

Can it be Rational to have Faith?¹

Lara Buchak

I. INTRODUCTION

My concern in this paper is the relationship between faith and 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 to have 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>I have faith in your abilities.</i>	<i>He has faith that his spouse won't cheat on him.</i>
<i>I have faith in you.</i>	<i>He has faith that you won't reveal his secret.</i>
<i>She acted on faith.</i>	<i>She has faith that her car will start.</i>
<i>It was an act of faith.</i>	<i>I have faith that God exists.</i>
<i>I have faith in God's goodness.</i>	<i>I have faith in God.</i>

¹ Forthcoming in Philosophy of Religion: An Anthology, 7th edition, eds. Louis P. Pojman and Michael Rea, Wadsworth. Reprinted from L. Buchak (2012) ‘Can it be Rational to Have Faith?’, Chapter 12 of *Probability in the Philosophy of Religion*, Oxford University Press, 225-247. Discussion has been edited to simplify and make accessible the primary mathematical results.

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*.² 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

² 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.

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.

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 that *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 that *X* (expressed by some act *A*) and the evidence he has for *X*. We make assertions of faith only when the outcome 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.³ Are there further

³ Although the following worry arises from the possibility of over-determination. 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

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 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 express 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 act of faith constitutes taking a risk on *X*. What is

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.

it for an act to constitute taking a risk on a proposition? We can say that doing *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* on the supposition that *X* holds and *B* is preferred to *A* on the supposition that $\sim X$ holds. For example, telling my friend a secret constitutes taking a risk that she will keep it because on the supposition that she will keep it, I would rather tell her, but on the supposition that she won't keep it, I would rather not tell her. Whether an act constitutes a risk on *X* will of course be relative to the individual performing the act. So, as a first pass, we might say:

A person has faith that *X*, expressed by *A*, only if that person performs act *A* and performing *A* constitutes taking a risk on *X*.

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 judgment about *Q* such that *P* has faith that *X*, expressed by *A*.⁴

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:

⁴ Notice that the judgment must be positive from the point of view of the agent, in the sense that the agent has a preference for *A* on the supposition that *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.

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 .

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.

First, however, let me briefly say something about the philosophical framework in which I am working. One way to talk about belief states is as “on” or “off”: you either believe something or you fail to believe it (this could be because you actively disbelieve it or because you withhold judgment from it). But notice that of the many things you believe, you are more confident of some than others; furthermore, of the many things you fail to believe, you take some to nonetheless be more likely than others. As evidence of these facts, we could ask you how much you are willing to bet on various propositions: in general, you will be willing to bet more (accept lower odds) on the ones in which you are more confident. To capture these facts, many philosophers have found it helpful to talk about belief states in a “graded” way: beliefs come in degrees (sometimes called ‘credences’ or ‘subjective probabilities’) between 0 and 1, and degrees of belief should behave like probabilities. For example, if you assign a degree of belief 0.7 to the claim that the Giants will win the World Series ($p(\text{GIANTS}) = 0.7$), then you should assign degree of belief 0.3 to the claim that the Giants won’t win the World Series. As we will see later, degrees of belief underlie a formal theory of rational decision-making, but for now we can just notice that choosing an action is like taking a bet whose outcome is determined by the state of the world. For example, when one assigns $p(\text{RAIN}) = 0.6$ and $p(\text{NOT-RAIN}) = 0.4$, and one leaves the house without one’s umbrella, one is taking a bet that involves a 60% chance of getting wet and a 40% chance of staying dry.

With this framework in mind, let us consider some attempts to analyse the way in which faith goes beyond the evidence. 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 that X , he must think that the evidence warrants believing X to some degree, say, r , but he nonetheless believes X to degree q , where q is greater than r . More particularly, one might think that faith requires believing X to degree 1 (being certain or fully confident that X), 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 prejudge the 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 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 is greater than r . For example, one would take the bets that would be rational to take if one had degree of belief r . Thus, one can maintain a 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. Again, more particularly, we might think that faith requires acting as if you have credence 1, that is, using $p(X) = 1$ when making decisions. This would amount to acting as if one is certain that X : 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.

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

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 particular 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 that 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 took 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

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 an 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, e.g., in the gamble that results in a million lives lost if God doesn't exist and one life improved mildly if God does exist.

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 analysis of faith holds that faith requires setting one's degree of belief to $p(X) = 1$ prior to 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 individual sets $p(X) = 1$ (either in his beliefs or for his actions) 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 credences before encountering any evidence. 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 initial credences.⁷ This vindicates William James' claim that one's non-rational or 'passional' nature must determine 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 insight of Søren Kierkegaard's pseudonymous Johannes

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

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 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 do 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 bought 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 holds, 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 H_2O and lose \$100 if it is not. One might reply that from an objective standpoint, doing A is a risk—because setting one's initial credences 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 that it is 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 initial credences 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 the credences he brings to the table before examining evidence: a person who starts out sceptical, but who then amasses evidence in favour of *X*, could indeed end up choosing to have faith that *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 that *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*, and performing *A* constitutes taking a risk on *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.*

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 he wants the evidence for purposes other than deciding between the available acts. 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 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 action 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

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.

⁹ Thanks to Sherrilyn Roush for this example.

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, and performing A constitutes taking a risk on 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}.¹⁰

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, have the same probability function, value the possible outcomes in the same way, and perform 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 $p(\text{Dan will keep a secret}) = 0.9$; and that both have the same utility

¹⁰ I formulate this condition in terms of preference because the theory of practical rationality that forms the basis of the discussion in the rest of this paper concerns preferences rather than choices. I note, however, that if the reader finds it more natural to formulate this condition in terms of a disposition to choose to commit to A, rather than formulating it in terms of a preference to commit to A, that is acceptable. (In any case, preference is strongly linked to a disposition to choose.)

functions (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 make some fairly standard assumptions about degrees of 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 initial degrees of belief. An important upshot of these assumptions is that a rational person can only change his degrees of belief 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.

Epistemic rationality concerns believing what the evidence suggests. Practical rationality, on the other hand, concerns selecting the means to achieve one's ends.

An informal way to characterize practical rationality is that it consists in taking the means to one's ends. And just as philosophers use a formal theory (probability theory) to characterize epistemic rationality more precisely, so too do they use a formal theory to characterize instrumental rationality more precisely: decision theory. Recall the earlier point that an action is essentially a gamble in which the state of the world determines which outcome you get. We already know that we can represent the likelihood you assign to various states – and thus the likelihood that an act will lead to each particular outcome – by a degree

of belief function. According to decision theory, we can also represent how much you value particular outcomes by a ‘utility’ 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. Here ‘mathematical expectation’ just means a weighted average: the value of each outcome is weighted by the likelihood of getting that outcome, and the resultant values are summed. So, for example, if $p(\text{RAIN}) = 0.6$ and $p(\text{NOT-RAIN}) = 0.4$, and if $u(\text{getting wet}) = -3$ and $u(\text{staying dry}) = 2$, then $EU(\text{don't bring umbrella}) = 0.6(-3) + 0.4(2) = -1$. If $EU(\text{bring umbrella}) = u(\text{carrying around an umbrella}) > -1$, then a rational decision-maker will bring her umbrella.

In my view, EU theory is too restrictive. However, since this is the widely accepted view, and since I agree that expected utility maximizers *are* practically rational (I merely hold that there are other ways to be rational as well), I will postpone discussion of an 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 degrees of belief, including any degrees of belief concerning the relationship between particular hypotheses and particular pieces of evidence, as long as one’s degrees of belief cohere with *each other*. 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

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

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 *given the agent's own attitudes* 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 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.

Finally, I want to be clear that I am not assuming here that there are any benefits to believing false propositions, or to holding certain beliefs regardless of whether they are true. In my analysis, I assume that knowing the truth in the matter of *X* is always beneficial to the agent who is making a decision based on *X*: it allows her to choose the action that in fact turns out better for her.

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—

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.

whether one has faith is completely separate from whether one is epistemically rational because it is separate from whether one has appropriately evaluated the evidence one has.¹³

But can faith also be practically rational? Faith requires two actions: choosing *A* rather than the alternatives and not looking for further evidence for the purposes of this choice (or committing to *A* before seeing further evidence rather than seeing further evidence before deciding). Therefore, faith can only be practically rational if both of these actions are practically rational. Assessing the rationality of the first action is fairly straightforward: choosing *A* rather than 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 action? Can it be rational to commit 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 there is an alternative *B* such that he prefers *A* to *B* on the supposition that *X*, and *B* to *A* on the supposition that $\sim X$. Let us simplify by assuming that *A* and *B* are the only two alternatives, and that *B* is such that its value does not depend on *X*, that is, the outcome of *B* is not affected by whether *X* holds.¹⁴ To put some concrete utility values in place:

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

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 commit to an act before seeing more evidence to postponing the decision.

¹⁴ We can represent this by stating that the individual is indifferent between *B* & *X* and *B* & $\sim X$. 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. As it turns out, the assumption that the agent is indifferent between *B* & *X* and *B* & $\sim X$ is not crucial to the results discussed here: see Buchak (2012) for details.

We could generalize by assuming 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 $H > M > L$. In other words: A performed when X holds yields the best outcome, A performed when $\sim X$ holds yields the worst outcome, and B performed either way yields an intermediate outcome.

Here are some examples of decisions that might include the relevant values. Consider an individual 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 individual 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 and that the individual lives roughly the same life as a non-monk whether or not God exists. To take another example, consider an individual 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 individual 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 individual 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.

Let us assume that given the individual's current credences and utility function, A has a higher expected utility than B , and so A is practically rational given the individual's current information. Now we want to know whether practical rationality requires that the individual 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

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.

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.¹⁶ In other words, if these two conditions are met, and if rational individuals must choose the act that maximizes expected utility, then 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 for the purposes of the choice at hand. 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 experiment yielding each possible result, and can say how he ought to update his credences given each possible result. For example, if the individual from our above example is considering the 'experiment' of turning on the weather channel to see whether the weatherman predicts rain, he may assign credence 0.65 to the weatherman predicting rain and credence 0.35 to the weatherman predicting no rain; and he may think that if the weatherman does predict rain, he ought to raise his credence in the hypothesis that it will rain to 0.9, and if the weatherman does not

¹⁶ For a further discussion of Good's setup and his result, see Buchak (2010) and Buchak (2012).

predict rain, he ought to lower his credence to 0.05.¹⁷ Furthermore, I assume that after seeing the results of the experiment, the individual will update his credence in RAIN and NOT-RAIN accordingly and choose the action that maximizes expected utility given his updated credences.

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 no matter what the results of the experiment, *A* still maximizes expected utility given his updated credences. In less technical terms, this will hold when none of the observations the individual is considering would tell against *X* conclusively enough to make him perform *B* instead of *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 I myself 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.

We can measure the value of an experiment by considering the difference between the expected utility of {performing the experiment and then choosing an act} and the expected utility of {choosing an act without further evidence}. If the former number minus the latter number is positive, rationality requires one to perform the experiment, and if it is negative, rationality requires one not to perform the experiment (if this number is 0, the agent will be neutral about performing the experiment, as in the case just discussed). We can examine which factors contribute positively to this value and which contribute negatively: these will roughly be the pros and cons associated with performing the experiment, and so we can precisely characterize in what situations the cons outweigh the pros. The most important positive factor is associated with the possibility that $\sim X$ holds and the evidence one gets leads one to do *B* instead of *A*. (For example, the third party says your friend is a gossip and so you don't tell her your secret, and your friend is in fact a gossip – you have benefitted by getting a

¹⁷ Note that these numbers need to satisfy the constraints that $p(\text{RAIN}) = p(\text{RAIN} \mid \text{PREDICT})p(\text{PREDICT}) + p(\text{RAIN} \mid \text{NOT-PREDICT})p(\text{NOT-PREDICT})$ and $p(\text{NOT-RAIN}) = p(\text{NOT-RAIN} \mid \text{PREDICT})p(\text{PREDICT}) + p(\text{NOT-RAIN} \mid \text{NOT-PREDICT})p(\text{NOT-PREDICT})$. The credences above do satisfy these constraints, with rounding.

payoff of 1 rather than 0, because you have refrained from telling your secret, rather than told your secret and have it revealed.) Another positive factor is associated with the possibility that X holds and the evidence one gets makes one more certain that A is the right choice. (For example, the third party confirms that your friend is not a gossip, and your friend is in fact not a gossip.) However, this factor will make less of an overall difference in most of the situations we are considering, since if one's credence in X is already high, additional confirmation will not make a substantial difference.

There are also factors which negatively impact the value of an experiment. As Good notes, it can be rational to eschew additional evidence if that evidence is costly. 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 the cost remains the same under all eventualities, and it is a simple matter to calculate it.

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 and when they can 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 constant 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.

VIII. THE COSTS 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 listen to one’s secret today, so if one does not reveal it, one will lose that option forever. Or it might be that one is deciding whether to get married and 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*.¹⁸

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 do these costs 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 doing *A* when *X* obtains without

¹⁸ Recall the supposition that one’s credence before gathering evidence is high enough to make *A* the rational choice.

gathering more evidence (or while committing to A regardless of the evidence) and doing A when X obtains after gathering more evidence.

We can now 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 (the value of the ‘middling’ act B) is lower, and H and L (the value of A if X obtains and if $\sim X$ obtains, respectively) are higher; that is, when there is less of a risk involved in doing A . Next, it will be lower to the extent that one of the following situations is likely to result (that is, to the extent that these situations are probable, given the evidence). Both situations are situations in which X in fact holds. The first is when X in fact holds and the experiment yields some result that will make it rational for the individual to do A . In this situation, performing the experiment hurt the individual because he paid a cost c that he otherwise would not have paid. (For example, the third party says your friend is trustworthy and so you stick with your original plan of sharing your secret, but you incur the cost of the time it took to listen to the third party – perhaps by the time you ask the third party, the friend is no longer available today.) The second situation is when X in fact holds but the experiment yields some result that makes it no longer rational for the individual to do A . This is a situation in which the evidence is ‘misleading’ in the sense that it leads to the rational performance of an action that in fact has a lower payoff than its alternative. (For example, the third party says your friend is a gossip, and so you don’t reveal your secret, but your friend is in fact trustworthy – here you get a payoff of 1 rather than 10, because you have refrained from telling your secret rather than telling it to a trustworthy friend.)¹⁹ Both of these situations are more likely the higher the credence one already assigns to X . The second of these situations is more likely to the extent that evidence against X is *non-conclusive*: to the extent that it will lower one’s credence enough to make it rational to do B instead of A , but not enough to *guarantee* that B is in fact the better choice.

The overall expected utility of gathering evidence depends on how likely each of the ‘beneficial’ situations is (the situation in which you get evidence that steers you away from your original choice when that would have in fact been the wrong choice, and, to a lesser extent, the situation in which you get evidence that confirms your original choice when that is in fact the right choice) as compared to how likely each of the ‘detrimental’ situations is (the situation in which you get evidence that doesn’t alter your choice but is costly, and the

¹⁹ See Buchak (2012) for specific mathematical results.

situation in which you get misleading evidence that steers you away from what would have in fact been the right choice). 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 one already has a high credence in X ; or (3) 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. This third possibility holds to the extent that the potential evidence that could tell against X does not tell conclusively against X .

The fact that costs associated with postponing a decision can make faith rational vindicates an observation made by William James, 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, provided one has the appropriate credences, and provided these are the correct credences to have given one’s evidence.

The upshot of this discussion is that, if we accept expected utility 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.²⁰

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 the character of the potential evidence one might encounter: in particular, faith is more apt to be rational if potential evidence against *X* will be inconclusive.

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 choice not to gather additional evidence—or more than a choice to commit to an action before the evidence comes in. It requires a choice to not gather 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 expected utility maximization is too strong a criterion of rationality, and that one can be practically rational without being an expected utility maximizer.

²⁰ 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.

While I am sympathetic to the general aim of decision theory, and hold that expected utility 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 modified.²¹ In my view, expected utility theory dictates a too-narrow way in which risk-considerations can play a role in an individual's choices. It dictates that rational individuals cannot, for example, care proportionately more about what happens in the worst-case scenario than the best, or vice versa. But most people do pay special attention to these features of their decisions. Moreover, rather than being an example of human irrationality, I think this tendency can be rational.

Above, I explained that decision theory formalizes means-ends reasoning: utility corresponds to how much an individual values particular outcomes, while credence corresponds to the likelihood with which the individual thinks some particular act will realize one of these outcomes. But even once we know an individual's credence and utility function, there is an additional question of how to move from considerations about how a choice will turn out under various circumstances to an overall evaluation of that choice. For example, how should one trade off a small chance of something great against a high chance of something fairly good? We need to know the relative priority the individual places on the worst-case and best-case scenario (and all the scenarios in between). We might think of this additional factor, the individual's risk-attitude, as corresponding to how he structures the potential realization of some of his aims. What is at issue here is, roughly speaking, the following: when an individual 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 'global' features of a decision like how spread out the utility values are?²² In an earlier paper, I demonstrated that individuals who care proportionately more about what happens in the worse-case scenario will sometimes rationally reject cost-free evidence (see Buchak 2010). In particular, in scenarios like those outlined in the previous section, it will be rationally required for these individuals to commit to an action *A* before looking at additional evidence

²¹ See L. Buchak (2013).

²² I say 'roughly speaking' because on certain interpretations of what utility amounts to, this will be a controversial way of putting things. But the exact formulation is not important to the point here, which is that there are some ways of evaluating gambles that expected utility theory cannot capture, that appear to stem from individuals caring proportionately more about certain possible outcomes of a gamble and less about others, and that (I argue elsewhere) are rational.

rather than to look at additional evidence and then decide. This will be the case in situations with the following properties, familiar from the previous section: (1) the individual already has a high credence in X (that is, $p(X)$ is antecedently high); and (2) the experiment is likely to result in *misleading* evidence against X (again, evidence that makes him ‘miss out’ on the possibility of doing A when X in fact holds) because potential evidence that tells against X will not be conclusive. This is the case because if an individual already has a high credence that X , more evidence in favour of X won’t be very helpful; and while evidence against X could be helpful if $\sim X$ holds, it will be harmful if X holds – and to the extent that evidence against X is inconclusive these two situations will be closer in likelihood.

In short, agents who care about global features of decisions are concerned with a particular risk involved in looking for additional evidence: the risk of coming across evidence that makes it rational to refrain from doing A even though X in fact holds. In other words, they are concerned, and rationally so, about the risk of getting evidence that is misleading. I already mentioned that the possibility of misleading evidence will make it rational for the EU maximizer to reject costly evidence. However, on the more permissive decision theory, the risk of misleading evidence makes faith rational even in cases in which there is no cost to looking for evidence. If one accepts the more permissive decision theory, then faith is rational in more cases than it is on standard decision theory, and precisely in cases in which the risks of getting misleading evidence outweigh the benefits of getting non-misleading evidence.

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) credence 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 potential piece of evidence would tell 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.²³ So, in a rough-and-

²³ 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).

ready way, we might say that faith that X (expressed by some particular act A) is practically rational to the extent that the individual's degree of belief in X is already based on a large body of evidence. Second, whether faith is rational depends on the kind of situation we find ourselves in. Faith will be rational to the extent that potential counterevidence wouldn't be very conclusive with respect to the position in question, 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 circumstances, perhaps circumstances some of us find ourselves in some of the time. 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.

Works Cited

- Adams, R. (1976) 'Kierkegaard's Arguments against Objective Reasoning in Religion', *The Monist*, 60, 2: 228–243.
- Buchak, L. (2010) 'Instrumental Rationality, Epistemic Rationality, and Evidence-gathering', *Philosophical Perspectives*, 24, 1: 85–120.
- Buchak, L. (2012) 'Can it be Rational to Have Faith?', Chapter 12 of *Probability in the Philosophy of Religion*, Oxford University Press, 225-247.
- Buchak, L. (2013) *Risk and Rationality*, Oxford University Press.
- Dreier, J. (1996) 'Rational Preference: Decision Theory as a Theory of Practical Rationality' *Theory and Decision*, 40: 249–276.
- Good, I. J. (1967) 'On the Principle of Total Evidence', *British Journal for the Philosophy of Science*, 17, 4: 319–321.
- 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: 153–178.
- Kierkegaard, S. (1846) *Concluding Unscientific Postscript*. Translated by D. F. Swenson. Princeton, N. J.: Princeton University Press, 1941.
- Rothschild, M. and J. Stiglitz (1970) 'Increasing Risk: I. A Definition', *Journal of Economic Theory* 2, 3: 225–243.