be no possible evidence that tells against X .

This third analysis has a number of advantages. For one, it sheds light on the fact that there seems to be no good answer to the question of how a rational person ought to set his priors. On this view, the reason that there is no good answer is that epistemic rationality stops just short of this question and faith takes over: one can't avoid having faith in something, because one can't avoid setting one's priors.⁷ This supports William James' claim that one's non-rational or 'passional' nature must play a role in what to believe when reason alone doesn't dictate an answer, and that the passional nature generally comes into play in figuring out how to interpret evidence (see James 1896). It also supports an argument of Søren Kierkegaard's pseudonymous Johannes Climacus that reason alone cannot produce faith; instead, faith requires an act of will (see Kierkegaard 1846). Roughly, Climacus argues that one can never get to religious faith by engaging in objective inquiry because religious faith requires total commitment to particular historical claims. Objective inquiry can never yield certainty in these matters: it always leaves room for doubt. On Robert Adams's interpretation of Kierkegaard, total commitment to a belief requires a commitment not to revise it in the future (see Adams 1976). Thus, it requires setting $p(X) = 1$ and interpreting any new evidence in light of this.

However, despite its attractiveness, this view is incorrect because it is vulnerable to similar phenomenological and ethical objections to the ones discussed above. Adams himself raises the ethical objection: 'It has commonly been thought to be an important

⁷ Technically, one could avoid having complete faith in anything, since one could avoid setting $p(X) = 1$ for any X . However, if we think that degrees of faith correspond to setting lower priors, then one would have some degree of faith in many things.

part of religious ethics that one ought to be humble, teachable, open to correction, new inspiration, and growth of insight, even (and perhaps especially) in important religious beliefs' (Adams 1976: 233). We might add that the ethics involved in friendship similarly does not seem to require that we remain determined not to abandon our belief in a friend's trustworthiness come what may.

The phenomenological objection can be brought out by considering that anyone who is acting on faith typically feels like she is taking a risk of some sort. The act *A* that you are performing on faith (that *X*) is supposed to be better than some alternative if *X* holds and worse than that alternative if *X* does not hold. But if one is certain that *X* is true, then doing *A* is not a risk at all! On the contrary, *A* is simply an act that, from your point of view, will undoubtedly turn out well. It is like the act of taking a bet on which you win \$100 if water is H₂O and lose \$100 if it is not. One might reply that from an objective standpoint, doing *A* is a risk—because setting one's priors is risky in some sense. But even if that is the case, the view still fails to explain the phenomenology of acts of faith, since they feel risky even from an internal perspective. What is distinctive about taking a leap of faith, so to speak, is that you are fully aware that it might turn out badly—even if you think it unlikely that it will.

An additional objection to this third analysis is that it cannot distinguish between cases of well-placed faith and cases of misplaced faith. Recall the above example of the person who knows that her friend has spilled all of her secrets so far but who has faith that he will not spill future secrets. We likely regard this as a case of misplaced faith. At any rate, when we compare this person to the person whose friend has never spilled a secret and who has faith that he will not spill future secrets, we think that this second person's faith has a lot more in its favour. But we cannot make sense of this on the present view, since rationality has no conclusions about which priors are laudable and since faith enters the picture before any evidence is interpreted.

So, although this third analysis was initially promising, it does not ultimately succeed. On my view, whether someone has faith is not determinable from his priors: a person who starts out sceptical, but who then amasses evidence in favour of *X*, could indeed end up choosing to have faith in *X* (consider the conversion of St Paul). Furthermore, a person who begins by assuming that *X* must be true doesn't thereby count as having faith in *X*: credulity and faith come apart. So do credence and faith, as we will see in the next section.

IV. Faith and Examining Further Evidence

There is something to Kierkegaard's idea that one can never arrive at faith by engaging in empirical inquiry—that faith instead requires an act of will. However, this is not because faith requires a kind of certainty that empirical inquiry cannot provide, nor because faith must precede inquiry. Instead, it is because engaging in an inquiry itself constitutes a lack of faith. That is, faith requires not engaging in an inquiry whose only

purpose is to figure out the truth of the proposition one purportedly has faith in. So the sense in which faith in *X* requires some response to the evidence aside from that normally warranted by epistemic norms is that it requires a decision to stop searching for additional evidence and to perform the act one would perform on the supposition that *X*.

Consider an example. If a man has faith that his spouse isn't cheating, this seems to rule out his hiring a private investigator, opening her mail, or even striking up a conversation with her boss to check that she really was working late last night—that is, it rules out conducting an inquiry to verify that his spouse isn't cheating. If he does any of these things, then she can rightfully complain that he didn't have faith in her, *even if* she realizes that, given his evidence, he should not assign degree of belief 1 to her constancy. Similarly, if I have faith that my friend will keep a secret, this seems to rule out asking a third party whether he thinks that friend is trustworthy. To use a religious example, when so-called 'doubting' Thomas asks to put his hand in Jesus' side to verify that he has been resurrected in the flesh, this is supposed to indicate that he lacks faith.

We can say something even stronger: faith seems to require not looking for further evidence even if one knows that the evidence is readily available. For example, consider a case in which a man simply stumbles across an envelope which he knows contains evidence that will either vindicate his wife's constancy or suggest that she has been cheating. He seems to display a lack of faith in her constancy if he opens it and to display faith in her constancy if he does not. And this seems true even if the evidence has been acquired in a scrupulous way: we might imagine the wife herself presents the envelope to the man, as a test of his faith.⁸

So we now have the following first pass at a full analysis of faith:

A person has faith that *X*, expressed by *A*, if and only if that person performs act *A* when there is some alternative act *B* such that he strictly prefers *A*&&*X* to *B*&&*X* and he strictly prefers *B*&& \sim *X* to *A*&& \sim *X*, and the person refrains from gathering further evidence to determine the truth or falsity of *X*, or would refrain, if further evidence were available.

To make this more precise, we might state this final condition in preference terms:

... and the person prefers to decline evidence rather than to view it.

This formulation has an unfortunate upshot, though: it implies that anyone who has faith that *X*, expressed by some act *A*, must decline evidence in the matter of *X* even if they want the evidence for purposes other than deciding between *A* and *B*. For example, consider the Christian apologist who has faith that Jesus was resurrected (expressed by, say, the action of going to church every week) but who combs through the historical evidence surrounding Jesus' resurrection in the hopes of finding evidence

⁸ Indeed, my account can easily explain why presenting him with the envelope could be a test of faith: it is a test to see whether he will choose to acquire further evidence.

to convince someone who does not believe. Or consider the person who intends to open the private investigator's envelope publicly, precisely to show that he has faith in his spouse's constancy.⁹ On the current analysis, neither of these acts can be acts of faith: indeed, performing them entails that the agent does not have faith in the proposition in question.

The reason that we would say that the apologist has faith in the resurrection even though he continues to look for evidence is that he doesn't consider his decision to attend church dependent on the outcome of his investigation. Indeed, if he had no desire to convince other people, he would not look for evidence. Similarly, the reason we know that the husband has faith in his spouse, expressed by, say, the act of remaining constant himself, is that his constancy doesn't depend on the contents of the envelope, even though it does depend on his (current) beliefs about whether his spouse is cheating. So what these examples show is that the claim that the faithful person does not look for evidence at all is too strong. Instead, the faithful person does not look for evidence *for the purposes of deciding whether to do A*. Thus, if he does look for evidence, he considers this search irrelevant to his decision to do *A*.

A precise way to spell out that the act doesn't depend on the evidence is that the faithful agent is willing to commit to *A* before viewing any additional evidence in the matter of *X*; indeed, he wants to commit to *A*. In preference terms, he prefers to commit to *A* before viewing any additional evidence rather than to first view additional evidence and then decide whether to do *A*. This covers both the case of the person who looks for no additional evidence and the person who does look for evidence, but not in order to decide whether to do *A*.

We can now formulate my final analysis:

A person has faith that *X*, expressed by *A*, if and only if that person performs act *A* when there is some alternative act *B* such that he strictly prefers *A*&*X* to *B*&*X* and he strictly prefers *B*& \sim *X* to *A*& \sim *X*, and the person prefers {to commit to *A* before he examines additional evidence} rather than {to postpone his decision about *A* until he examines additional evidence}.

I purposely leave it ambiguous whether this second preference should be strict or weak, as I think a case can be made for either view, and I return to this issue in Section VII. As mentioned above, my analysis vindicates part of Kierkegaard's insight that faith does require total commitment, and that looking for evidence reveals that one is not totally committed. But what one must commit to is an *act*, not a belief: specifically, one must commit to performing an act regardless of what the evidence reveals. My analysis also vindicates the idea that faith requires an act of will—on my account one consciously chooses not to look for more evidence (even though doing so might be tempting!)—which is difficult to explain if faith requires a certain degree of belief and belief is not directly under one's volitional control.

One upshot of my analysis is that it is possible for two people to have the same evidence, the same credence function, and the same utility function, and perform

⁹ Thanks to Sherrilyn Roush for this example.

the same act, and yet one of these acts displays faith while the other doesn't. So, for example, assume Ann and Erin have the same evidence about Dan's secret-keeping ability; that both have $p(\text{Dan will keep a secret}) = 0.9$; that both have the same utility function (that is, the stakes are the same for both of them). Now assume that each has a choice whether to ask a third party what he thinks about Dan's secret-keeping ability before deciding whether to tell Dan her secret. Ann decides to simply tell her secret; Erin decides to ask the third party, and then ends up telling her secret to Dan on the advice of this third party. Here, Ann displays faith that Dan will keep a secret (expressed by the act of revealing her own secret), whereas Erin does not display faith, even though she also performs this act. So the same act in the same circumstances can be done with or without faith.

The argument so far has told us nothing about the circumstances, if indeed there are any, in which faith can be rational. I now turn to this question. First I briefly explain the distinction between epistemic rationality and practical rationality, beginning with the former.

V. Epistemic and Practical Rationality

I will assume a broadly evidentialist conception of epistemic rationality: one should proportion one's beliefs to one's evidence. One should not, for example, simply believe what one likes or believe what would make one happy. More generally, one should not take non-truth-conducive reasons as reasons for belief. I will also assume a subjective Bayesian account of partial belief: degrees of belief obey the probability calculus; one updates one's beliefs by conditionalizing on new evidence; and two people can (rationally) have different degrees of belief in a proposition if and only if they have different evidence that bears on that proposition, or they believe the same evidence bears on that proposition differently, or they have different priors.¹⁰ An important feature of this account for our purposes is that a rational person can only change his credences in response to evidence—and, in fact, *must* update them in response to new evidence, at least in matters he cares enough about to form beliefs.

To count as epistemically rational, you must proportion your beliefs to the evidence and your beliefs must be coherent (in other words, your credences must obey the probability calculus). Practical rationality, on the other hand, concerns selecting the means to achieve one's ends.

Practical rationality is the kind of rationality that decision theory is meant to formalize: the values of the ends are made precise by means of a utility function,

¹⁰ Depending on how one defines evidentialism, it might be strictly speaking incompatible with subjective Bayesianism, since the latter allows any set of prior degrees of belief to count as rational, so that two people could share a body of evidence and still (rationally) have different degrees of belief. On the other hand, a natural version of evidentialism implies that there is only one rational way to respond to a given body of evidence. We could alleviate this problem by pointing out that rationality is a feature of belief *change*, or that evidentialism and subjective Bayesianism agree about what to do once the priors have been set.

and the likelihoods of any particular acts achieving those ends are made precise by means of a credence function.¹¹ On the standard view, subjective expected utility (EU) theory, rational agents maximize expected utility: they prefer the act with the highest mathematical expectation of utility, relative to their utility and credence functions. So if we think of an act as a gamble that yields a particular outcome in a particular state of the world—for example, $g = \{O_1, E_1; O_2, E_2; \dots; O_n, E_n\}$ is the act that yields O_i if E_i is true, for each i —then the value of this act is:

$$EU(g) = \sum_{i=1}^n p(E_i)u(O_i)$$

In my view, EU theory is too restrictive. However, since it is the widely accepted view, and since I agree that expected utility maximizers are practically rational, I will postpone discussion of my alternative view to Section IX.

Before turning to the question of whether faith can be rational, it is worth clearing up a worry: that the definitions of rationality that I've adopted might not be strong enough. On the present definition of epistemic rationality, one may adopt any prior degrees of belief, including any conditional priors: priors about the relationship between particular hypotheses and particular pieces of evidence. For example, a person may rationally believe he has been abducted by aliens, as long as he also believes that the evidence he has supports this to the degree that he believes it. A similar point holds about preferences in the case of practical rationality. Both epistemic rationality and practical rationality, as I define them here, are notions of *consistency*: the only restriction on degrees of belief is that they are consistent with one another, and the only restriction on preferences is that they are consistent with one another, given one's degrees of belief in each possible state of the world. However, there is another notion of rationality, which rules out believing one has been abducted by aliens and rules out certain preferences, which we might call *reasonableness*.

I cannot fully respond to this worry in depth here. But it is important that the consistency notion of rationality and the reasonableness notion come apart quite readily. Consistency restrictions are structural: they rule out particular patterns of belief or desire, regardless of the content of these attitudes. On the other hand, reasonableness restrictions are substantial: they rule out particular beliefs or desires, regardless of which other beliefs or desires one has. Therefore, we can talk about what they require separately, and this project is about the requirements of rationality in the consistency sense. Or, since this sense of rationality exhausts the subjective sense of rationality, this project addresses the question of whether it can ever be rational *from an agent's own point of view* to have faith. If we want to answer the further question of whether it is reasonable to have faith—that is, of whether a person has objective reasons to have faith—we can address this separately. Indeed, nothing in my argument relies on the

¹¹ On the debate between realism and functionalism about the utility function, see J. Dreier (1996).

content of the propositions believed or desired. Hence the question of whether faith is reasonable can be answered by asking whether there are any contents for which it is reasonable to have the patterns of belief and desire presented in my examples.

VI. Practical Rationality and Evidence-gathering

It should be clear that on the first two analyses of faith, considered in Section III above, faith is irrational. On the first analysis, which requires the agent's credences to be higher than those the evidence suggests, faith is epistemically irrational, and on the second analysis, which requires the agent to act as if his beliefs are different than they are (that is, to take something other than his credences as 'what he believes' for the purposes of decision making), faith is practically irrational, however we spell out practical rationality. On the third analysis, faith is always rational, provided one has consistent credences and preferences (though one can't, for example, have faith in two contradictory propositions).

On my analysis faith can also be epistemically rational: that one has faith in X implies nothing about one's degrees of belief or the consistency thereof.¹² Therefore, one can clearly meet the requirements of epistemic rationality, as I've stated them, while having faith—whether one has faith is completely separate from whether one is epistemically rational.¹³

But can faith also be practically rational? Faith requires two preferences: a strict preference for A over the alternatives and a preference not to look for further evidence for the purposes of this decision (or to commit to A rather than seeing further evidence before deciding). Therefore, faith can only be practically rational if both of these preferences are practically rational. Assessing the rationality of the first preference is fairly straightforward: strictly preferring A to the alternatives is rational just in case A has a higher expected utility than the alternatives, given the agent's credences. So the first (fairly obvious) restriction on when faith is rational is this: one's credence in X must be sufficiently high as to make A the practically rational act.

What about the second preference? Can it be rational to prefer committing to A rather than seeing more evidence before deciding? To address this question, it will help to have a canonical example of the kinds of situations in which the question of faith arises. Again, these situations involve an agent performing an act A in a situation in which the status of some proposition X is in question, and in which he prefers $A \& X$ to $B \& X$ and prefers $B \& \sim X$ to $A \& \sim X$ for some alternative act B . Let us simplify by assuming that there are only two alternatives, and that the alternative act is such that

¹² Technically, it implies that $p(X) \neq 1$, but if $p(X) = 1$ then X is not an appropriate object of faith, so this is not a restriction on one's credences, but on what can count as an object of faith.

¹³ Even if epistemic rationality requires that one look for further evidence, it will not conflict with faith on my definition, since having faith doesn't forbid one from looking for evidence; it only dictates that one must prefer committing to an act before seeing more evidence to postponing the decision.

its value does not depend on X , that is, the agent is indifferent between $B\&X$ and $B\&\sim X$.¹⁴ To put some concrete utility values in place:

$$\begin{aligned} u(A\&X) &= 10 & u(B\&X) &= 1 \\ u(A\&\sim X) &= 0 & u(B\&\sim X) &= 1 \end{aligned}$$

To generalize, let us assume there is a high value, a middle value, and a low value such that $u(A\&X) = H$, $u(B\&X) = u(B\&\sim X) = M$, $u(A\&\sim X) = L$, and $H > M > L$.

Here are some examples of decisions that might include the relevant values. Consider an agent who is deciding whether to become a monk and does not have conclusive evidence that God exists. If God exists, then becoming a monk is very good—the agent will experience all the goods of the religious life. But if God does not exist, then becoming a monk will result in the agent living a life that is ultimately wasted. On the other hand, failing to become a monk is fine, but not great, either way, if we assume that from a religious point of view becoming a monk is supererogatory: the agent lives roughly the same life as a non-monk whether or not God exists. To take another example, consider an agent who is deciding whether to use his van to transport 10 critically injured patients to the hospital rather than using his car to transport 1 and who does not have conclusive evidence that his van works (but he has, say, near-conclusive evidence that his car works, or a backup plan in case it doesn't). Or consider an agent who is deciding whether to reveal a secret to someone else, and who does not have conclusive evidence that the friend will keep it; or an agent who is deciding whether to marry a particular person, and does not have conclusive evidence that this person will make a good spouse. In each case, performing the act could turn out very well or poorly, whereas not performing the act is the same either way.¹⁵ We could think of the act as an opportunity for something great—but a risky opportunity—that one might take or pass up.

¹⁴ What if X is something that the agent strongly prefers to be true, so that $A\&X$ and $B\&X$ both have a high utility? One might think, for example, that faith in God is typically like this, since the agent often prefers worlds in which God exists to worlds in which God does not exist, regardless of the agent's choices. In fact, we could build this into the example without changing any of the results. In our discussion, there are only two options and two possible outcomes for each option, and we are only concerned with which of the two options has a higher EU and by how much, that is, with the difference between $EU(A)$ and $EU(B)$. We know that $EU(A) - EU(B) = p(X)(u(A\&X) - u(B\&X)) + p(\sim X)(u(A\&\sim X) - u(B\&\sim X))$. Therefore, the only thing that matters is the difference between $A\&X$ and $B\&X$ on the one hand, and between $A\&\sim X$ and $B\&\sim X$ on the other hand. So if we uniformly increase the value of the outcomes in which X holds, say by replacing the above values with $u(A\&X) = 110$, $u(A\&\sim X) = 0$, $u(B\&X) = 101$, $u(B\&\sim X) = 1$, to symbolize that the truth of X has inherent value to the agent, this equation will have the same value, and consequently, all of the results of sections VI–VIII will still hold.

¹⁵ One might argue that the act couldn't be the same either way, since, e.g. passing up the possibility of a faithful spouse (or an opportunity to save 10 lives) is worse than passing up the possibility of an unfaithful spouse (or an opportunity not to save any lives). I think the question of how decision theory should handle these nuances is an interesting one, but for now, I will just assume that facts about what might have happened do not make a difference to the agent's utility function—or at least that they make a negligible difference.

Let us assume that A has a higher EU than B , and so A is practically rational given the agent's current information. Now we want to know whether practical rationality requires that the agent gather more information before she makes her decision. There is a theorem that bears directly on this question: I. J. Good (1967) showed that gathering additional evidence (in Good's terminology, making a new observation) and then using it to update one's beliefs before making a decision always has a higher expected utility than making a decision without doing so, provided the following two conditions are met:

- 1) It is not the case that the agent will perform the same act regardless of the evidence she receives.
- 2) Gathering additional evidence is cost-free.

If condition (2) holds but, contrary to (1), the agent will perform the same act regardless of the evidence, then gathering the evidence will have the *same* expected utility as not doing so.

It is helpful to say something about Good's setup and about how he calculates the expected utility of gathering and using additional evidence and that of not doing so. The expected utility of not gathering additional evidence is simply the EU of the act with the highest expected utility: in our case, $EU(A)$. (So, for example, if $p(X) = 0.59$, then $EU(\text{don't gather}) = 5.9$.) The expected utility of gathering the evidence is more complicated. We consider, for each piece of evidence the agent might receive, which act will have the highest EU relative to the agent's new credence function after receiving that piece of evidence. The EU of this act is the agent's utility upon receiving that evidence, since we assume that the agent will pick the act that maximizes expected utility. (So, for example, the EU upon receiving a bit of evidence E that will result in $p(X) = 0.95$ will be $EU_{\text{new}}(A | E) = 9.5$, and the EU upon receiving a bit of evidence $\sim E$ that will result in $p(X) = 0.05$ will be $EU_{\text{new}}(B | \sim E) = 1$.) We then need only to weight each of these values—the value of each piece of evidence, so to speak—by the probability of receiving that piece of evidence, to determine the expected utility of gathering additional evidence. (So, for example, if $p(E) = 0.6$ and $p(\sim E) = 0.4$, then $EU(\text{gather}) = 6.1$.) I note that, throughout, Good assumes that gathering evidence itself has no utility costs.

Formally, where $O_i(A_j)$ is the outcome of action A_j in state S_i and $p(S_i | E_k)$ is the agent's conditional credence for S_i given E_k :

$$\begin{aligned} EU(\text{don't gather}) &= \max_j \sum_i p(S_i) u(O_i(A_j)) \\ EU(\text{gather}) &= \sum_k (p(E_k) \max_j \sum_i p(S_i | E_k) u(O_i(A_j))) \end{aligned}$$

In our scenario, with two possible actions whose utility values depend only on the status of X , we have:

$$EU(\text{don't gather}) = p(X)H + p(\neg X)L$$

$$EU(\text{gather}) = \sum_k p(E_k) \max\{M, (p(X|E_k)H + p(\neg X|E_k)L)\}$$

Good proves that unless the agent will do the same act regardless of the result of the experiment—that is, unless the same act maximizes expected utility for each possible piece of evidence E_k —then $EU(\text{gather})$ is always higher than $EU(\text{don't gather})$. If the agent will do the same act regardless of the result, then these values are the same. In other words, it is always rationally permissible to make a new observation and use it, and it is rationally *required* that one do so if some piece of evidence that might result from doing so will lead one to do B instead of A .

VII. Commitment and Interpersonal Cost

We can now consider under which circumstances, if any, one is rationally permitted, or rationally required, to have faith. In this section and the next I spell out when it is rational to refrain from examining additional evidence. I assume throughout that the other conditions for the rationality of faith have been met: that the agent's credences are coherent and are such that doing A rather than B maximizes expected utility. I also assume, as Good does, that for each 'experiment' (or bit of searching) the agent is considering performing, the agent can assign credences to the possible results E_i and conditional credences $p(X | E_i)$ to the hypothesis given each possible result; and that if the agent performs the experiment, he will update his credence in X on the result that he obtains and choose the action that maximizes expected utility on his updated credence.

As we saw in the previous section, one case in which it would be rationally permissible not to examine additional evidence is the case in which the agent will do the same thing no matter which result obtains. According to expected utility theory, doing so will be rational just in case the conditional probabilities $p(X | E_i)$ are all such that A maximizes expected utility after updating on each result. In our example, this will be the case in which $p(X | E_i) > 0.1$ for all E_i , or, more generally, the case in which $p(X | E_i) > (M - L)/(H - L)$ for all E_i . In less technical terms, this will hold when none of the observations the agent is considering would tell against X conclusively enough to dissuade the agent from A .

However, in these cases the agent will never *strictly* prefer not examining the evidence; instead, he will be indifferent to whether he examines the evidence or not (provided, again, that doing so is cost-free). So if we think that faith requires a strict preference against examining additional evidence—and as I mentioned in Section IV, I am not sure which stance to take—then pointing to situations in which the agent will do the same regardless of the evidence will not help us.

The other condition under which it would be rational to eschew additional evidence is the situation in which the evidence is costly. Since one of the conditions of Good's Theorem is that costs are negligible, he doesn't explain how to proceed when there are

costs. But there is a fairly obvious way to determine whether the benefit of performing the experiment outweighs the cost, in expected utility terms. Specifically, we can measure the benefits of the experiment as the difference between $EU(\text{gather})$ and $EU(\text{don't gather})$, where S_k is a relation defined to hold of all evidential results E_k such that the agent will want to choose A if he sees that result:

$$EU(\text{gather}) - EU(\text{don't gather}) = \sum_{k \in S_k} p(E_k)(M - (p(X|E_k)H + p(\neg X|E_k)L))$$

The multiplicand in each term of the summation is the EU of performing B minus the EU of performing A , relative to the new evidence E_k . If the experiment has some monetary or cognitive cost that is measurable in utility and that does not depend on which action the agent chooses, the result of the experiment, or the state of the world, then it is a simple matter to determine whether the benefit outweighs the cost.¹⁶

However, the costs of gathering evidence needn't be monetary or cognitive. Indeed, there are two types of (non-monetary) costs that seem particularly relevant to situations involving faith: interpersonal costs and the costs of postponing a decision. We need to examine each of these, to see whether they could make faith rational.

The first kind of cost comes up primarily in contexts in which one has faith in another person. In these cases, lacking faith might in itself cause harm to the relationship. For example, one's spouse or friend might be upset if one doesn't have faith in her, or one might miss out on certain goods that mutual faith is a prerequisite for (a feeling of connection or security, perhaps). In the religious case, it might be that one's relationship with God will be lacking if one does not have faith and will be lacking *for that fact*. We might think of these costs as intrinsic costs to lacking faith.

Before concluding that these costs can play a role in rationalizing faith, though, we should consider *why* it should be upsetting to someone that another person lacks faith in them. I suspect that the most common reason is that a lack of faith *indicates* that the agent is not as he ought to be with respect to his beliefs and desires regarding the other person. For example, a husband assigns low credence to the possibility that his wife is faithful even though she's given him evidence through her actions and character that should be sufficient for a high credence, or he assigns a low utility to continuing the marriage with her. But these aren't complaints against the husband's lack of faith per se. In one case, it is a complaint that he is being irrational in a particularly hurtful way. In the other, it is a complaint that he has the wrong values given their relationship. So there is no 'additional' cost to lacking faith, beyond the costs of these actions.

I tentatively conclude that the intrinsic cost of lacking faith *might* be a way in which evidence could be costly for a rational agent, but that this position would need to be further supported.

¹⁶ Since we can only calculate the utility of fully specified outcomes, the cost will technically be a utility difference c such that $u(A \& X \& \text{don't gather}) - u(A \& X \& \text{gather}) = u(B \& X \& \text{don't gather}) - u(B \& X \& \text{gather}) + c$. $u(A \& \sim X \& \text{don't gather}) - u(A \& \sim X \& \text{gather}) = u(B \& \sim X \& \text{don't gather}) - u(B \& \sim X \& \text{gather}) = c$.

VIII. The Cost of Postponement

The second kind of cost that might be associated with gathering evidence is the cost of postponing the decision. In the most extreme case, looking for further evidence amounts to losing the option of doing *A*. For example, it might be that one's friend is only available to hear one's secret today, so if one does not reveal it, one will lose that option forever. Or it might be that one's potential spouse has given one an ultimatum. Or it might be that one needs to choose a vehicle to drive critically injured patients to the hospital, and any delay will result in their certain death. In these cases, the overall value of gathering further evidence will be negative: it will be the difference between the expected value (on one's current credences) of doing *A* and that of doing *B*, since deciding to gather further evidence is equivalent to deciding to do *B*.

In a less extreme version, *A* might be an action that provides more good to the decision-maker the earlier he chooses it (in the event that *X* holds), so the utility of choosing *A* tomorrow might be slightly lower than the utility of choosing *A* today. If we imagine that the agent always prefers a day of marriage with a faithful spouse to a day of bachelorhood, or that in the event that God does exist, the agent prefers a day spent as a monk to a day spent as an ordinary citizen, then each day of postponing the decision is costly.

For decisions where postponement is costly but does not prevent the agent from eventually choosing *A*, under what conditions does this cost outweigh the benefit of gathering additional evidence? To answer this question, let us assume that the only cost associated with postponing the decision occurs in the event that one eventually does *A* and that *X* obtains. Then let us fix the cost of postponing the decision while one does a particular experiment as *c*: specifically, *c* is the difference between $u(A \& X \& \text{don't gather})$ and $u(A \& X \& \text{gather})$, that is, the difference between doing *A* when *X* obtains without gathering more evidence (or while committing to *A* regardless of the evidence) and doing *A* when *X* obtains after gathering more evidence. Then the benefit of performing the experiment—according to the agent's current credences—is:

$$\text{EU}(\text{gather with cost}) - \text{EU}(\text{don't gather}) = \sum_k p(E_k)(M - (p(X|E_k)H + p(\neg X|E_k)L)) - c \sum_k p(E_k)p(X|E_k)$$

Here, C_k is defined to hold when, if the agent were to receive evidential result E_k , he would still want to choose *A*, even in light of the additional cost of doing *A* if *X*. Note that the first term will always be positive, so the magnitude of the second term must be sufficiently high to offset it in order for the value of performing the experiment to be negative.

We can say something about when the value of performing the experiment is apt to be negative. It will be lower when *c* is higher; that is, when the costs of doing the experiment in the circumstances in which it is costly are higher. It will also be

lower when *M* is lower, and *H* and *L* are higher; that is, when there is less of a risk involved in doing *A*. Next, it will be lower when $\sum_k p(E_k)p(X|E_k)$ is high. This

corresponds to the possibility of the agent doing *B* as a result of the experiment, even though *X* in fact holds: the agent might, for example, get 'misleading' evidence that leads to the rational performance of an action that in fact has a lower payoff than its alternative. Finally, it will be lower when $\sum_k p(E_k)p(X|E_k)$ is high. This corresponds to the probability of the agent doing *A* when *X* holds, and thus having to pay a cost he otherwise would not have had to pay. In conjunction with the previous fact, the experiment will have a lower value when $p(X)$ is antecedently high.

Thus, holding fixed *H*, *M*, and *L*, in situations in which gathering the evidence proves costly in the event that *X* obtains and the agent does *A*, refraining from gathering further evidence is more likely to be rationally required (1) when this cost is high; (2) when the experiment is likely to result in *misleading* evidence against *X*, that is, evidence that makes one 'miss out' on the possibility of doing *A* when *X* in fact holds; or (3) when one already has a high credence in *X*.

The fact that costs associated with postponing a decision can make faith rational vindicates an observation made by William James (1896), though he did not express it in these terms. James argued that when a decision about what to believe is momentous—in that it involves a once-in-a-lifetime opportunity, for example—then it must be made by the will, and that postponing the decision is a decision in itself. He used this observation to argue that it is rationally permissible to choose to believe in God even when one does not have conclusive evidence for God's existence. I don't think that it is rationally permissible to *believe* that God exists when one does not have conclusive evidence, if this means setting one's credences differently from what one has evidence for (though I'm not saying that this is what James is suggesting). However, I do think that it is sometimes rationally permissible (and indeed, sometimes rationally required!) to *have faith* in God—as evidenced by doing some particular religious act without looking for further evidence—in circumstances in which postponing the decision to act is costly.

The upshot of this discussion is that, if we accept EU theory as the correct theory of practical rationality, then faith can be rational—depending, of course, on one's credences and the situation in which one finds oneself. We have seen three important results in this regard. First, if we think that faith requires only a weak preference for not gathering additional evidence—that is, if you count as having faith when you are indifferent between making the decision on current evidence and postponing the decision—then faith is rationally permissible, but not rationally required, in cases in which no piece of evidence that one could potentially gather would alter the agent's decision about what to do. This will hold when no piece of evidence will tell conclusively enough against *X* such that doing *A* will no longer maximize expected utility. However, if we think that faith requires a strict preference for no

additional evidence—that is, you must strictly prefer making the decision on current evidence—then faith will not be rationally permissible in these circumstances.¹⁷

Second, faith (under both the strict and weak reading of preference) will be rationally required in circumstances in which there is an interpersonal cost to looking for more evidence; that is, in which lacking faith is intrinsically worse than having faith. However, it is unclear whether such circumstances obtain. In my opinion, the right explanation for the fact that there are relational goods one can't get unless one has faith isn't that faith is in itself valuable, but rather that there are some goods that one can't get if one is more suspicious of another person than the evidence warrants, or if one hesitates to act on a matter involving the relationship.

Third, and most crucially, faith (under both readings) is rationally required in circumstances in which the costs of delaying the decision are high enough to outweigh the benefit of additional evidence. Holding fixed the costs of delay, whether these costs outweigh the benefits, depends both on one's credence in the proposition one has faith in and on one's beliefs about the potential evidence one might encounter.

IX. Risk Aversion and the Possibility of Misleading Evidence

There are two reasons to think that our results so far are incomplete. First, one might think that faith requires more than a preference for not gathering additional evidence—or for committing to an action before the evidence comes in. It requires a preference for not gathering additional evidence *even when this evidence is cost-free*. For example, we may think that the person who examines the private investigator's envelope even when there are no 'postponement costs' lacks faith. Second, one might think that EU maximization is too strong a criterion of rationality, and that one can be practically rational without being an expected utility maximizer.

While I am sympathetic to the general aim of decision theory, and hold that EU theory is largely correct in its analysis of practical rationality, I nonetheless think that expected utility theory employs too narrow a criterion of rationality and should therefore be rejected. There are two motivations for this. The first is that EU theory incorrectly analyzes risk-averse behaviour for rational agents. An agent is risk-averse if he prefers a sure-thing sum of money to a gamble that yields that sum on average: for example, if he prefers \$50 to a coin-flip between \$0 and \$100.¹⁸ On EU theory, the only way an agent can rationally have this preference is if he values monetary

outcomes in a particular way: in our example, if the utility difference between \$50 and \$0 is greater than the utility difference between \$100 and \$50. Intuitively, though, there are two different reasons for having this preference. On the one hand, one might genuinely value increments of money in this way (one's values for money might 'diminish marginally'). On the other hand, one might think that an increment of \$50 adds the same value regardless of whether one already has \$50, but care about other properties of the gamble besides its average utility value: for example, the minimum value it might yield, the maximum, or the spread of possible outcomes. On EU theory, anyone who is risk-averse for this latter reason is thereby irrational. However, as I argue elsewhere, I think that both kinds of considerations could figure into a rational agent's preferences.¹⁹

The second motivation for rejecting EU theory is that it misses an important subjective component of practical rationality. Once we know the values an agent attaches to the relevant ends and the credences he assigns to the relevant states, we are still missing an important factor in determining what he ought to do—namely, how he weighs the prospect of a high probability of realizing some ends he values against a small probability of realizing some ends he values much more. Above, I explained that decision theory formalizes means-ends reasoning: utility corresponds to how much an agent values particular ends, while credence corresponds to the likelihood with which the agent thinks some particular means will realize one of these ends. We might think of the additional factor, which is missing from EU theory, as the factor corresponding to whichever means the agent considers most effective, not in realizing some particular end or another but in realizing his ends as a whole, given that there are many particular ends at stake. In other words, we might think of this factor as corresponding to how the agent structures the realization of his goals. This involves deciding whether to prioritize definitely ending up with something of *some* value or instead to prioritize *possibly* ending up with something of extraordinarily high value, and by how much. That is, it is up to the agent how the minimum value, the maximum value, the spread, and other 'global properties' factor into his decision.

What is at issue here is the following: when an agent is making a single decision, ought he to care only about how a decision would turn out on average if it were to be repeated, or can he place some weight on, for example, how bad the worst-case scenario is? In an earlier paper, I demonstrated that agents who conform to an alternative theory rather than EU theory—agents who are risk-averse in the second sense described above—will sometimes rationally reject cost-free evidence (see Buchak 2010). In particular, in scenarios like those outlined in Section VI above, it will be rationally required for these agents to commit to an action *A* before looking at additional evidence rather than to look at additional evidence and then decide. This will be the case in situations with the following properties: (1) the agent already has

¹⁷ Perhaps we could argue that faith is rational in these circumstances by stipulating that every experiment has some cost. However, when we consider that faith requires not just a (strict) preference for avoiding evidence in the matter of *X* when deciding whether to do *A* or *B*, but more precisely a (strict) preference for committing to a decision before seeing the evidence, we realize that we would have to stipulate that not committing before performing the experiment always has a cost, and this is less plausible.

¹⁸ For a more general definition of risk aversion, see M. Rothschild and J. Stiglitz (1970).

¹⁹ See L. Buchak, 'Risk and Rationality' (MS).

a high degree of belief in X (that is, $p(X)$ is antecedently high); and (2) evidence that lowers the agent's credence in X sufficiently to make her do B does not tell conclusively against X (for example, for an experiment with two possible outcomes, evidence E in favour of X and evidence $\sim E$ against X , $p(X | \sim E)$ is low but still significant).

In short, agents who care about global properties are concerned with a particular risk involved in looking for additional evidence: the risk of coming across evidence that makes it rational to do B even though X is true. In other words, they are concerned, and rationally so, about the risk that additional evidence will be misleading. I already mentioned that the possibility of misleading evidence will make it rational for the EU maximizer to reject costly evidence. However, on my view, the risk of misleading evidence makes faith rational even in cases in which there is no cost to looking for evidence. If one accepts my view, then faith is rational in more cases than it is on standard decision theory.

X. Conclusion

We have seen that whether faith that X , expressed by A , is rational depends on two important factors: (1) whether one has a high enough (rational) degree of belief in X , and (2) the character of the available evidence. Specifically, faith in X is rational only if the available evidence is such that no possible piece of evidence tells conclusively enough against X .

There are two interesting practical upshots of this conclusion. First, notice that in a standard class of cases, when one has a high degree of belief in a proposition, the odds of any particular experiment being such that it could drastically lower one's degree of belief decreases the larger the collection of evidence the agent already has.²⁰ So, in a rough-and-ready way, we might say that faith in X (expressed by some particular act A) is practically rational to the extent that the agent's belief in X is already based on a large amount of evidence. Second, whether faith is rational depends on the kind of world we are in. Faith will be rational to the extent that evidence isn't ordinarily very conclusive, or to the extent that our decisions usually do have postponement costs. We won't be able to vindicate the claim that faith is rational regardless of the circumstances. But we can explain why having faith is rational in certain worlds, perhaps worlds like ours. Individuals who lack faith because they insist on gathering all of the available evidence before making a decision stand to miss out on opportunities that could greatly benefit them.²¹

²⁰ As James Joyce avers, it is 'usually the case that the greater volume of data a person has for a hypothesis the more resilient her credence tends to be across a wide range of additional data' (2005: 161).

²¹ I am indebted to Matthew Benton, Jake Chandler, Joshua Hershey, John Pittard, Jada Twedt Strabbing, and two anonymous reviewers, for extensive written comments on earlier drafts of this paper. This paper has also benefited from discussions with David Baker, Brandon Fitelson, Marco Lopez, Tristram McPherson, Lydia McGrew, and Sherrilyn Roush. Finally, I am grateful for helpful comments from participants at the Conference on Formal Methods in Philosophy of Religion at the Katholieke Universiteit in Leuven,

References

- Adams, R. (1976) 'Kierkegaard's Arguments against Objective Reasoning in Religion', *The Monist*, 60, 2: 228–43.
- Buchak, L. (2010) 'Instrumental Rationality, Epistemic Rationality, and Evidence-gathering', *Philosophical Perspectives*, 24, 1: 85–120.
- 'Risk and Rationality'. Unpublished MS.
- Dreier, J. (1996) 'Rational Preference: Decision Theory as a Theory of Practical Rationality', *Theory and Decision*, 40: 249–76.
- Good, I. J. (1967) 'On the Principle of Total Evidence', *British Journal for the Philosophy of Science*, 17, 4: 319–21.
- James, W. (1896) *The Will to Believe: And Other Essays in Popular Philosophy*. London: Longmans, Green & Co.
- Joyce, J. (2005) 'How Probabilities Reflect Evidence', *Philosophical Perspectives*, 19: 153–78.
- Kierkegaard, S. (1941) *Concluding Unscientific Postscript*. Trans. D. F. Swenson. Princeton, NJ: Princeton University Press. Originally published in 1846.
- Rothschild, M. and J. Stiglitz (1970) 'Increasing Risk: I. A Definition', *Journal of Economic Theory*, 2, 3: 225–43.

Belgium; the Philosophy of Religion Society at San Francisco State University; and the Cal Undergraduate Philosophy Forum, Brandon Fitelson's Epistemology seminar, and PhilFemme, all at UC Berkeley.