UC Berkeley

UC Berkeley Previously Published Works

Title

Faith and steadfastness in the face of counter-evidence

Permalink

<https://escholarship.org/uc/item/9x9649v4>

Journal

International Journal for Philosophy of Religion, 81(1-2)

ISSN

0020-7047

Author

Buchak, L

Publication Date

2017-04-01

DOI

10.1007/s11153-016-9609-7

Peer reviewed

Faith and Steadfastness in the Face of Counter-evidence

Lara Buchak, UC Berkeley

1. Introduction

It is sometimes said that faith is recalcitrant in the face of new evidence, in some sense or other,¹ but it is puzzling how such recalcitrance could be rational or laudable. In this paper, I use my account of faith—that faith requires committing to a risky act before examining further evidence—to explain this aspect of faith and why faith is not only rational, but in addition serves an important purpose in human life. Where I have elsewhere focused on defending the claim that deciding without further evidence can sometimes be rational, I here focus on the commitment itself. On my account, faith requires that one maintain a commitment to act on the claim one has faith in, even in the face of counter-evidence. I show why maintaining such a commitment can be rational over time, even if it is not rational at some particular times. Thus, one important role of faith is to allow us to carry out long-term, risky projects—to maintain *integrity* over time. This makes clear one way in which faith is an important virtue.

2. Rational Faith and Commitment to a Risky Act

I here briefly rehearse my account of faith—the *risky commitment* account—and when and why such faith is rational.² This account covers both mundane and religious faith. In the case of religious faith, it is intended to explicate the religious *virtue* of faith—what it is to *have faith*, what attitude Abraham was exemplifying when he journeyed *by faith*—rather than, say, what it is to be an adherent of a particular religion—what it is to have *a* religious faith, what attitude one must have to count as an adherent of a religion. This account is an account of *propositional* faith, and only certain propositions are candidates for faith at all. In order for a proposition to be a potential object of faith for a subject, the subject must care whether the proposition is true, must have a positive attitude towards the proposition, and must not be certain of the proposition on the basis of his evidence alone—his evidence must leave it open that the proposition is false.³

¹ Howard-Snyder (2013), Pace (ms.), Wolterstorff (1990). Wolterstorff interprets Calvin (1536) as holding that steadfastness and tenacity—perseverance in the struggle against unbelief—are central to Christian faith.

² For more detail, see Buchak (2012).

³ At least initially and with respect to ‘objective’ standards of evidence—one might hold that cases of faith in which one eventually becomes certain can continue to count as faith.

Faith is tied to action, and having faith is a matter of being disposed to act in certain ways: in particular, being willing to take risks on the basis of the claim one has faith in. When we have faith that a particular individual will act in a certain way—keep our secret, do what is in our best interests, fulfill her promises—we take a risk that the individual will let us down. We are vulnerable to the individual not acting as we have faith that she will act, in the sense that we will bear some cost if she lets us down. When we have faith that the world is a certain way, we bear some cost if we are wrong.

But not every case of risk-taking will be an act of faith. Faith involves a willingness to commit to acting on the proposition one has faith in, without first looking for further evidence for or against that proposition. An individual with faith in her friend's ability to keep a secret must be willing to confide in her friend without first verifying with a third party that the friend isn't a gossip. A man who has faith that his wife is true to her vows must commit to his marriage without reading her private correspondence. An individual who has faith that a particular bridge will hold his weight doesn't test the bridge before stepping onto it. Not only do individuals with faith not need further evidence, they will choose not to obtain it if it is offered to them, when their only interest in obtaining it is in how it bears on the decision to act. For example, I must decline if a third party offers to tell me about her experiences with my friend's secret-keeping abilities. (A person might choose to obtain evidence for other reasons—consider the example of the evangelist who looks for evidence for God's existence in order to convince others or who reads the Bible to discover what God requires of him while already committed to fulfilling these requirements—but in order for this to be consistent with faith his decision to act must be independent of that evidence.) Individuals with faith will commit to the risky act without looking for further evidence.

Thus, a subject has faith in some candidate proposition if he is willing to commit to taking risks on the proposition without examining additional evidence. More formally, Act A is an act of faith that X, for S, if and only if:

- (i) S performing act A constitutes S taking a (subjective) risk on X
- (ii) S chooses {to commit to A before he examines additional evidence} rather than {to postpone his decision about A until he examines additional evidence}.

And the degree or strength of the subject's faith is a matter of which such acts he is disposed or willing to perform—S has a higher degree of faith in X to the extent that he is willing to take more risks on X without looking for further evidence. (More on an individual's strength of faith in the next section.)

It is important that condition (2) concerns *commitment* rather than simply action—this is what allows the faithful individual to in fact intake additional evidence, whether for other purposes or accidentally, while

continuing to act as faith requires—and the rationality of that commitment over time will be the focus of this paper. In this section I will just focus on the rationality of the act at the initial time, since that was the focus of previous work.

An initial thought is that such faith is always irrational—even if it has other virtues—because acting on the basis of more evidence better conduces to getting what you want. However, as it turns out, there are situations in which, because of the structure of the available acts and the available evidence, it is instrumentally rational—rational *from the point of view of getting what you want*—to act rather than to examine additional evidence.⁴

For example, let's say that Mary and Andy have been colleagues for a long time, and Mary has abundant evidence that Andy is upstanding. However, she receives word that Andy is under investigation for fraud, and the results of the investigation are soon to be released. Mary does not have conclusive evidence that Andy didn't commit fraud, but she has faith that he didn't. She is venturing out on her own and has to select a business partner. Because she has faith in Andy, she does not wait until the investigation is concluded to select Andy as a partner, even knowing that it would be a disaster to partner with someone capable of fraud. (Many such examples abound, of both a religious and mundane nature: when Jesus called Simon Peter and Andrew, they immediately stopped what they were doing, without question, to follow him; when one has faith that a prospective spouse would be a good marriage partner, one commits to marriage without needing to see their behavior under every relevant circumstance.)

We can analyze the rationality of such faith using the framework of formal epistemology. Three concepts from formal epistemology are relevant here: credences, utilities, and risk-attitudes. For purposes of this paper, I will understand credences—an individual's subjective degrees of belief—to express how much support she thinks various propositions receive from her evidence or how likely she thinks various propositions are to be true, given her evidence. The argument here will be neutral between a purely subjective view, on which one can have any internally coherent credence assignment, and a view according to which one must meet additional credence-to-evidence standards in order to count as rational. I will understand utilities—an individual's desires—to be a measure of how valuable a consequence is for the individual. Again, I will be neutral between a "Humean" view of utility, on which one can have any internally coherent utility assignment, and a view according to which one must meet additional standards. Finally, I will understand risk-attitudes to express the extent to which an individual facing a risky gamble

⁴ See Buchak (2012) and especially Buchak (2010).

cares about what happens in relatively worse states rather than what happens in relatively better states, or the extent to which an individual is willing to prioritize security over additional potential payoff. An individual is risk-avoidant to the extent that she cares more about worse scenarios or prioritizes security more highly, and risk-inclined to the extent that she cares more about better scenarios or prioritizes additional potential payoff.⁵

To see whether it is preferable for Mary to partner with Andy without first hearing the results of the investigation or to instead hear the results before deciding, we compare three different acts: partnering with Andy; partnering with some safer option, say, Bob; and hearing the results of the investigation and then deciding who to partner with in light of what one learns. And we compare these acts by considering the utility values they yield under the relevant unknowns: what the investigation will uncover, and whether Andy is in fact innocent or guilty.

Assume for the sake of example that if Andy is innocent then partnering with him will yield a payoff of 9 utils and that if he is guilty then partnering with him will yield a payoff of -1 utils; and that partnering with Bob will yield a payoff of 7 utils no matter what. Let “positive” or “negative” evidence be evidence that tells in favor of or against Andy’s innocence (“positive” evidence is a conclusion by the investigators that he is innocent; “negative” evidence is a conclusion that he is guilty). Let “correct” evidence be evidence that tells in the correct direction (it tells in favor of Andy’s innocence when he is in fact innocent or against Andy’s innocence when he is in fact guilty), and “misleading” evidence be evidence that tells in the wrong direction. Finally, let c be the cost of postponing the decision to partner with Andy, should Andy in fact be innocent.⁶

There are two relevant cases. The first is one in which Mary’s current credences are high enough to make it rational to partner with Andy, and furthermore they will still be high enough to partner with Andy if the investigation finds him guilty.⁷ The payoffs for this case are:

⁵ Credences and utilities are standard in decision theory, whereas the existence and rationality of risk-attitudes is more controversial—they are particularly associated with Buchak’s (2013) risk-weighted expected utility theory.

⁶ There may also be costs to postponing the decision if one eventually partners with a guilty Andy or with Bob, but I will follow Buchak (2012: 242) in ignoring these costs, which can only widen the set of situations in which postponing the decision is irrational.

⁷ I am assuming that Mary will update her credences in response to the new evidence and form preferences on the basis of these credences. The conclusion of this paper does not ultimately depend on the correctness of this assumption, but the assumption makes it easier to frame the discussion.

	Found innocent & actually innocent (Positive correct evidence)	Found innocent & actually guilty (Positive misleading evidence)	Found guilty & actually innocent (Negative misleading evidence)	Found guilty & actually guilty (Negative correct evidence)
A: partner with Andy	9	-1	9	-1
B: partner with Bob	7	7	7	7
C: wait for the report and then choose a partner	$9 - c$	-1	$9 - c$	-1

The payoffs of C were determined by noting that if Mary waits, she will still choose to partner with Andy and get the associated payoffs.

It is easy to see that if there are no postponement costs, then A and C are equivalent, and if there are, then A dominates C. So faith in Andy is rational if one's credences are antecedently high enough to make partnering with Andy rational and if the potential evidence, even if it tells against Andy, will still make it rational to partner with Andy—and faith in Andy is rationally *required* if these conditions obtain and there are costs to postponing the decision. (These conditions are apt to attain to the extent that Mary has a *lot* of evidence about Andy, so that very little could lower her credences, or to the extent that the investigation itself inherently provides weak evidence: it is sloppily conducted or hopelessly corrupt.)

The more interesting case is one in which Mary's current credences make it rational to partner with Andy, but more evidence might change her credences in such a way that she prefers to partner with Bob. Thus, assume that if the investigation report concludes that Andy is guilty then it will be rational to partner with Bob. Then we have the following payoffs:

	Found innocent & actually innocent (Positive correct evidence)	Found innocent & actually guilty (Positive misleading evidence)	Found guilty & actually innocent (Negative misleading evidence)	Found guilty & actually guilty (Negative correct evidence)
A: partner with Andy	9	-1	9	-1

B: partner with Bob	7	7	7	7
C: wait for the report and then choose a partner	$9 - c$	-1	7	7

The payoffs of C were determined by noting that if the investigation reports that Andy is innocent, then Mary will partner with Andy and get the payoffs associated with doing so (minus the postponement cost), and if it reports that he is guilty, then she will partner with Bob and get the payoffs associated with doing so.

Mary will be rationally required to have faith that Andy is innocent, in the sense of partnering with Andy rather than postponing until she gets more evidence, if partnering with Andy not only has a higher instrumental value than partnering with Bob, but also has a higher instrumental value than waiting for the evidence before deciding. Instrumental value here is analyzed as maximizing expected utility or maximizing the more general risk-weighted expected utility, which allows for the rationality of various degrees of risk-avoidance and risk-inclination.⁸ Several key features of the act and evidential situation determine whether partnering with Andy indeed has the highest instrumental value. First, as noted, it must be the case that Mary's current evidence recommends partnering with Andy over partnering with Bob. Second, the likelihood of the states in which partnering with Andy is better than waiting to decide must be sufficiently high relative to the likelihood of the states in which waiting to decide is better than partnering with Andy, where what is "sufficiently high" depends on how much better these states are. There are three important differences between partnering with Andy and waiting: partnering with Andy is better if the evidence is positive and correct, in that doing so yields 9 rather than $9 - c$; partnering with

⁸ Let $g = \{E_1, x_1; \dots; E_n, x_n\}$ be an ordered gamble that yields outcome x_i in event E_i , where $x_1 \leq \dots \leq x_n$. Then the expected utility of g is $EU(g) = \sum_{i=1}^n p(E_i) u(x_i)$. The risk-weighted expected utility of g is

$$REU(g) = \sum_{i=1}^n \left[r \left(\sum_{j=i}^n p(E_j) \right) (u(x_i) - u(x_{i-1})) \right]$$

where r is "risk function" from $[0, 1]$ to $[0, 1]$, with $r(0) = 0$, $r(1) = 1$, and r non-decreasing. Risk-avoidance corresponds to a convex risk function and risk-inclination to a concave risk function. If $r(p) = p$, then REU reduces to EU. See Buchak (2013).

Andy is better if the evidence is negative and misleading, in that doing so yields 9 rather than 7; and waiting is better if the evidence is negative and correct, in that doing so yields 7 rather than -1.

The features that make partnering with Andy superior are all present to the extent that the following conditions hold:⁹

- (1) Mary is antecedently fairly confident that Andy is innocent.
- (2) If the investigation says that Andy is guilty, then (even though it will lower Mary's credence enough to prefer partnering with Bob) it won't tell *conclusively* in favor of his guilt.¹⁰
- (3) Postponing the partnership with Andy would be costly OR Mary is risk-avoidant.

Call conditions (1) and (2) the *credence* conditions and the elements of condition (3) the *value* and *attitude* conditions. Notice that the credence conditions are more apt to obtain if Mary antecedently has a large body of evidence in favor of Andy's innocence—if the *weight* of the evidence is high.¹¹ For example, if Mary is already well-acquainted with Andy, then even if the investigation says he is guilty, her credence in his innocence will only drop somewhat, since the guilty verdict must be weighed against her own large body of evidence. When all three of these conditions obtain, Mary is in a situation in which looking for more evidence can harm her, from the point of view of her taking actions that are best for her in the actual world. Mary is not merely rationally permitted to have faith in Andy under these conditions; she is rationally required to have it.

The downsides of looking for more evidence before acting are that one might get “misleading” negative evidence—one is talked out of a risk that would have paid off—and that one might bear a cost for postponing the decision (for looking for evidence but then taking the risk anyway, when it would have been at least as good to take it immediately). The upside of looking for more evidence before acting is

⁹ Buchak (2012).

¹⁰ Evidence is more conclusive to the extent that it is highly correlated with the truth of the proposition it is evidence for. So, for example, if negative evidence would result in credence $p(\text{Andy is innocent}) = 0.5$, it is less conclusive than negative evidence that would result in credence $p(\text{Andy is innocent}) = 0.1$.

¹¹ Terminology originates with Keynes (1921). See Joyce (2005) on the difference between the dimensions of strength, weight, and specificity, and on how they are each reflected in an individual's credence. Joyce notes (following Skyrms (1980)) that “the weight of the evidence for a proposition X often manifests itself not in X 's unconditional credence, but in the *resilience* of this credence conditional on various potential data sequences”(161). Joyce later notes that weight and resilience can come apart, particularly in cases when a subject's credence is mediated by chance hypotheses, but this case won't concern us here.

that one might get “correctly leading” negative evidence, so that one is talked out of a risk that would not have paid off. The downsides are apt to outweigh the upsides to the extent that one has a large body of evidence in favor of the proposition one is considering taking a risk on, the evidence out there in the world is sparse, and either postponing the decision is costly or one is risk-avoidant.

To get a sense of how one might meet the credence conditions, here is an example. Mary’s current credence that Andy is innocent is 0.9. If the report says he is guilty (a possibility to which she assigns credence 0.12), her credence will drop to 0.83—low enough to make her prefer partnering with Bob instead, given the risk of partnering with Andy if he is guilty, but not very low. There are no postponement costs, but Mary is moderately risk-avoidant. In this case, it will be rational for Mary to partner with Andy rather than waiting.¹² (A different example of meeting the credence conditions focused on a case in which the risky act was less risky and more rewarding, and where the potential counter-evidence would drop the actor’s credence in the relevant proposition from 0.89 to 0.05.¹³) When thinking about the character of future possible evidence that makes it better to act now, the key is that there is a window in which counter-evidence would be negative but inconclusively so: the negative result would lower Mary’s credence in Andy’s innocence enough to make her want to take the safe course rather than the risky course, but not enough that Andy’s guilt is definitive. When counter-evidence wouldn’t lower Mary’s credence very much, then counter-evidence wouldn’t change Mary’s decision, so it will be just as good to wait as to act now, absent postponement costs. When counter-evidence would lower Mary’s credence a lot, then it will be better to wait. So faith is both ‘active’ (in that it can make a difference to one’s action) and rational in a particular strength-of-counter-evidence window. I will return to this point.

This summarizes the synchronic aspect of the rationality of acting on faith—when it is good, from the agent’s point of view, to take a “leap of faith” and act without looking for further evidence. I turn now to

¹² Let X be the proposition that Andy is innocent and E be the proposition that the report says he is innocent. Then set $p(XE) = 0.8$, $p(X\bar{E}) = 0.1$, $p(\bar{X}\bar{E}) = 0.02$, $p(\bar{X}E) = 0.08$. This yields $p(X) = 0.9$, $p(X | \bar{E}) = 0.83$. If we assume that Mary is risk-avoidant with $r(p) = p^2$ and that $c = 0$, then $REU(A) = -1 + (0.9)^2(9 - (-1)) = 7.1$, $REU(B) = 7$. If Mary learns $\bar{E}$, then $REU(A) = -1 + (0.83)^2(9 - (-1)) = 5.94$ and $REU(B) = 7$, so she will prefer to do B upon learning $\bar{E}$. Therefore, $REU(C) = -1 + (p(\bar{X}\bar{E}) + p(X\bar{E}) + p(XE))^2(7 - (-1)) + ((p(XE))^2(9 - 7) + (0.8)^2(2)) = 7.05$. So A is rationally preferred to both B and C.

¹³ Buchak (2010).

the diachronic aspect of faith: to committing to act even in the face of counter-evidence, and to following through on that commitment.

3. Rational Faith and Tenacity

Recall that condition (2) mentions commitment: the faithful individual must *commit* to the act of faith without looking for more evidence, which is to say that he must intend to follow through with the act regardless of what evidence comes in, even if he doesn't go looking for the evidence himself. For explicitness, let us divide the required attitude into two aspects in our analysis of faith, and say that Act A is an act of faith that X, for S, if and only if:

- (i) S performing act A constitutes S taking a (subjective) risk on X:
- (ii) S chooses {to perform A before he examines additional evidence} rather than {to postpone his decision about A until he examines additional evidence}.
- (iii) S would follow through on performing A even if he were to receive counter-evidence.¹⁴

Informally: a subject has faith in some candidate proposition if he is willing to take risks on the proposition without examining additional evidence—and to take these risks even if he receives counter-evidence. In order to count as having faith, the subject needn't be willing to follow through on A in the event of any counter-evidence whatsoever. How much counterevidence the individual would be willing to tolerate (tenacity) will be one dimension along which his faith can be stronger or weaker. The other dimension along which an individual's faith can be stronger or weaker is the size of risk he is willing to take: the bigger risks he is willing to take, the stronger his faith is in this dimension.¹⁵

On this way of dividing the conditions, what I've done so far is to argue for the rationality of (ii), but I haven't said much about the rationality (iii): of intending to follow through or actually doing so upon receiving counter-evidence.

¹⁴ Condition (iii) applies after the initial choice but before A has been performed or completed. Perhaps there is a more general requirement here: S would follow through on performing A even if his preferences about performing A change. For purposes of this paper, I will just focus on the evidential requirement.

¹⁵ In earlier work, explicated in section 2, I implied that there is a single dimension along which an individual can have faith: the extent to which he is willing to commit to risks without further evidence. I am now convinced by Pace (ms.) that resilience is a dimension of strength of faith that is distinct from the dimension of what risks one is willing to take. What types of evidence one does not need to look at before acting might also be considered another dimension, though for the rational agent this might go hand-in-hand with resilience. See also Wolterstorff (1990: 412), who points out that believing something with a high level of confidence and believing something tenaciously are two distinct phenomena—he thinks the latter is a central requirement of faith, whereas the former is desirable primarily “for the peace of mind which that brings.”

This aspect of faith is particularly important if the act of faith is one that takes place over a long time-scale—if it is a long-term *project*—during which one is likely to receive additional evidence even if one is not looking for it. Being married is one such project, and so are other interpersonal endeavors, such as maintaining a deep friendship. Working to bring about some moral good is often such a project. So are personal endeavors such as training for a marathon or completing graduate school. Finally, most serious religious commitments are of this nature: picking a religious vocation (such as becoming a nun), tithing, engaging in spiritual disciplines or practices, or even simply committing oneself to live according to the commands expressed by Moses, Jesus, or the Buddha. All of these projects involve serious risks and so could be acts of faith—acts of faith that one’s spouse or friend is loyal, that one’s work will succeed, that one will complete the marathon or graduate school, that God exists, that one has correctly interpreted one’s vocation, that spiritual disciplines are worth the effort, and so forth. (Indeed, these seem to be the paradigm cases, perhaps indicating that the diachronic aspects of faith are often more central than the synchronic aspects.)

We will use our above example: let us assume that Mary will not be able to help finding out the results of the investigation. When should Mary commit to the partnership, and should she in fact maintain the commitment if the results of the investigation are against Andy? There are three questions to consider. Let t_1 be the time at which Mary commits to the partnership, and let t_2 be the time at which new evidence has come in and Mary has updated her credences to reflect the new evidence.¹⁶ The first question is whether, from Mary’s point of view at time t_1 , it is better to plan to maintain the partnership with Andy regardless of what new evidence reveals, rather than to take any of the alternatives. (The alternatives include the non-risky act of partnering with Bob; postponing the decision until the evidence comes in; and deciding that one will begin a partnership with Andy but back out and switch to partnering with Bob if one gets counter-evidence.) The second question is whether, from Mary’s point of view at time t_2 , if she has in fact learned that the investigation found Andy guilty, it is better for Mary to maintain her commitment to Andy or to back out and switch to partnering with Bob. The third question is whether, when we consider Mary over time, it is better for her to make the commitment and follow through on it, rather than to not make the commitment at all, or to make the commitment and back out if negative information comes to light. We will only treat the case in which the results of the investigation would make Mary prefer to back out (i.e., the second case in section 2). The options at time t_1 are:

¹⁶ Again, we are assuming that Mary will conditionalize on her new evidence and form preferences on the basis of her new credences.

	Found innocent & actually innocent (Positive correct evidence)	Found innocent & actually guilty (Positive misleading evidence)	Found guilty & actually innocent (Negative misleading evidence)	Found guilty & actually guilty (Negative correct evidence)
A: partner with Andy, regardless of what the report says	9	-1	9	-1
B: partner with Bob	7	7	7	7
C: wait for the report and then choose a partner	$9 - c$	-1	7	7
D: partner with Andy but revise choice according to new credences	9	-1	7	7

The options at time t_2 are represented by only considering the first two columns or the last two, depending on what is learned.

The analysis in section 2 considered the relative merits of A and C, at time t_1 . We can now consider the relative merits of A and D, both at time t_1 and time t_2 . Notice that the payoffs of D are exactly the same as those of C when there are no postponement costs. Thus, A will be preferable to D, at time t_1 , if the individual is risk-avoidant and if the credence conditions (conditions (1) and (2) above) hold. **If learning the negative evidence will make Mary prefer to back out, but if the negative evidence will not be conclusive and she is already fairly confident that Andy is innocent, then if she is risk-avoidant she should choose A over D: she should *commit* to A regardless of what future evidence will tell her.** Again, what is driving the result is the possibility of misleading evidence: the possibility that Mary might be talked out of an action that is in fact good for her given the actual state of the world. Commitment buffers against this possibility. Of course, the only way the evidence could be misleading in this way is if it will make her prefer backing out, given her total evidence, at t_2 —thus, there will be a conflict between what Mary-at- t_1 wants to do if she gets negative evidence and what Mary-at- t_2 wants to do given that she's already gotten the negative evidence. More on this conflict shortly.

For a risk-avoidant individual, committing to an act is a way to mitigate the effects of potentially misleading evidence. If risk-avoidance is rational, then perfectly rational individuals can have reason to

commit. If risk-avoidance is irrational, then commitment to an act is a way for an irrationally risk-avoidant individual to mitigate against the effects of misleading evidence given his irrational risk-avoidance.

Mitigating against the possibility of misleading evidence is particularly important when we are engaged in long-term risky projects where evidence that tells against the success or worth of the project isn't apt to be conclusive. This is because if the time-scale is long enough, then even if the proposition one is taking a risk on is in fact true, we can expect to encounter misleading evidence at some point. Thus, faith allows us to adhere to an act over time—to complete a risky long-term project—in a way that is decoupled from future evidence that the project will fail or isn't worth it. *Faith keeps us from being blown about by the changing winds of evidence.*

The Mary/Andy example is just a toy example, but we can see that the long-term projects mentioned above tend to be like this. They involve a risk on a proposition that (we may assume) you already have strong and deep evidence for; evidence out there in the world is not strongly correlated with the truth of the proposition you are taking a risk on; but you can nonetheless expect evidence to come in before the risk-taking is complete. If one is involved in a friendship or marriage for long enough, one may expect to encounter evidence at some point that the friend or spouse is not a good partner. If one works hard for a cause, one may expect to encounter evidence at some point that one's efforts are ineffective or wrong-headed. And so forth. While there are perhaps many irrational temptations to give up on one's projects, faith can guard against even the rational temptation to give up on one's projects because of the evidence one in fact has.

Indeed, for many of these phenomena, the (informally) accepted procedure is first to gather a lot of evidence, and then close the search for evidence and make a formal commitment that it would be costly (interpersonally or personally) to back out of. In the case of contemporary marriage, two people date each other until they have good, though imperfect, evidence that they are suitable partners—they then have a formal ceremony to cement their commitment. In many Christian traditions, when one joins a religious order or becomes a priest or minister, the potential candidate first undergoes a discernment process and then takes vows. And there is evidence that religious conversion itself follows a typical pattern that begins with forming strong beliefs (evidence-based beliefs, at least subjectively) and ends with an individual making a formal commitment to adhere to the religion.¹⁷ So the analysis here can

¹⁷ Rambo (1993).

explain why these procedures are in fact rational. In addition, it can explain the perhaps surprising facts that commitments are sometimes necessary to staying the course even when one does not expect to be irrational in the future, and that it can make sense to make commitments even when one does not expect to believe these commitments were for the best in the future.

4. Faith and Integrity

Making a commitment and sticking with it is preferable, from the t_1 perspective, to allowing for the possibility of backing out, only if there is a possibility of genuinely misleading evidence—of evidence that makes one want to back out, from the t_2 perspective. (Otherwise the commitment would be inert.) Thus, in cases where faith must be exercised—cases that differ from just acting in accordance with one's credences at the time of action—there will be a conflict between the t_1 and t_2 perspectives. What should we make of this conflict?

First, it is worth saying that for purely epistemic purposes—for having accurate or appropriate credences without regard to which things are practically more important—we can allow that the t_2 perspective is superior. (We can allow this because nothing about the commitment to act requires that one do anything epistemically out of the ordinary.)¹⁸ Second, though, neither perspective is necessarily a better one *from which to make a decision*. If the evidence is in fact misleading, then the t_1 perspective is better, but if the evidence is not misleading, then the t_2 perspective is better, where a “better” perspective here is just a perspective from which a rational individual will make a decision that is apt to lead to a higher payoff for her.¹⁹ Given this, there is no neutral perspective the agent can adopt to choose between her t_1 preferences and her t_2 preferences. For the risk-avoidant individual, learning what the investigation did conclude won't necessarily put one in a better position to make a decision.

So, the individual potentially faces a conflict between the t_1 and t_2 perspectives, a conflict that cannot be resolved by any facts the individual has access to—“potentially” because one faces the conflict only if

¹⁸ One might worry that the t_2 perspective is not necessarily epistemically better, for reasons analogous to those that the t_2 perspective is not necessarily practically better. See Campbell-Moore and Salow (ms.). Note that the result that the t_2 perspective is not practically speaking superior rests on an asymmetry in the value of information: some propositions make more of a difference to one's choice than others do. See Buchak (2013: Ch. 6). So there would have to be a similar asymmetry in the value of information to credences to justify an analogous result.

¹⁹ See Buchak (2010) and Buchak (2013: ch. 6).

one receives the negative evidence. And if one thinks that such a conflict is likely occur, as it might be in the case of a long-term project where evidence of success comes in somewhat randomly, then adopting a policy of choosing according to whatever perspective one has at the time *might rule out* completing A at all. (In the case of Mary and Andy, if we were to assume that Mary will get succeeding bits of evidence that “bounce” her credence around a lot—the conclusion of the investigation, remarks by other colleagues of Andy’s, Andy’s demeanor when he talks about the accusations—so that there is a possibility of her credence being in the relevant ‘negative-but-inconclusively-so’ window, then it is likely that she will at some time have a low enough credence that continuing the partnership with Andy no longer makes sense.) Thus, when we adopt a *holistic* perspective on the decision—when we consider the individual over time—we notice that the only way to make it likely that one will complete A in the event that X in fact holds is to adopt a policy of letting one’s commitments make a difference in what one chooses, and to commit when it is rational to do A now but might not seem so in the future.²⁰

Thus, we can see why commitments make a difference—even if they are just commitments to oneself—and how faith, considered as a commitment to act even in the face of counter-evidence, makes long-term, risky projects possible. Faith allows an agent to make choices that are *integrated* over time, even if these choices sometimes seem irrational from her current perspective. Mary’s faith in Andy allows her to adhere to a single course of action even if she sometimes rationally entertains doubts about whether this course is best.

I want to be clear about how both credence and action go from the perspective of the individual after she has received the counter-evidence. As mentioned, we are assuming that the individual updates her credences in response to counter-evidence. (Indeed, it is not clear that she could do otherwise, since it is not clear that one can have credences that are different from those the evidence suggests to one.) However, her earlier commitment requires her to stick with her chosen action. Thus, the individual with counter-evidence will feel a tension between her credences and her actions. Faith in a proposition is primarily “active”—makes the individual’s behavior different than it otherwise would have been—in cases of doubt. And, again, since action is up to an individual even though credences are not, faith allows an individual to stay the course when engaged in a long-term risky project where the evidence that might come in meets the requirements discussed above.

²⁰ Paul (2015), in the course of arguing that individuals should exercise *self-control* in not reconsidering their decisions—a conclusion similar to that in this paper but argued for on a different basis—notes that self-control only makes sense for an agent who is unified over time, i.e. who can adopt the holistic perspective.

The (potential) divergence between action and credence at t_2 raises an objection to the idea that I've correctly described a sense of faith which is laudable. Although faith in this sense allows one to maintain integrity over time, one must sacrifice a different kind of integrity: integrity, at a time, between one's actions and one's credences. And, one might think, this kind of integrity is particularly important when we are dealing with fundamental moral and religious commitments, as in the cases here. In response, note that what the results here show is that one cannot always avoid making a choice between the two kinds of integrity. I want to emphasize that either choice is rationally permissible, for all I've said—but if one thinks the over-time narrative of an individual is more important than any of his time-slices, and notices that the divergence between credence and action at t_2 is apt to be temporary, then one can see that this is not so bad. If one wants the possibility of completing certain risky long-term projects, then one cannot always have at-a-time integrity, even though eschewing this integrity also opens one up to a certain kind of risk.

Interestingly enough, if the proposition that one has faith in is true, and if evidence is somewhat regular (if weight and credal resilience go together and if the evidence coming in isn't of a character to swamp previous evidence), then we can expect a characteristic pattern to emerge. The initial period of evidence-gathering is characterized by credences that are highly non-resilient—one can easily get negative evidence that (even from the prior point of view) makes it irrational to perform the risky act. The middle period is characterized by credences that are somewhat resilient—one can get negative evidence that would make it rational from the prior point of view to perform the risky act (because the evidence might be misleading) but would not make it rational from the later point of view to do so. This is the period in which faith is 'active' in telling the individual to do something that she wouldn't otherwise want to do, and this is the period in which faith is felt to be trickiest to adhere to. The last period, finally, is characterized by credences that are very resilient—the risky act will still be rational even if one encounters negative evidence. During this period, faith will be easiest to maintain and also least 'active'. So, in many cases, faith will not require one to act in a way that is at odds with one's current preferences indefinitely. One "passes through" active faith—faith that is more resilient than one's credences—to get to practical certainty. In these "regular" situations, faith primarily protects against misleading evidence in the middle-period, before becoming rational but inert.

5. The Nature and Value of Commitment

So we have an individual, Mary, who might find herself in the following situation at some time. She made a commitment to partner with Andy; if she hadn't made the commitment she would instead partner with Bob; and she chooses to continue her partnership with Andy because of her commitment. In the language of reasons, she has a reason given by her current credences to choose a partnership with Bob—doing so is the instrumentally rational action given her credences, not taking her commitment into account—but she also has a reason given by her commitments to continue the partnership with Andy, a reason which outweighs her reason to partner with Bob.

There are two ways to think about how Mary's commitment makes a difference to what she chooses. One way to think about it is that she prefers to partner with Bob, but nonetheless chooses to partner with Andy: she chooses *against* her preferences at the time, in order to respect her commitments. The other way to think about it is that she prefers to partner with Andy because of her commitment: her commitment *influences* her preferences, because she prefers to honor her commitments. Which story one is attracted to will depend on how one thinks about the relationship between preferences—typically represented by utility—and choices. If one thinks that utility captures one's preferences at a time before one employs substantive rules to get from what one prefers to what one should choose, then one will be attracted to the former picture. On this picture, preferences (utility) will provide one kind of reason among many, where others include moral reasons, reasons based on one's own commitments, and reasons based on what other people prefer. If one instead thinks that utility captures all that goes into making a choice—if the link between preference and choice is analytic or nearly so—then one will be attracted to the latter picture. On this picture, each reason will correspond to a utility value added to some action (e.g., that a friend prefers something increases its utility, that an action is immoral decreases its utility). Without taking a stand on which picture is ultimately better, I will adopt the latter, since it is more widely adopted among decision theorists.

On the latter picture, making a commitment amounts to transforming act D into:

	Found innocent & actually innocent (Positive correct evidence)	Found innocent & actually guilty (Positive misleading evidence)	Found guilty & actually innocent (Negative misleading evidence)	Found guilty & actually guilty (Negative correct evidence)
D': partner with Andy but revise choice according to new credences	9	-1	7 - d	7 - d

where d is the cost of backing out of the choice. Thus, if A is better than D from the t_1 perspective but might not seem so from the t_2 perspective, then a rational agent should want to transform D into D' —she should want to make an action strictly worse for herself to avoid being tempted by it. Indeed, this fact explains why it is beneficial for long-term risky commitments, like marriage or joining a convent, to involve public vows that it would be costly to break. The cost of breaking a commitment can't make *making* the commitment rational, but it can help the agent maintain integrity over time, in the sense discussed in the previous section.

Additionally, while the picture here accounts for the value of faith in instrumental terms—faith allows us to secure goods that we would be less likely to secure if we were blown about by the evidential winds—one might also hold that the committed life is intrinsically better than the non-committed life.²¹ If so, then instead of a picture on which the t_2 self grudgingly carries out a commitment of faith against her better judgment (grudgingly because from her perspective, she wishes the commitment hadn't been made), we have a picture on which the t_2 self is unsure whether the thing she's committed to is the best, but is happy that she is committed to something, and moreover should be happy that she was once in an evidential position in which commitment to that particular thing was rational. Moreover, if commitments become more valuable the longer they last, then over time she will get stronger and stronger reason to remain committed, and so will be able to endure stronger counter-evidence.

One might wonder how far a faith-commitment needs to go. I am not suggesting that it would be rational to stick with a commitment *regardless* of the evidence one encounters. In particular, if Mary comes across conclusive evidence of Andy's guilt, she ought to break the commitment. Commitment provides a strong reason in favor of an action, but it need not provide an overriding one, particularly if you come across information you didn't think you would get or if your initial reasons for having faith were undermined. This would be a case of losing or eschewing faith, and one's evidence might make it rational to do so. And importantly, this is not something one would expect to do when one makes the commitment in the first place.

²¹ See also Kelly (2004), who argues that it can be rational to honor sunk costs: to perform some action that one wouldn't perform if choosing 'anew', on the grounds that this action can bring it about that one's past efforts were not in vain.

6. Credences and Misleading Evidence

I've shown that it can sometimes be rational to fail to take new evidence into account when acting, because of the possibility that that evidence might be misleading. Could this point about not acting on what might potentially be misleading evidence already be taken into account in our credences and in at-a-time recommendations on the basis of them? To appreciate this question, notice that the advice I've given here can be stated in ordinary language: if ahead of time you think that there is a chance of getting evidence that will mislead you—enough of a chance that you want to avoid basing your decision on new evidence—then when you get the new evidence you should continue to recognize this possibility and not take the evidence into account. Why can't this advice also be stated in the language of credences and decisions there-upon, so that it is already built into the automatic recommendations of probabilism that the evidence is simply not taken into account when acting?

There are two potential responses. One is to hold that credence is the thing that follows the norms given by probabilism, and it is simply a surprising fact that these norms do not prevent one from falling afoul of evidence whose possible misleadingness one did not want to risk ahead of time. Credence is just about what the evidence at a given time reflects. (Note that in order even to formulate the idea of misleading evidence we need a distinction between what the evidence reflects and what is actually the case.) And it is a surprising fact that when we are making risky choices, different evidential positions suggest different courses of action, even when the evidence at one position is a superset of the evidence at another. It is a surprising fact that more evidence does not always lead to better action, so there is a genuine choice about which temporal point of view to ratify.

Of course, there might be ways to capture the phenomenon here (or similar phenomena), while adhering to the view that you ought to act on your credences and that credences-at-a-time capture all the relevant considerations. This brings us to the second response. A main point of formalism in philosophy is to precisify ordinary phenomena and use the precisification to discover facts that would have been unknown or contentious without the precisification; and then to translate the results back to ordinary phenomena. (These may be results which seem obvious once stated, but might not have been arrived at or believed without the formalism.) Regardless of the correct formal description, what I am ultimately interested in is in showing that there are situations with the following features: just doing what I have evidence for at a time (considered independently of my history as an agent) can be a worse strategy—from the point of view of obtaining what I care about—than ignoring some of my evidence when I act. Thus, it is good that we have an attitude—*faith*—that guards against precisely this mistake.

7. Belief

We have so far been framing our discussion in terms of credence, which represents the extent to which an individual takes a proposition to be supported by her evidence. But an attitude that is important to ordinary people—and to most philosophers—is *belief*. One might wonder how the diachronic requirement of faith discussed here bears on the relationship between faith and belief.²² There are a number of views about how belief is related to credence, and the answer to the question will depend on which view one adopts. While I don't want to take a stand on the correct view, the risky-commitment account of faith has particularly important upshots for two popular views of this relationship.

The first of these views holds roughly that what you believe is a matter of what you're willing to act on. Although details of this account differ, the basic idea is that if your credence in X is sufficiently high to justify accepting that X, for the purposes of action, then you believe that X.²³ If we accept this account and the picture of faith articulated here, then since faith that X requires a commitment to act on X, faith that X thereby requires a commitment to believe that X—and *a commitment to maintain one's belief even in the face of counter-evidence*. Again, such a commitment will be rational in the circumstances articulated above.

The second of these views sees belief as itself a kind of action, evaluable by decision theory.²⁴ This is perhaps easiest to appreciate by illustration. Following Levi (1967), consider an individual who is trying to form beliefs about which of three candidates (X, Y, or Z) will win an election. She can be committed to propositions of various logical strengths: for example, she could believe that X will win, or she could believe that X will win or Y will win. We can identify what she believes with the position she is committed to that is logically strongest. And in deciding what to believe, there are two desiderata, one corresponding to her goal of believing truth, and one to her goal of avoiding error: she wants to believe the logically strongest true proposition she can, but she wants to refrain from believing a false proposition. Levi represents this by assigning utility to having each belief in each state of the world,

²² See Buchak (2014) for a discussion of how faith and belief are related, given the synchronic requirement of faith.

²³ See Harsanyi (1985). Other versions of this view are articulated by Fantl and McGrath (2002), Weatherson (2005), Pace (2011), and Ross and Schroeder (2014). Two important dimensions that these views differ on are whether the individual who believes must explicitly adopt X in decision-making or simply act as if X holds, and whether pragmatic factors affect whether a credence suffices to *give rise to* a belief or whether a credence suffices to *justify* a belief.

²⁴ See Levi (1967). Other versions of this type of account include Maher (1993) and Kaplan (1996).

where utility is higher when a true belief is logically stronger but lower when a belief is false. A rational individual then selects the belief that maximizes expected utility (we revise this: maximizes risk-weighted expected utility). In a simple case in which there are only two propositions to believe (Andy is innocent, and Andy is guilty), we might have:

	Andy is guilty	Andy is innocent
Believe that Andy is innocent	0.9	-0.1
Believe that Andy is guilty	-0.1	0.9
Believe that {Andy is innocent or Andy is guilty}, i.e. suspend judgement	0.7	0.7

And, if we consider the question of whether to decide what to believe now or to look for more evidence, and whether to commit or to be open to revising in response to new evidence, we have a decision analogous to that discussed above (omitting for sake of brevity the options that correspond to believing that Andy is guilty):

	Found innocent & actually innocent (Positive correct evidence)	Found innocent & actually guilty (Positive misleading evidence)	Found guilty & actually innocent (Negative misleading evidence)	Found guilty & actually guilty (Negative correct evidence)
A: Believe that Andy is innocent and continue to believe even if counterevidence arises	0.9	-0.1	0.9	-0.1
B: Suspend judgment	0.7	0.7	0.7	0.7
C: Gather evidence, update, and then decide what to believe based on new credences	$0.9 - c$	-0.1	0.7	0.7
D: Believe that Andy is innocent but revise choice according to new credences when new evidence comes in	0.9	-0.1	0.7	0.7

Notice that the only relevant values here are the truth and strength of what is believed—this is not a case like Pascal’s Wager, in which believing one proposition if true is more valuable than believing its negation if true.

As above, it will be rational to believe that Andy is innocent and (from the initial perspective) to maintain this belief in the face of counter-evidence if the credence conditions obtain and one is risk-avoidant.²⁵ This is to say: belief as an act of faith requires a commitment to believe without looking for more evidence, even in the face of current uncertainty and future counter-evidence, and such a commitment can sometimes be rational.

So, on both of these pictures, faith requires stopping the search for evidence, and committing to maintain a belief even in the face of counter-evidence. And doing so can be rational in the conditions articulated earlier in this paper.²⁶ If we accept one of these two pictures of belief, then the tension between credence and action at a time becomes a tension between credence and belief: the weight of the evidence at a time can be at odds with what one accepts for the purposes of action or what one believes for purposes of holding the strongest true beliefs one can.

In addition to making good on a folk theory about faith—that it requires maintaining a belief in the face of counter-evidence—this upshot can be used to respond to an additional worry that arises from the divergence of credence and action, namely that faith in the context of personal relationships sometimes requires a cognitive stance, not just action. Consider Andy’s friend Sarah, who is not considering partnering with him, but does form beliefs about him. Andy denies having committed fraud, and Sarah has faith in him. The report has not yet come out, and it may contain evidence of Andy’s guilt. If faith merely requires a commitment to action, then Sarah can remain agnostic about Andy’s guilt (though she would have to be willing to act if there were relevant acts to perform). And this seems to miss something important in the value of faith in the testimony of people who are close to you. In response, if we couple the picture here with the second picture of belief above, then faith requires something stronger: that Sarah *believe* in Andy’s innocence despite the potential counter-evidence, and would continue to believe in his innocence even if she lowered her credence in response to actual counter-evidence. Again, the

²⁵ In addition, it will be rational to believe that Andy is innocent without looking for more evidence, if the credence conditions obtain and either there are costs to postponing a decision about what to believe or one is risk-avoidant.

²⁶ Though notice that if we accept the Levi-type view, then the exact situations in which acting on faith and believing on faith are rational may come apart, since the risks might not be exactly the same.

explanation for the rationality of such faith is that believing Andy to be innocent is a risk that Sarah could be lured away from taking by misleading evidence. Thus the role in Sarah's beliefs played by Andy being her friend isn't that her values (loyalty etc.) affect her beliefs but rather that her friendship is a concomitant of the fact that Sarah has strong and deep enough evidence to form a committed belief in Andy's innocence.

8. Conclusion

Previous work of mine has concentrated on the synchronic requirement of faith that one act without looking for further evidence. This paper has explored the diachronic requirement of faith that one remain committed to acting regardless of what evidence comes in. Which of these aspects is more central may depend on the particular case in question, though both are important to faith. Furthermore, both of these requirements can be rational to adhere to, from an individual's perspective before encountering evidence. Even though the commitment might make things worse from the individual's perspective after encountering evidence, there is no view from nowhere to assess which perspective is better—and since one can only guard against misleading evidence when carrying out long-term risky projects if one commits from the initial perspective, faith makes an important contribution to human life, by making it possible for us to carry out such projects.

Thus, for a rational individual, the process of acquiring and maintaining faith looks like this. One first explores a question in detail, gathering whatever evidence one can and updating one's credences on the basis of that evidence. If the evidence gathered has a certain character—if it tells strongly and deeply in favor of a particular proposition—and if future evidence is not likely to be conclusive, then one adopts faith in the proposition in question. In adopting this faith, one is willing to take risks on the proposition without first seeking out more evidence, and to continue to do so even if counter-evidence comes in. The downside of such faith is that once adopted, it is more difficult to eschew it if the proposition is false. But the upside of such faith is that if the proposition one has faith in is true, then one can successfully complete long-term projects that depend on its truth while withstanding the changing winds of evidence.

ACKNOWLEDGEMENTS

I am grateful for the insights of Nate King, John Pittard, Ryan Preston-Roedder, and Dan Speak. I am also grateful for comments and questions from audiences at Georgetown, Baylor, the Association for the Philosophy of Judaism, the Pacific-Mountain SCP 2013, Westmont, UC Merced, the Midwest SCP 2015, the Faith Summer Seminar 2015, the Pacific APA 2016, Azusa Pacific, and UCLA.

BIBLIOGRAPHY

- Buchak, L. (2010), 'Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering,' *Philosophical Perspectives* 24: 85–120
- (2012), 'Can it be rational to have faith?', in J. Chandler and V.S. Harrison (eds.), *Probability in the Philosophy of Religion* (Oxford): 225-247.
- (2013), *Risk and Rationality* (Oxford).
- (2014), 'Rational Faith and Justified Belief,' in T. O'Connor and L. Goins (eds.), *Religious Faith and Intellectual Virtue* (Oxford), 49-73.
- Calvin, J. (1536), *Institutes of the Christian Religion*.
- Campbell-Moore, C. and B. Salow (ms.), 'Avoiding Risk and Avoiding Evidence'
- Fantl, J. and McGrath, M. (2002), 'Evidence, Pragmatics, and Justification', *The Philosophical Review* 111(1): 67-94.
- Harsanyi, J.C. (1985), 'Acceptance of Empirical Statements: A Bayesian Theory Without Cognitive Utilities', *Theory and Decision* 18 (1):1-30.
- Howard-Snyder, D. (2013), 'Propositional faith: what is it and what it is not', *American Philosophical Quarterly* 50(4): 357-372.
- Joyce, J.M. (2005), 'How Probabilities Reflect Evidence', *Philosophical Perspectives* 19(Epistemology): 153-178.
- Kaplan, M. (1996), *Decision Theory as Philosophy* (Cambridge).
- Kelly, T. (2004), 'Sunk Costs, Rationality, and Acting for the Sake of the Past', *Noûs* 38(1): 60-85.
- Keynes, J.M. (1921), *A Treatise on Probability* (London: MacMillan).
- Levi, I. (1967), *Gambling with Truth* (Alfred A. Knopf and Routledge & Kegan Paul).
- Maher, P. (1993), 'The Concept of Acceptance', *Betting on Theories* (Cambridge), 130-161.
- Pace, M. (2011), 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' "Will To Believe"', *Noûs* 45 (2):239-268.

--- (ms.), 'Oh Ye of Little Faith: On Measuring and Evaluating the Strength of Faith and Trust'

Paul, S. (2015), 'Doxastic Self-Control', *American Philosophical Quarterly* 52(2):145-58.

Skyrms, B. (1980), *Causal Necessity* (Yale University Press, London).

Rambo, L.R. (1993), *Understanding Religious Conversion* (Yale University Press).

Ross, J. and Schroeder, M. (2014), 'Belief, Credence, and Pragmatic Encroachment', *Philosophy and Phenomenological Research* 88(2): 259-288.

Weatherson, B. (2005), 'Can We Do Without Pragmatic Encroachment?', *Philosophical Perspectives* 19: 417-43.

Wolterstorff, N. (1990), 'The Assurance of Faith', *Faith and Philosophy* 7(4): 396-417.