

When is Faith Rational?

Lara Buchak

Forthcoming in *Norton Introduction to Philosophy 2nd edition* (eds. Alex Byrne, Josh Cohen, Liz Harman, Gideon Rosen).

Can it be rational to have faith? In order to answer, we need to know what having faith actually means. In the first section of this essay, I give an analysis of faith. Faith is not ‘belief without reasons’. Rather, to have faith in a claim is to be willing to *act* on that claim without further evidence, and to remain committed to acting on that claim even when counterevidence arises. In the remainder of this essay, I argue that faith can be rational in certain circumstances, and that those who lack faith stand to miss out on important goods.

1. What is Faith?

Let’s begin with a few preliminaries.¹ First, this essay is about faith both in the mundane sense (you have faith that a friend is trustworthy or that a bridge will hold your weight) and in the religious sense (you have faith that God exists, or that God is asking you to take some action). The account explains religious and mundane faith as instances of the same general attitude.

Second, although we talk both about what is it to have faith *in* someone and faith *that* something is the case, this essay will only give an account of the latter. (They are, of course, related: to have faith in someone requires having faith that some things are the case; for example, to have faith in a friend requires having faith that she is trustworthy.) Notice that only certain claims are candidates for faith at all. For example, while it makes sense to say that you have faith that a friend will quit smoking, it does not make sense to say that you have faith that a friend will continue smoking (assuming you think that smoking is bad). It is not that you *lack* faith that your friend will continue smoking; rather, this claim is not an appropriate object of faith at all. You can only have or lack faith in a claim that you have a positive attitude towards. Similarly, it does not make sense to say that you have (or lack) faith that $2+2 = 4$. You

¹ This essay draws heavily on L. Buchak (2012), ‘Can It be Rational to have Faith?’, *Probability in the Philosophy of Religion*, eds. J. Chandler and V.S. Harrison (Oxford: Oxford Univ Press), 225-247; and L. Buchak (2017), ‘Faith and Steadfastness in the Face of Counter-Evidence’, *International Journal for Philosophy of Religion* 81(1-2): 113-133. See also L. Buchak (2014), ‘Rational Faith and Justified Belief,’ *Religious Faith and Intellectual Virtue*, eds. T. O’Connor and L. Goins (Oxford: Oxford Univ Press), 49-73.

can only have or lack faith in a claim if you're not certain of the claim on the basis of your evidence alone: your evidence must leave it open that the claim is false.

Claims that will be particularly good candidates for faith, then, are claims concerning interpersonal, religious, or moral matters. Take the claim that my friend will keep my secrets. I care a great deal that this claim is true, and my evidence will always leave it open that the claim is false, since she could at any time decide not to keep my secrets. Or take the claim that God exists. Many people have a positive attitude towards this claim—they hope it is true—and while they may have some evidence in its favor (evidence from their experiences within a religious community, direct personal experiences, philosophical arguments, or the testimony of others), all of these sources typically leave room for doubt. Similar points hold for particular instances of faith in Jewish and Christian texts: Abraham's faith that his descendants would be greater than the number of stars; Moses's faith that God would lead his people out of Egypt into the Promised Land; Paul's faith that Jesus was raised from the dead.

Preliminaries out of the way, we can now turn to the account of what it is to have faith that something is the case. Some people think that to have faith is to believe something without any evidence at all, or even to believe something when the evidence is squarely against it. But these ideas don't fit with the way in which 'faith' is used, both in interpersonal and religious contexts. Moses had some evidence that God would lead his people into the Promised Land, and yet he still had faith that this was so. So what is it to have faith?

Start with the observation that faith is tied to action. To have faith in a claim is to be willing to take risks on that claim. It is risky to tell your friend a secret, to dedicate your life to following God's commands, to leave Egypt to seek the Promised Land. You have a higher degree of faith in a claim to the extent that you are willing to take more risks on that claim.

But not every case of risk-taking will be an act of faith. Faith requires a willingness to act on the claim you have faith in without first looking for additional evidence. Furthermore, faith requires a willingness to continue to act, even if counter-evidence arises. In the Christian Bible, when Jesus called Simon Peter and Andrew, they immediately stopped what they were doing to follow him—and Jesus's disciples continued to remain committed to him even after his death gave them some reason to think he was not the messiah they were looking for. (Simon Peter and Andrew were Jews, and 1st century Jews did not expect their messiah to die.) Mundane faith works the same way. If you have faith that a particular person would be a good marriage partner, you will marry him and you will stick with him even if the marriage

isn't what you expected. If you have faith that you will complete your college degree, you will continue in the program even after failing an exam.

Thus, you have faith in a claim if you are committed to taking risks on the claim without examining additional evidence and would maintain that commitment even if counter-evidence arose. More formally:

A proposition X is a *candidate for faith* for a person S if S cares that X holds and is uncertain that X holds on the basis of his evidence alone.

S *has faith* that X if and only if:

- (1) X is a candidate for faith for S.
- (2) S is willing to take a risk on X without looking for additional evidence; and
- (3) S is willing to follow through on such risky actions even when he receives evidence against X.

To the extent that S is willing to perform riskier acts that express his faith, and remain steadfast in the face of stronger counter-evidence, he has a higher degree of faith that X.

To tell your secrets to a friend is to take a risk on the claim that she will keep your secrets: you gain something if she keeps the secrets and lose something if she doesn't. Whether you have faith in this claim is a matter of whether you are willing to tell your secrets without first double-checking that your friend is trustworthy and to continue to tell your secrets even if you overhear a piece of gossip that your friend is bad at keeping secrets.

So too for religious faith. People with faith that God is real act in ways that make sense on this assumption without waiting for more evidence, and they continue to act in these ways even when they encounter evidence and arguments against their belief. Someone might take these ways of acting to show that people with religious faith are dogmatic and irrationally confident in their beliefs; but on the present view, that is not what faith consists in. As we will see, you can be disposed to *act* faithfully in these ways even if your underlying beliefs are non-dogmatic and responsive to evidence.

2. What is Rationality?

Now that we know what faith is, we can ask whether faith is rational.

There are two important senses of rationality that are relevant to the main question of this essay: rationality in what you believe, and rationality in what you do. Philosophers use formal tools for thinking about each kind of rationality.

In the case of rationality in what you believe (*epistemic rationality*), philosophers commonly think of belief not as an all-or-nothing phenomenon, but instead as a matter of degree. Your degree of belief in a claim—sometimes called your *credence*—can be thought of as the probability you assigns to that claim. Degrees of belief must obey two requirements to count as rational. First, they must cohere with each other. You can't believe to degree 0.8 that it will rain and also believe to degree 0.8 that it won't rain. (If you're rational, your credences in *X* and *not-X* must add up to 1.) If you believe to degree 0.5 that it will rain tomorrow, then your credence that it will *either rain or snow* must be at least 0.5. (If you're rational, your credence in *X or Y* must be at least as great as your credence in *X*.)² Second, your degrees of belief must fit with the evidence you possess: if you live in a sunny climate and have no reason to suppose that today will be any different from previous days, then you cannot believe to a high degree that it will rain. Finally, your degrees of belief must be sensitive to new evidence: when the weatherman predicts rain you must raise your degree of belief that it will rain. (If you're rational, then if you get some good evidence for *X* that you didn't have before, you should take it into account by raising your credence in *X*, and lowering your credence in *not-X*.)

In the case of rationality in what you do (*practical rationality*), the rough idea is that to be practically rational, you must take the means that you believe will lead to your ends. Taking on board the above idea that belief is a matter of degree, philosophers commonly hold that practical rationality is formalized by *decision theory*. According to decision theory, we can represent how much you value particular outcomes by a 'utility' function that assigns high values to things you want a lot and lower values to things you want less. On the classical view, a rational individual must choose the act that maximizes **expected utility**, relative to his utility values and degrees of belief. Here 'expected' utility just means a weighted average: the utility value of each outcome is weighted by the likelihood of getting that outcome, and the

² For an overview of some of the concepts in this section, see M. Resnik (1987), *Choices* (Univ of Minnesota Press).

resultant values are summed.³ For example, if you think that the probability of rain is 0.5 and of not-rain is 0.5, and if the utility of getting wet is -3, the utility of staying dry but having to carry an umbrella is 1, and the utility of staying dry without having to carry an umbrella is 6, then the expected utility of not bringing an umbrella is $0.5(-3) + 0.5(6) = 1.5$ and the expected utility of bringing an umbrella is $0.5(1) + 0.5(1) = 1$. So decision theory says that a rational agent with these credences and utilities will leave his umbrella at home.

In addition to the classical view, we will be concerned with a variant on the classical view that takes into account how an individual treats risk. According to the classical view, you simply average over all the possible utility values. This means that all equally likely scenarios weight equally in your evaluation of an action. However, contra the classical view, it is plausible that you might care about making sure that things won't go too poorly, so that in thinking about the value of an act, you place more weight on what happens in worse scenarios; or someone else might care a great deal about making sure there's some chance that things go really well, so that he places more weight on what happens in better scenarios.⁴ If you give more weight to worse scenarios than better ones, we will say that you are *risk-avoidant*.

3. When is it Rational to have Faith?

Anna has known Bates for a long time, and he has proposed marriage.⁵ She has abundant evidence that he is morally upstanding; thus she believes to a high degree that he is not capable of murder. However, he is currently under investigation for murder. If he is innocent, then he would make an excellent husband, but it would be a disaster to marry a murderer. She could accept his proposal, or decline it and continue her fairly happy life, which is better than marrying a murderous Bates but worse than marrying an innocent Bates. (We may assume that marrying an innocent Bates is preferable even if Bates is *found*

³ For a brief introduction to decision theory, see Alan Hájek, "Pascal's Ultimate Gamble," in Chapter 2 of this anthology.

⁴ See L. Buchak (2013), *Risk and Rationality* (Oxford: Oxford Univ Press). In the example in the previous footnote, if you place twice as much weight on what happens in the worst 50% of states than the best 50% of states, then the *risk-weighted expected utility* of not bringing an umbrella is $0.67(-3) + 0.33(6) = 0$ and of bringing an umbrella is $0.67(1) + 0.33(1) = 1$. Thus, you should bring your umbrella. I note that the idea that risk-avoidance is rational is controversial. Still, risk-weighted expected utility maximization is more descriptively accurate than expected utility maximization, so if one does not think that it is rational, then one can read this essay as a comment on how to act if one is in fact (irrationally) risk-avoidant.

⁵ This example comes from the television show *Downton Abbey*.

guilty and thrown in jail—his actual innocence is the basis for the risk Anna would take in marrying him.) If she has faith that Bates is innocent, what will she do? She will agree to marriage before the investigation is concluded, and stick with the decision even if the investigation finds him guilty.⁶

Anna's decision has the same structure as many cases of religious faith. Moses, we are to suppose, has interacted with God personally. He has ample reason to believe that God is trustworthy and cares about his people. God has told him to journey out of Egypt with his people. If such a journey would indeed be successful—if God will indeed lead his people to the Promised Land—then embarking on it would be the best thing for Moses and his people, but if it would not be successful, then embarking on it would be a disaster. He could embark on the journey, or simply stay in Egypt. Moses had faith that God would lead his people to the Promised Land, so what did he do? He embarked on the journey without first verifying that it would be successful, and he stood firm even when the Egyptians pursued his people to the Red Sea.

The key question for this section is whether such faith is rational.

3.1 Epistemic Rationality

Let us begin with epistemic rationality. To be epistemically rational, Anna needs to have degrees of belief that are coherent and supported by the evidence. If Anna's evidence supports a high degree of belief in Bates's innocence, and if her other degrees of belief cohere, then she will indeed be epistemically rational. So too with Moses: if his evidence indeed supports a high degree of belief in the truth of God's utterances⁷—in this case, that God will lead his people to the Promised Land—and if his other degrees of belief cohere, then he will be epistemically rational.

Faith does not require either individual to alter their evidentially-supported degrees of belief. More generally, faith does not require you to do anything 'special' with your beliefs: it does not, for example, require you to believe a claim more strongly than the evidence suggests, or believe a claim in the absence of any evidence at all. Furthermore, faith does not require you to refrain from looking for evidence for purposes of knowing more about the world, nor does it ask you to ignore new evidence in regulating your degrees of belief. It simply means that your *decisions* won't depend on what new evidence comes in.

⁶ This is not to say that a faithful Anna will stick with her decision *come what may*. If the investigation provides *conclusive* evidence that Bates is a vicious killer, all bets are off. See section 3.3 below for more on this.

⁷ Or in the claim that the utterer really is God.

Thus, as long as you meet the criteria for epistemic rationality—as long as your beliefs are coherent and respect the evidence—you can have faith while remaining epistemically rational.

Anna’s faith in Bates’ innocence, for example, does not require her to ignore the evidence of his guilt as it comes in. Her degree of belief in his innocence may drop as the damning evidence accumulates. Her faith simply requires her to stick by her decision to marry him nonetheless: acting on the premise that he is innocent even when this action seems too risky given her degree of belief. Similarly for Moses: his faith does not require him to ignore evidence that the journey will be unsuccessful. When he approaches the Red Sea, his degree of belief that they will make it through alive might drop. His faith requires him to stay the course nonetheless.

As I’ve been emphasizing, faith is not belief without reasons. It is a commitment to action, even when you’re not sure, and to seeing the action through, even when doing so seems too risky. Thus, faith is perfectly compatible with having beliefs that respect the evidence.

3.2 Practical Rationality at the Beginning of an Act of Faith

The question of practical rationality is more interesting, since faith requires doing two things: taking a risk without looking for more evidence, and sticking with that risk even if evidence arises that makes it no longer seem like a good one. We will ask about each of these elements of faith separately: when is it practically rational to stop your search for evidence and make a decision, and when is it practically rational to stick with a decision even though new evidence no longer supports it.

Let us begin with the question of looking for further evidence. I will use the Anna/Bates example, but the reader is invited to keep in mind that analogous claims hold for Moses. To see whether Anna should answer Bates’s proposal before hearing the verdict, we compare her three options, and the utility payoffs they yield under the four possible states:⁸

⁸ When assigning utility numbers, the actual numbers don’t matter, just the relationship between them: so what these numbers represent is that Anna’s life without Bates would be nearly as good as her life married to an innocent Bates, and much better than her life married to a murderous Bates.

| | Bates is actually innocent & the investigation finds him innocent (Positive correct evidence) | Bates is actually innocent & the investigation finds him guilty (Negative misleading evidence) | Bates is actually guilty & the investigation finds him innocent (Positive misleading evidence) | Bates is actually guilty & the investigation finds him guilty (Negative correct evidence) |
|---|---|--|--|---|
| A: say yes to Bates now *the faithful act* | 9 | 9 | -1 | -1 |
| B: say no to Bates now | 7 | 7 | 7 | 7 |
| C: wait for the verdict and then decide | | | | |

I've labeled the states according to whether the evidence is *positive* or *negative*—whether it tells in favor of Bates's innocence or not—and whether it is *correct* or *misleading*—whether it tells in the direction of the truth. Let us assume that Anna's current degree of belief that Bates is innocent is high enough that A has a higher utility than B (otherwise, faith in the claim will already be irrational).⁹

How we fill in the last line of the matrix will depend on what Anna would do upon hearing the results of the investigation, which will depend on how she expects the investigation to affect her degrees of belief. If she thinks the evidence from the investigation will be weak compared to the rest of her evidence, either because she already has a *lot* of evidence about Bates's character or the investigation itself is flawed (sloppily conducted or hopelessly corrupt), then she knows that even if she were to learn that the investigation found Bates guilty, this will not lower her degree of belief in his innocence very much. Thus, if she waits for the verdict, then whatever it is, she will still say yes to Bates, and get the corresponding payoffs, minus whatever cost c there is to postponing the decision (not getting to be married as quickly, Bates's disappointment in her dithering, etc.):

⁹ For example, if Anna is an expected utility maximizer, then she believes to degree greater than 0.8 that Bates is innocent.

| | Bates is actually innocent & the investigation finds him innocent (Positive correct evidence) | Bates is actually innocent & the investigation finds him guilty (Negative misleading evidence) | Bates is actually guilty & the investigation finds him innocent (Positive misleading evidence) | Bates is actually guilty & the investigation finds him guilty (Negative correct evidence) |
|---|---|--|--|---|
| A: say yes to Bates now *the faithful act* | 9 | 9 | -1 | -1 |
| B: say no to Bates now | 7 | 7 | 7 | 7 |
| C: wait for the verdict and then decide | $9 - c$ | $9 - c$ | -1 | -1 |

In this case, A (the faithful act) and C will yield the same payoffs, except that C might be worse if there are postponement costs. So, if there are no postponement costs, then A and C are equally choiceworthy, and if there are, then A is to be preferred to C.

Thus, our first result is that *faith is rationally permissible if you know that whatever negative evidence you get, it won't be strong enough to discourage you from taking a risk on the faith-claim—and faith is rationally required if, in addition, there are costs to postponing the decision.* If your evidence for a claim is already deep and convincing—even if it doesn't yield certainty—then you *may* have faith in that claim, and indeed you *must* have faith in that claim if there are costs to postponing the decision.

The more interesting case is the one in which a negative result from the investigation would lower Anna's degree of belief enough that marrying Bates would no longer be the utility-maximizing option. In this case, waiting for the verdict will lead to saying yes to Bates if the verdict is positive (innocent) and to saying no if the verdict is negative (guilty); therefore, the payoffs for waiting for the verdict are as follows:

| | Bates is actually innocent & the investigation finds him innocent (Positive correct evidence) | Bates is actually innocent & the investigation finds him guilty (Negative misleading evidence) | Bates is actually guilty & the investigation finds him innocent (Positive misleading evidence) | Bates is actually guilty & the investigation finds him guilty (Negative correct evidence) |
|---|---|--|--|---|
| A: say yes to Bates now *the faithful act* | 9 | 9 | -1 | -1 |
| B: say no to Bates now | 7 | 7 | 7 | 7 |
| C: wait for the verdict and then decide | $9 - c$ | 7 | -1 | 7 |

Under what circumstances should A be chosen rather than C? Notice that there are two potential benefits to saying yes to Bates rather than waiting. The first is that if the evidence is positive and correct, doing so yields 9 rather than $9 - c$: this is the *potential benefit of acting now rather than later*. The second is that if the evidence is negative and misleading, doing so yields 9 rather than 7: this is the potential benefit of *not backing out on bad information*. On the other hand, there is also a potential benefit to waiting. If the evidence is negative and correct, doing so yields 7 rather than -1: this is the potential benefit of *backing out on good information*. How these benefits trade off against each other will depend on the likelihood of each of the states in which the benefits occur. And, using decision theory, we can say exactly when the benefits of saying yes to Bates now outweigh the benefits of waiting for more evidence:¹⁰ Consider the following version of the case:

- (1) Anna is already fairly confident that Bates is innocent.
- (2) If the investigation says that Bates is guilty, then (even though it will lower Anna's degree of belief enough to prefer saying no) it won't tell conclusively in favor of his guilt.
- (3) Postponing the decision to marry Bates is costly OR Anna is risk-avoidant.

Call conditions (1) and (2) the *credence conditions* and the elements of condition (3) the *value and attitude conditions*. When these conditions obtain, Anna is *rationally required* to have faith that Bates is innocent: she is rationally required to stop her search for evidence and decide to marry Bates.

¹⁰ For details, see Buchak (2012) and L. Buchak (2010), 'Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering', *Philosophical Perspectives* 24: 85-120.

On the side of the faithful act, we have the potential benefit of acting now (you might take the risk anyway, when it would have been at least as good to take it immediately) and the potential benefit of not backing out on bad information (you might be talked out of a risk that would have paid off); and on the side of waiting until more evidence comes in, we have the potential benefit of backing out on good information (you might be talked out of a risk that would not have paid off). If the cost of postponing the decision is high, then the potential benefits of the faithful act are higher, particularly if you are likely to get positive correct evidence—which is more likely if you already have a large body of positive evidence. If you are risk-avoidant, then you require a higher degree of belief to take a risk; therefore, you are more easily talked out of taking a risk. Consequently, if the negative evidence wouldn't be conclusive, so that backing out on bad information and backing out on good information are both possible, you are in particular danger of the former. Negative evidence is less likely to be conclusive if the source of the evidence is unreliable or if you already have a large body of positive evidence.

Putting these facts together, the benefits of the faithful act are apt to outweigh the benefits of waiting for more evidence when you have a large body of evidence in favor of the proposition you are considering taking a risk on, the evidence that you might still encounter is sparse or unreliable, and either postponing the decision is costly or you are risk-avoidant. When your evidential situation has these features, the benefits you stand to lose by getting more evidence (you won't act now, you might back out on bad information) outweigh the benefits you stand to gain by getting more evidence. We might say: faith guards against unnecessary dithering and misleading evidence.

Again, the credence conditions are more apt to obtain if Anna antecedently has a large body of evidence in favor of Bates's innocence. For example, if Anna has a long history of observing Bates's character, then even if the investigation says he is guilty, her degree of belief in his innocence will only drop somewhat, since the guilty verdict must be weighed against her own large body of evidence. Furthermore, if her evidence comes from many sources, then her degree of belief is more likely to remain high in the face of counter-evidence. So, we can say more generally: a larger and more diverse body of evidence is more apt to make faith practically rational. An individual who commits prematurely is not rational, *but neither is an individual who insists on seeing all the evidence before committing.*

If Anna has no experience with Bates, it would be irrational for her to have faith in his innocence, but since she does, it would be irrational for her not to have faith. If we accept the supposition that Moses had a long history of trusting God and had even interacted with him personally, then it would have been

irrational for him not to have faith that God would lead his people to the Promised Land. If you have no evidence that God exists or no experience with a particular religion, then it would be irrational for you to have faith that God exists or to have faith in the claims of that religion,¹¹ but if you've had a series of religious experiences or if you are embedded in a religious community that has in your experience led to truth and flourishing, it would be irrational for you not to have faith. Similar points holds of reliance on testimony, which can be considered a special case of faith (faith that some testifier is telling the truth): it might not be rational to take the testimony of a stranger on faith, but it will be rational to take the testimony of someone who you know has a good track record or about whom you know other facts which imply that he has access to the truth and honestly shares it. You must get evidence first, but you must not postpone acting forever.

3.3 Practical Rationality during an Act of Faith

Let us now turn to the other requirement of faith: if Anna has faith that Bates is innocent, then she must stick with her decision to marry him, even if the verdict is that Bates is guilty. (If Moses has faith that God will lead his people to the Promised Land, then he must not turn back when Pharaoh's army pursues him.) Again, this case will only be interesting if a guilty verdict lowers her degree of belief enough to make her not want to risk marrying Bates. When will sticking with her decision anyway be practically rational?

This case has roughly the same features as the 'more interesting case' above, with one twist. If the credence and attitude conditions are satisfied, then *before* Anna gets the evidence it will be rational for her to plan to stick with her decision no matter what the evidence says.¹² But if she learns that the verdict is that Bates is guilty, then she will, *upon hearing this*, want to back out of her commitment, and rationally so (in the sense that this action will be recommended by her new degrees of belief, given her utilities). Thus, there might be a conflict between what Anna should plan to do at an earlier time and what it would be rational for her to do at a later time after new information has come in, at least if we set

¹¹ This stands in contrast to claims made by Pascal, 'The Wager,' and James, 'The Will to Believe' (both excerpted in this volume), at least on plausible interpretations of these authors.

¹² The choice between sticking with her decision and backing out is represented by the choice between A and C when $c = 0$. (There will be no costs to postponing the decision, because she has already made it, and all four possibilities are live, because we are asking about what, *before* she gets the evidence, she should plan to do in the future.)

her past plan aside. It looks like faith will tell Anna to stick with a plan it was rational for her to make given what she knew then, but which it would not be rational for her to make now given what she's learned. *How can that ever be rational?*

Recall one feature that made the don't-wait-for-more-evidence aspect of faith rational: faith guards against the possibility that misleading evidence will make you want to back out of a course of action that would actually be best. Such a safeguard is needed when you are apt to encounter evidence that will make you less sure of the claim you are taking a risk on, but will still leave you in a position in which that claim is highly probable. And you are more apt to encounter evidence of this form when you are engaged in a long-term risky project where a variety of 'low-quality' evidence comes in that pushes your degree of belief up and down. When you are engaged in a project on a long enough time-scale, and with irregular enough evidence, you are apt to encounter misleading evidence. Thus, this feature also explains why remaining committed, despite counter-evidence, is rational.

Take the claim that you will successfully complete the degree and training required for your career. Completing a degree takes years, and you can't be completely certain ahead of time you will finish (in part because doing so depends on your own efforts). You may begin with a high degree of belief—based on good evidence—that you will complete the degree. But during the time required to complete it, evidence will come in that will knock your degree of belief around. You do well as a freshman, and your degree of belief rises; you struggle as a sophomore and your degree of belief plummets; you are praised by a professor and your degree of belief rises again; and so forth. You won't be able to complete the degree if you drop out whenever your degree of belief is low enough to justify your doing so. You won't be able to complete it unless your commitment to completing it is resilient in the face of counter-evidence.

Or consider again Moses's faith in the claim that God will lead his people to the Promised Land. Pharaoh at first refused to let Moses's people go; then he agreed; then he took back his agreement; then he agreed again to let them go; then he pursued them with his army; then God parted the Red Sea; then there was no food in the desert; then God gave them manna to eat; and so forth. Moses wouldn't have been able to lead his people to the Promised Land if he turned back whenever his degree of belief was low enough to justify doing so—he needed a resilient commitment to completing the journey.

Other cases have this structure as well. Consider the claim that your efforts towards justice will succeed, and the risk of dedicating your life to such an effort: anyone who pursues justice will likely experience

setbacks that make her reasonably doubt whether her efforts will succeed. Or consider the religious adherent: he can expect to reasonably doubt when he encounters setbacks, when his expectations are thwarted, or when he encounters reasonable disagreement with an adherent of a different religion.

Even if you are rationally confident that the faith-claim is true when you embark on a long-term risky project, if the project involves a long enough time-scale, then you should ahead of time expect to encounter misleading evidence at some point. You can expect, even if the faith claim is in fact true, that at some point you will rationally doubt enough to want to back out of the project. Given this, in order to have the possibility of completing such a project at all your commitment must be resilient: you must not waver when your degree of belief drops.

We might have thought that a rational person is someone who always does what makes most sense in light of his current information and his current preferences. But as these examples show, if you adopt a policy of always acting on your degrees of belief at a given time—if Anna adopts a policy according to which she will back out of her commitment to Bates if the verdict says he is guilty—then you will likely not be able to complete long-term risky projects at all, because you can expect to encounter misleading evidence at some point. If instead, you adopt a policy of making commitments and following through on them—if Anna maintains faith in Bates—then you will be able to complete these projects *if* the faith-claim is true. Thus, faith allows us to complete a risky long-term project in a way that is decoupled from evidence that the project will fail or isn't worth it. *Faith keeps us from being blown about by the changing winds of evidence.*

Of course, you do lose something by adopting faith: you leave yourself open to acting at particular times against what your evidence suggests, and so leave yourself more open to taking a bad risk. It is important to note that while faith recommends sticking with an action in the face of counterevidence, there are mechanisms for losing faith: if your degree of belief in the faith-claim drops drastically, or stays low for a long period of time, or if your initial reasons for adopting faith are undermined, then it may be wise, all things considered, to give up the faith-commitment. Rational faith is not blind, dogmatic faith. Although someone who initially has faith should maintain that faith as the evidence starts to go against the claim in question, it can be rational to change one's mind if negative evidence keeps piling up. At some point it could be rational for Anna to give up on Bates, and maybe even for Moses to give up on God. But in the meantime, especially if they start out with convincing and deep evidence, their faith can withstand a lot of negative evidence.

4. Conclusion

Faith is not belief without evidence or reason. To have faith in a claim is to be willing to take risks on that claim without looking for more evidence, and to continue to take these risks even if you encounter evidence against the claim.

Such faith can be rational. In particular, embarking on an act of faith—acting without gathering more evidence—is rationally required when your evidence is strong enough to warrant a high degree of belief in the claim in question and in the claim that further evidence will not tell conclusively against it, and when there are costs to postponing the decision or you are risk-avoidant. Seeing the act through—maintaining your commitment to act even if you get evidence against the claim in question—is rationally required when your evidence is strong enough in the above sense, when you are risk-avoidant, and when you would rather allow for the possibility of completing long-term risky projects than act as you think best at every given time.

Taking these two points together, the rationality of faith arises from the possibility that you might get evidence that is misleading—evidence that talks you out of a risk that would have paid off. Faith is a bulwark against misleading evidence. It is the attitude that allows you to start and finish risky long-term projects, particularly when these projects depend on claims about which you cannot be certain before acting.