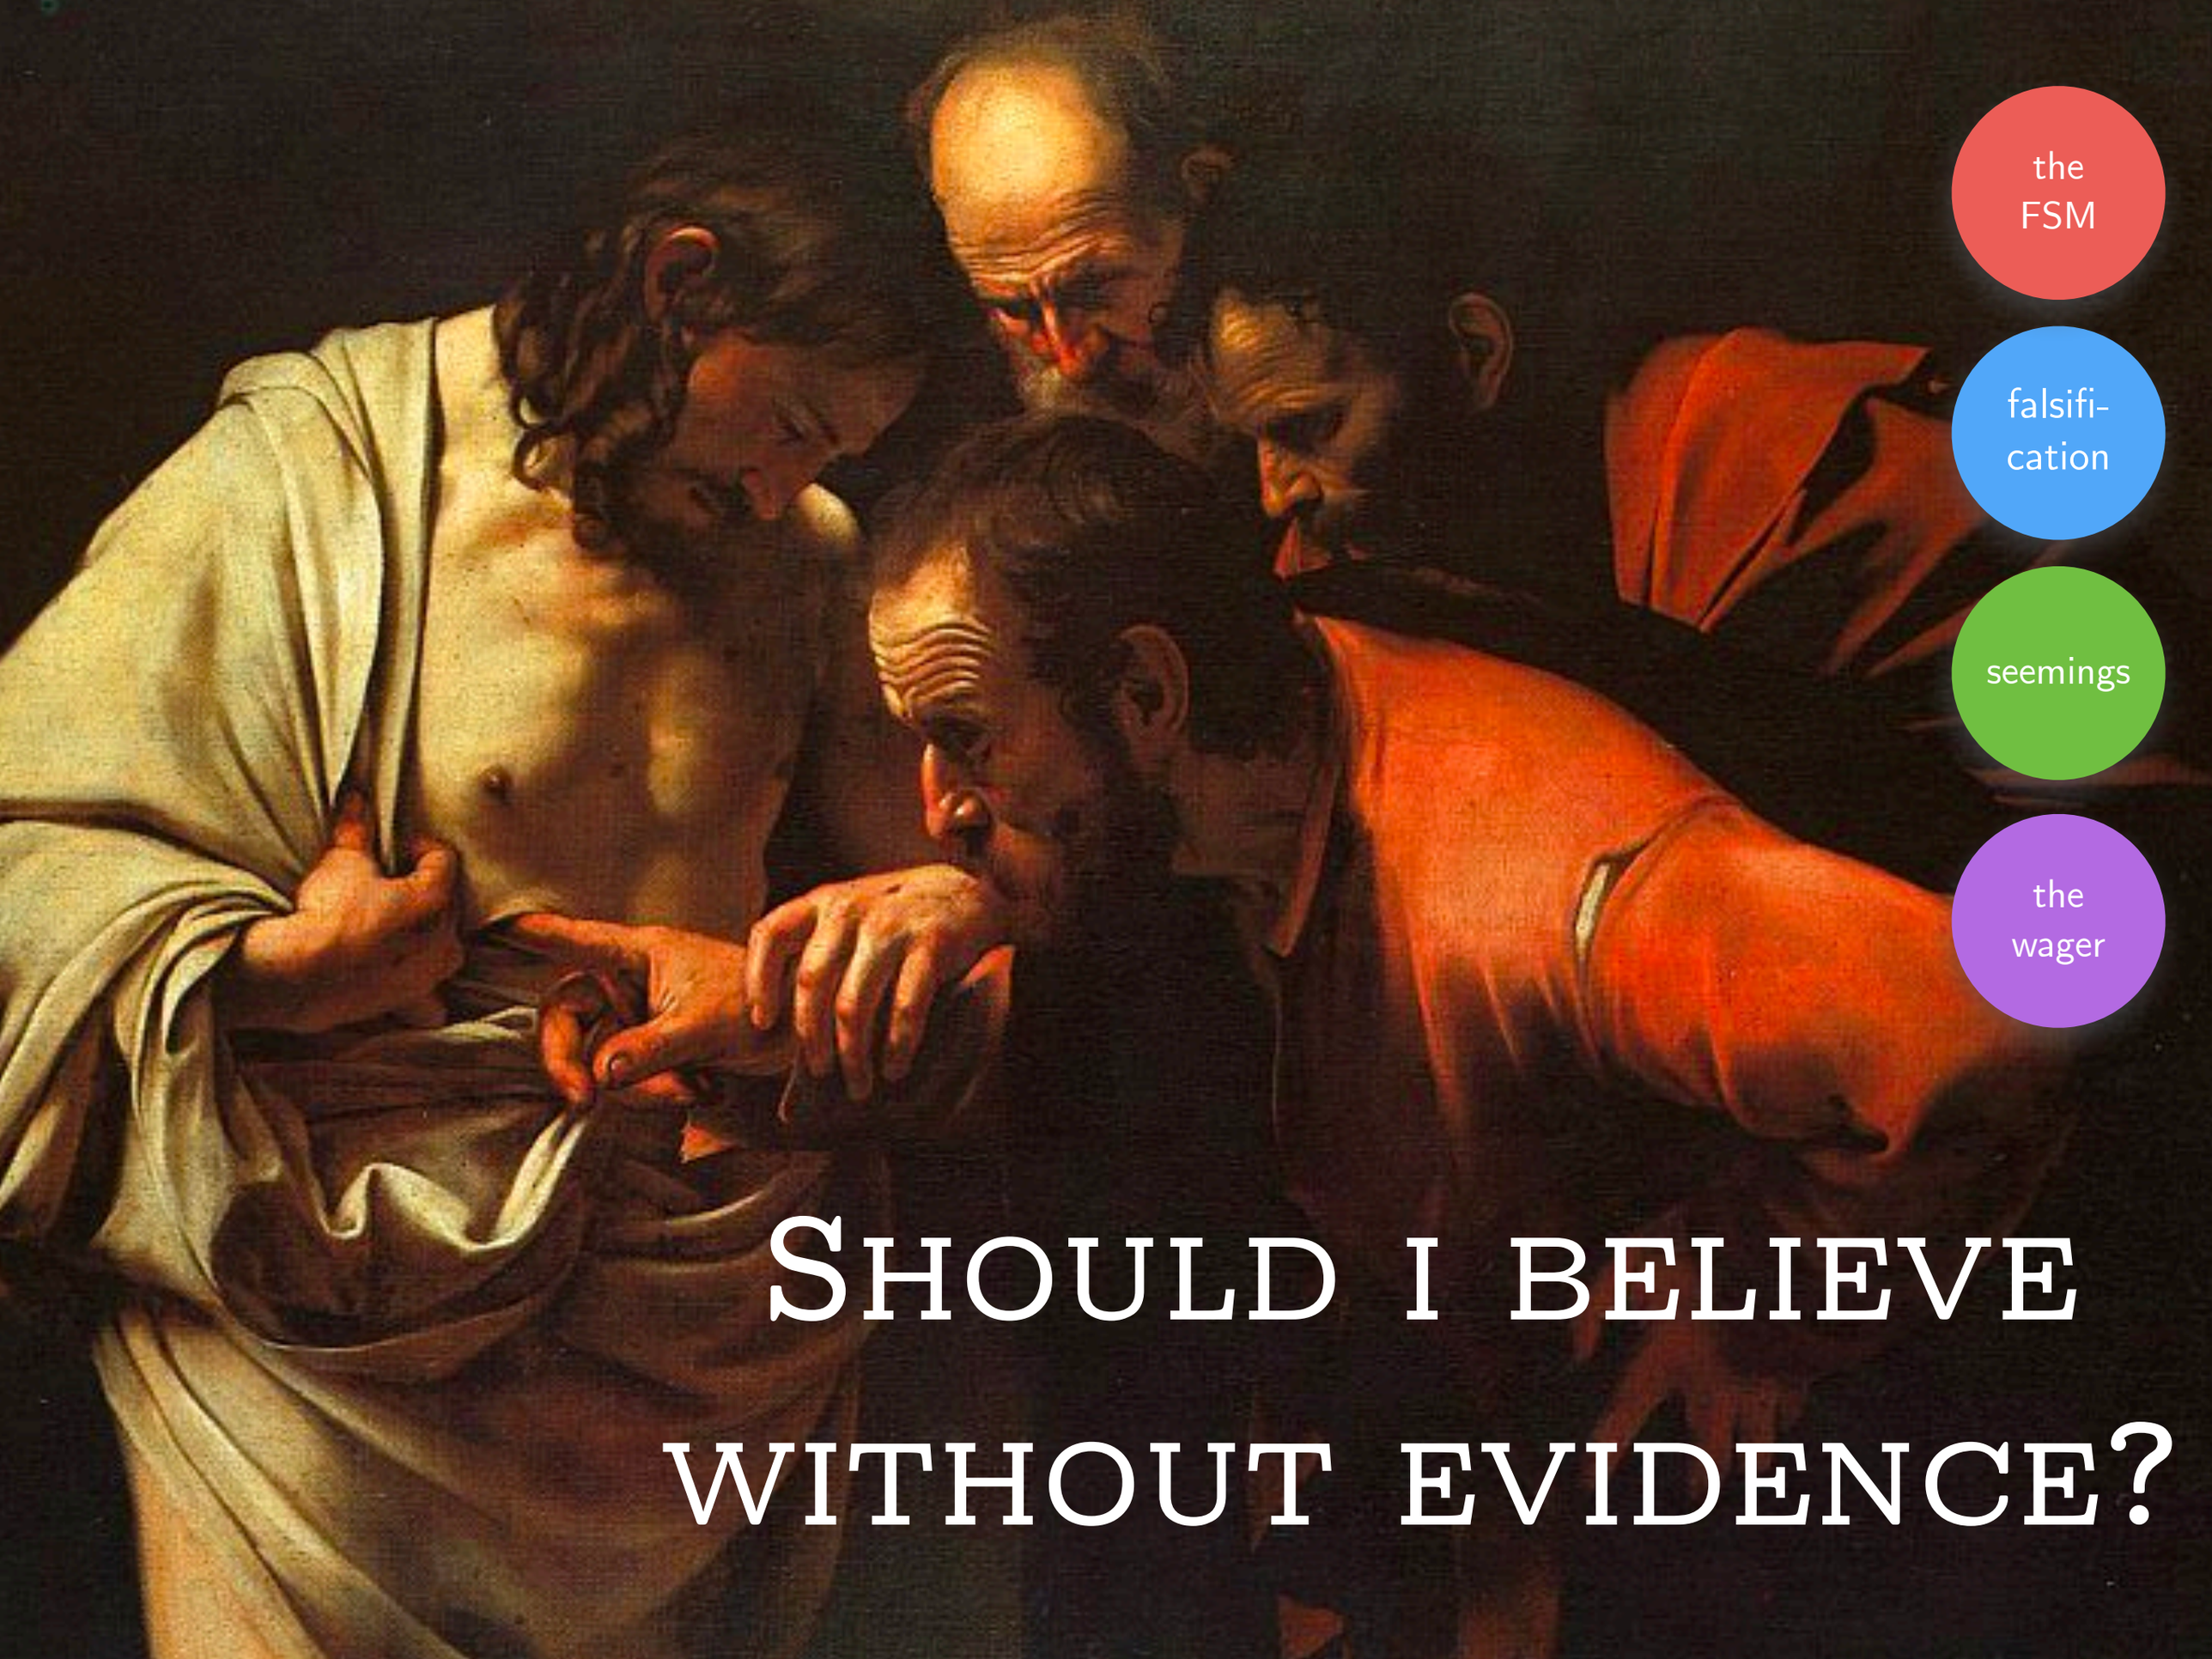
the FSM
falsifi-cation
seemings
the wager
Should i believe without evidence?

Many have thought that the following principle about what we should believe is plausible:

> "It is wrong always, everywhere, and for anyone to believe anything on insufficient evidence."

This slogan, from W.K. Clifford, captures the idea that there is something wrong with forming a belief on **no evidence at all**.

A rule of belief we discussed last time makes this more precise by saying what "evidence" is. Our evidence is what our senses tell us and what we can be certain of.

No Foundations → No Belief
If you can't be certain that P and your senses don't tell you that P and you can't give a good argument for P, you should not believe P.

No Foundations → No Belief
If you can't be certain that P and your senses don't tell you that P and you can't give a good argument for P, you should not believe P.

We have encountered this idea twice already in this course. One place it came up was in the following argument against the existence of immaterial souls:

THE EVIDENTIALIST ARGUMENT AGAINST BELIEF IN IMMATERIAL SOULS

1. We have no sensory experience of immaterial souls.
2. You can't be certain that there are immaterial souls.
3. We have no good argument for the existence of immaterial souls.
4. If you can't be certain that P and your senses don't tell you that P and you don't have a good argument for P, you should not believe P.

C. You should not believe in the existence of immaterial souls. (1,2,3,4)

No Foundations → No Belief
If you can't be certain that P and your senses don't tell you that P and you can't give a good argument for P, you should not believe P.

THE EVIDENTIALIST ARGUMENT AGAINST BELIEF IN IMMATERIAL SOULS

```
1. We have no sensory experience of immaterial souls.
2. You can't be certain that there are immaterial
   souls.
3. We have no good argument for the existence of
   immaterial souls.
4. If you can't be certain that P and your senses don't
   tell you that P and you don't have a good argument
   for P, you should not believe P.
-------------------------
C. You should not believe in the existence of
   immaterial souls. (1,2,3,4)
```

The last premise of this argument just is our proposed No Foundations → No Belief rule.

No Foundations → No Belief
If you can't be certain that P and your senses don't tell you that P and you can't give a good argument for P, you should not believe P.

The last premise of this argument just is our proposed No Foundations → No Belief rule.

It also came up earlier in the course. On the second day, I said that there were two main kinds of arguments against belief in God.

The first are the various versions of the argument from evil which we discussed at length.

The second is the argument that you should not believe that God exists because there is no evidence that God exists.

We can now put that second argument in a more precise form.

No Foundations → No Belief
If you can't be certain that P and your senses don't tell you that P and you can't give a good argument for P, you should not believe P.

We can now put that second argument in a more precise form.

THE EVIDENTIALIST ARGUMENT AGAINST BELIEF IN GOD

```
1. We have no sensory experience of God.
2. You can't be certain that God exists.
3. We have no good argument for the existence of God.
4. If you can't be certain that P and your senses don't
   tell you that P and you don't have a good argument
   for P, you should not believe P.
-------------------------
C. You should not believe in God. (1,2,3,4)
```

THE EVIDENTIALIST ARGUMENT AGAINST BELIEF IN GOD

1. We have no sensory experience of God.
2. You can't be certain that God exists.
3. We have no good argument for the existence of God.
4. If you can't be certain that P and your senses don't tell you that P and you don't have a good argument for P, you should not believe P.

C. You should not believe in God. (1,2,3,4)

One might of course reject premise (3), if you found one of the arguments for the existence of God we discussed in class convincing. And you might reject (1) if you have had certain kinds of mystical experiences.

But many religious believers have not had mystical experiences, and don't take themselves to be in possession of good arguments for God's existence. For them, the belief that God exists is a **basic belief** despite not fitting into either of the two categories of basic belief allowed by foundationalism.

One might of course reject premise (3), if you found one of the arguments for the existence of God we discussed in class convincing. And you might reject (1) if you have had certain kinds of mystical experiences.

But many religious believers have not had mystical experiences, and don't take themselves to be in possession of good arguments for God's existence. For them, the belief that God exists is a **basic belief** despite not fitting into either of the two categories of basic belief allowed by foundationalism.

Last time we discussed a problem for foundationalism: it is hard for the foundationalist to explain why we should believe that the sun will come up tomorrow. This may seem to be good news for the religious believer; if belief in the uniformity of nature can be a basic belief, why not belief in God?

One way to put some pressure on this move uses the example of one of the world's fastest growing religions: Pastafarianism.

One way to put some pressure on this move uses the example of one of the world's fastest growing religions: Pastafarianism.

Q: **How do Pastafarians believe our world was created?**
A: We believe the Flying Spaghetti Monster created the world much as it exists today, but for reasons unknown made it appear that the universe is billions of years old (instead of thousands) and that life evolved into its current state (rather than created in its current form). Every time a researcher carries out an experiment that appears to confirm one of these "scientific theories" supporting an old earth and evolution we can be sure that the FSM is there, modifying the data with his Noodly Appendage. We don't know why He does this but we believe He does, that is our Faith.

<— A ChrisFSMas tree

Pastafarianism has its uses. For example, it can be used to get a religious exemption from the rule that one cannot wear a hat in a driver's license photo:

As you might guess, many Pastafarians take a somewhat less than serious attitude toward the tenets of Pastafarianism. But it can be used to make a serious philosophical point.

As you might guess, many Pastafarians take a somewhat less than serious attitude toward the tenets of Pastafarianism. But it can be used to make a serious philosophical point.

But suppose that someone were a serious Pastafarian. We would, I take it, be inclined to think that there is something irrational about their beliefs.

And this might be so even if we could not come up with any decisive argument against Pastafarianism. After all, Pastafarianism is designed so as to avoid arguments against it. When presented with such an argument, the Pastafarian will simply say that the evidence on which the argument is based is misleading, and was planted by the Flying Spaghetti Monster.

Is there something wrong with the Pastafarian who takes belief in the FSM to be a basic belief?

Let's look at a few different attempts to state rules which will help us decide what basic beliefs to have.

One idea is that what's wrong with Pastafarianism is not that there is some convincing evidence against the FSM, but rather that the belief in the FSM is designed to be immune from counter evidence. In that sense, the claims of Pastafarianism are not **falsifiable**.

By contrast, the claims made by science characteristically are falsifiable. For example, the belief that tomorrow will be like the past is falsifiable — I could have sense experiences tomorrow which show me that it is very unlike the past in various ways. Scientific generalizations — like the claim that water freezes at 0 degrees Celsius — also seem to be falsifiable.

Some have thought that falsifiability is what distinguishes genuine scientific claims from 'pseudo-science.' This might suggest the following rule of belief:

Not Falsifiable → No Belief
If you have no good argument for P and P is not falsifiable, you should not believe P.

This rule allows us to have beliefs which go beyond our experience. It just requires that those beliefs be in some way answerable to a possible experience.

Some have thought that falsifiability is what distinguishes genuine scientific claims from 'pseudo-science.' This might suggest the following rule of belief:

Not Falsifiable → No Belief
If you have no good argument for P and P is not falsifiable, you should not believe P.

This rule allows us to have beliefs which go beyond our experience. It just requires that those beliefs be in some way answerable to a possible experience.

This looks reasonably plausible. It also seems to show what is irrational about Pastafarianism without ruling out the existence of basic beliefs which go beyond sense experience in various ways.

Is belief in God falsifiable?

Not Falsifiable → No Belief
If you have no good argument for P and P is not falsifiable, you should not believe P.

But this rule of belief is not without its problems.

Here's an argument by dilemma that we should not believe this principle:

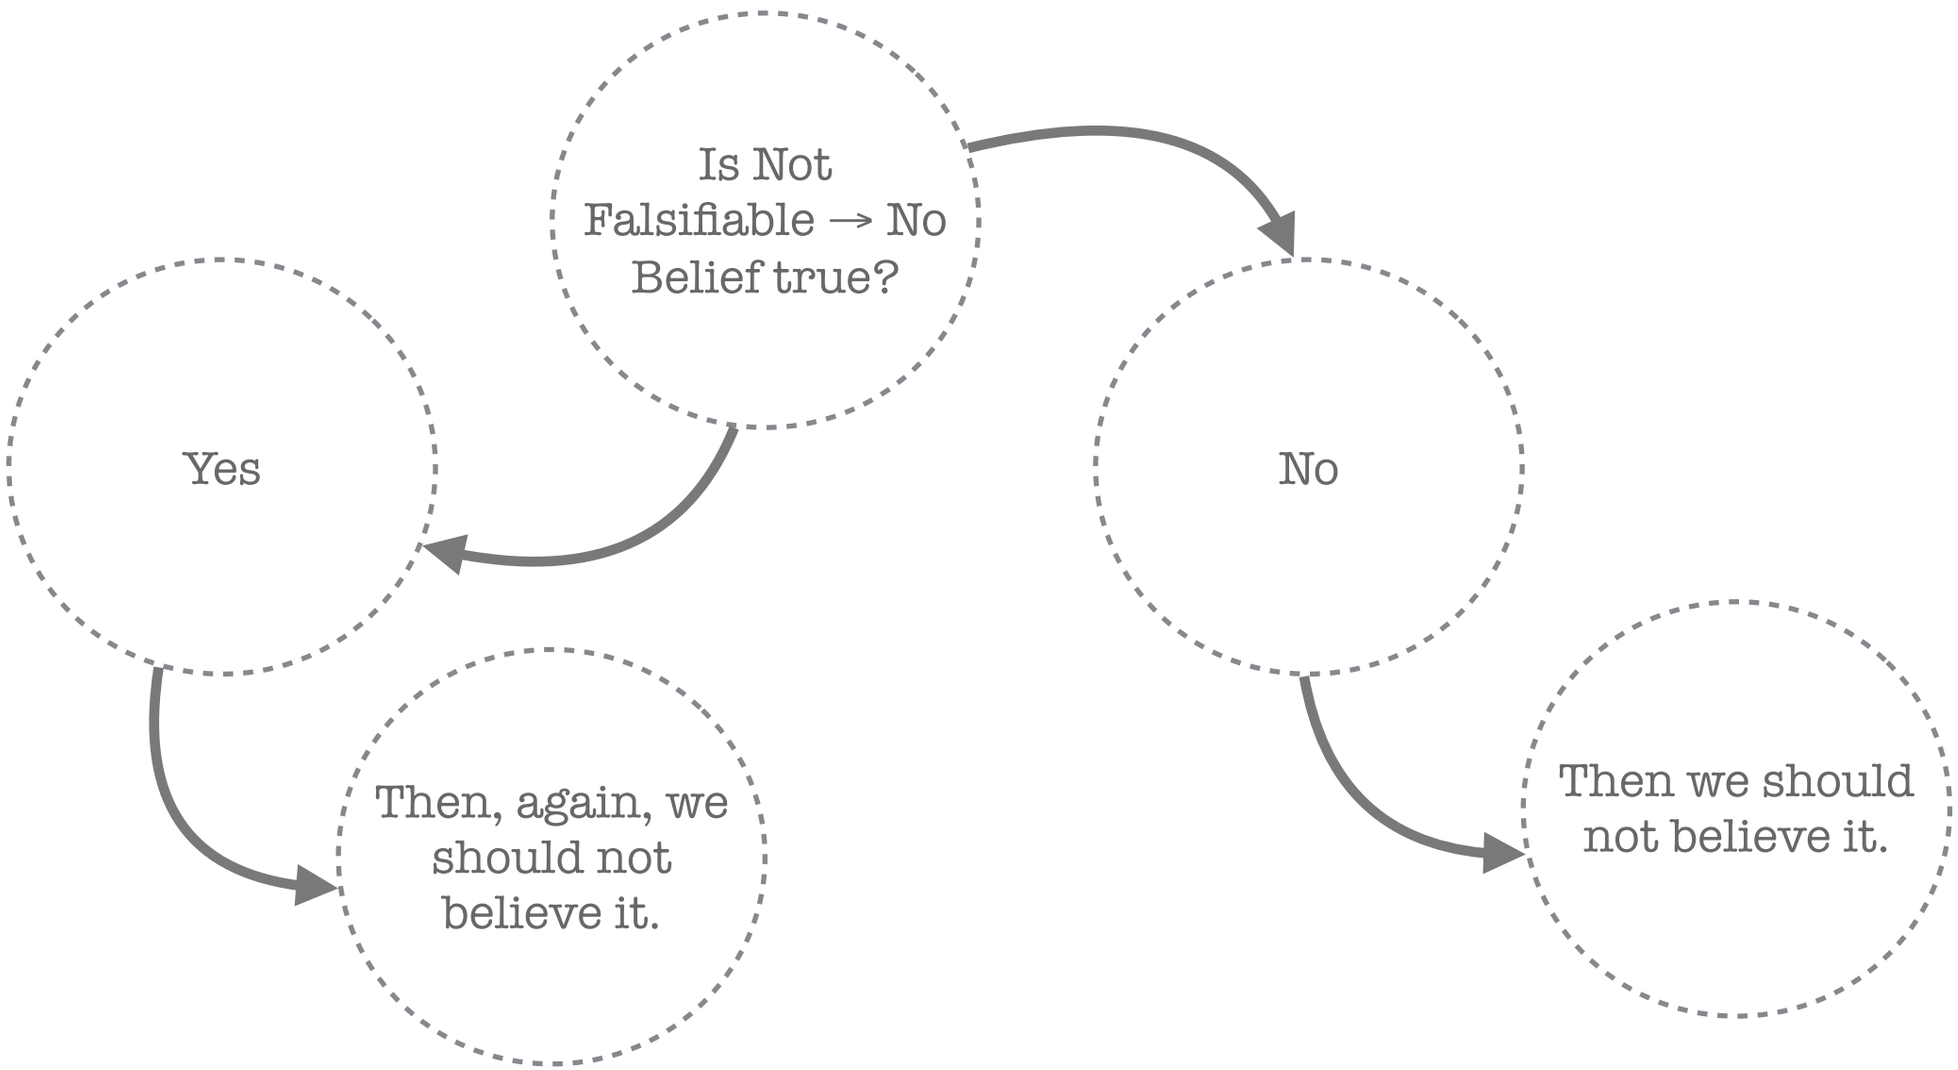

Let's use "Falsificationism" as a name for our rule of belief. Then we can turn this dilemma into an argument as follows.

THE SELF-REFUTATION ARGUMENT AGAINST FALSIFICATIONISM

```
1. If Falsificationism is true, you should believe it only
   if you have an argument that it is true or if it is
   falsifiable.
2. Falsificationism is not falsifiable.
3. You have no good argument that Falsificationism is true.
4. If Falsificationism is true, you should not believe it.
   (1,2,3)
5. If Falsifications is not true, you should not believe it.
6. Falsificationism is true or not true.
-----------------------
C. You should not believe Falsificationism. (4,5,6)
```

This argument seems to show that Falsificationism is self-refuting in a certain sense: it itself implies that you should not believe it.

Note that we could give an exactly parallel argument against Foundationalism.

Not Falsifiable → No Belief
If you have no good argument for P and P is not falsifiable, you should not believe P.

A second type of argument against Falsificationism tries to show that it (like some other negative rules of belief we have considered) is too restrictive.

Remember from our discussion of consciousness the idea of a zombie: someone who is physically just like you, but who lacks consciousness.

Now consider your beliefs that the person next to you is not a zombie. Surely this is a belief that you should have. But is it falsifiable?

Perhaps it is falsifiable for them — they have experiences which tell them that they are not a zombie. But it is unclear, from a falsificationist perspective, why that makes it ok for you to have the belief.

Or consider the belief that all solids have a melting point. Is this falsifiable? What experience could show that some solid has no melting point?

Let's return to the idea of a basic belief you should have. If you think that we should believe that other people are conscious, then it seems like we need to think that there are some basic beliefs you should have of which we cannot be certain and for which we have no sensory evidence. After all, our belief that other people are conscious seems to be (for all we have said) a basic belief you should have.

Here's one idea we might try out. Certain claims just **seem** true to you. It isn't that you can be certain that they are true, or that have sensory evidence that they are true. They just seem true.

Here's one idea we might try out. Certain claims just **seem** true to you. It isn't that you can be certain that they are true, or that have sensory evidence that they are true. They just seem true.

Maybe just the fact that it seems true to us that other people are conscious, and we have no argument against that claim, is good reason for us to believe it. That suggests the following positive rule of belief:

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

One can think of this as a kind of generalization of Experience → Belief. The idea is that experiential seemings are just one kind of seeming.

We'll turn in a second to some challenges to this idea. But let's first ask what would follow if this were correct.

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

We'll turn in a second to some challenges to this idea. But let's first ask what would follow if this were correct.

First, it would seem to permit a rather straightforward response to the challenge of Pastafarianism. After all, the following two claims both certainly seem true.

Spaghetti is a human invention. There is no spaghetti (anywhere in the universe) which was not made by a person.

No person has ever made any magical spaghetti.

But these two claims would seem to rule out the existence of the FSM.

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

Spaghetti is a human invention. There is no spaghetti (anywhere in the universe) which was not made by a person.

No person has ever made any magical spaghetti.

But these two claims would seem to rule out the existence of the FSM.

So perhaps the religious believer should just say that the difference between Pastafarianism and more ordinary religious belief is just that there is a good argument against the former, but not the latter.

Of course, the Pastafarian could then turn to the argument from evil against traditional religious views. But that would be to just give up on the evidentialist argument — we already knew that the religious believer had to have something to say about the different versions of the argument from evil.

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

So perhaps the religious believer should just say that the difference between Pastafarianism and more ordinary religious belief is just that there is a good argument against the former, but not the latter.

We should not overstate the case here. Some Pastafarians might reply that the preceding claims about spaghetti do **not** seem true to them, and that it **does** seem true to them that the FSM exists. If this is really true, then Seems → Belief would seem to imply that this Pastafarian **should** believe in the FSM.

The point is just that the believer in God who accepts Seems → Belief would seem to have a way to defend the view that, by their own lights, belief in God is reasonable in a way that belief in the FSM is not.

Seems → Belief
If it seems to you that P is true,
and you have no argument
against P, you should believe P.

A second reason why this rule of belief might seem attractive to the religious believer is that it promises to make sense of the reason many people would give for their belief in God.

Many religious believers cite as a reason for their belief things like the beauty of nature.

This can seem puzzling; it isn't as if there is a straightforward argument from the beauty of nature to the existence of God. A more plausible understanding of what they mean is that when they contemplate nature it just seems to them that God must exist.

If the above rule of belief is right, then — again presuming a satisfactory reply to the arguments from evil — this can be a reasonable basis for religious belief.

Seems → Belief
If it seems to you that P is true,
and you have no argument
against P, you should believe P.

I want to look now at three challenges to this kind of rule of belief.

The first goes back to Descartes. His idea was that if we don't carefully examine our whole structure of belief, we can allow error to slip in. Surely just going by how things seem is not going to be a foolproof way to escape error! So why is this not just an irresponsible way to go about forming beliefs?

One response to this is suggested by the following quote from William James.

One response to this is suggested by the following quote from William James.

> "There are two ways of looking at our duty in the matter of opinion ... We must know the truth, and we must avoid error. These are our first and great commandments as would-be knowers; but they are not two ways of stating an identical commandment, they are two separable laws."

In this spirit, one might say that Descartes' advice is the best one if we only care about minimizing error. But this is not our only aim: we also want to believe the truth. If we limit ourselves to the beliefs we can be certain of, we will in so doing prevent ourselves from believing many truths.

Seems → Belief
If it seems to you that P is true,
and you have no argument
against P, you should believe P.

A second worry about Seems → Belief is that it is too permissive — it implies that we should have beliefs which we shouldn't.

Remember the example of the belief that my dog is a poodle. Suppose that, after you learn about the other breeds it could be, it still just seems to you like it is a poodle. Could that really be enough for this to be a belief you should have?

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

The last worry is the problem that not all "seemings" are created equal. For example, to many people individuals of another race "seem" more threatening or untrustworthy. To many men, women "seem" less intelligent or capable.

On its face, Seems → Belief would seem to license these people to form all kinds of beliefs about people of another race or gender that they surely should not form. What should we say about cases like this?

Seems → Belief
If it seems to you that P is true, and you have no argument against P, you should believe P.

Let's look first at how things seem to us in our visual experiences. Some interesting studies have been done which seem to show that our background beliefs, expectations, and desires can have an effect on how things visually appear to us.

In one well-known study, white Americans were first shown a picture of either a white man's face or a Black man's face, and then shown a picture of either a tool or a gun. Under time pressure, they had to categorize what they were shown. Participants primed with a Black man's face mischaracterized tools as guns significantly more than those primed with a white man's face.

The best way to interpret this study is controversial. But what seems reasonably clear is that whether the participant saw a white face or a Black face affected whether it seemed to them that they were being shown a gun or a tool.

In one well-known study, white Americans were first shown a picture of either a white man's face or a Black man's face, and then shown a picture of either a tool or a gun. Under time pressure, they had to categorize what they were shown. Participants primed with a Black man's face mischaracterized tools as guns significantly more than those primed with a white man's face.

The best way to interpret this study is controversial. But what seems reasonably clear is that whether the participant saw a white face or a Black face affected whether it seemed to them that they were being shown a gun or a tool.

Similar results have been obtained in less politically charged contexts. In one case, people are given two beers, one of which has some balsamic vinegar in it, and asked to pick which one they liked better. A majority chose the one with balsamic vinegar in it. The experiment was then repeated with the change that participants were told in advance that one of the beers had some vinegar in it (but not which one). A majority chose the one with ought vinegar in it. Some infer that the expectation of a vinegar taste changed the way the liquid tasted to the subjects.

One reason why these cases are interesting is that they call into question Seems → Belief. If our background beliefs can affect the way things seem to us, then it is tempting to say that we should trust the seemings only if we should have the belief on which the seemings are based.

One reason why these cases are interesting is that they call into question Seems → Belief. If our background beliefs can affect the way things seem to us, then it is tempting to say that we should trust the seemings only if we should have the belief on which the seemings are based.

Here's an interesting case, from the contemporary philosopher Susanna Siegel:

> Jill, for no particular reason, has the belief that Jack is angry. This is a belief which Jill should not have.
>
> When Jill sees Jack, Jill's belief that Jack is angry at her makes Jack look angry to her — it causes it to seem to her that Jack looks angry.
>
> On the basis of the fact that it visually seems to her the Jack is angry, Jill's belief that Jack is angry at her is strengthened.

At the start, Jill should not believe that Jack is angry. If Seems → Belief is true, it looks like at the end she **should** believe that Jack is angry. But can this be right? Does Jill really have a better reason for her belief at the end than at the start?

At the start, Jill should not believe that Jack is angry. If Seems → Belief is true, it looks like at the end she **should** believe that Jack is angry. But can this be right? Does Jill really have a better reason for her belief at the end than at the start?

We have been focusing on perceptual cases. But the moral of the above cases would seem to apply even more strongly to cases of non-perceptual seemings.

Consider the way in which your political beliefs can affect what claims seem true to you. This is an instance of the well-known phenomenon of **confirmation bias**.

Cases of confirmation bias are structurally the same as the Jack/Jill case: one begins with a belief (which might well be a belief one should not have), that belief causes other claims to seem true, and those other claims support the original belief.

If Seems → Belief is true, this kind of thing is perfectly ok. One's belief in P can be justified by one's belief in Q, even if one believes Q because Q seems true and Q seems true because one believes P.

But doesn't this seem like the kind of circular reasoning we would reject in other contexts?

But doesn't this seem like the kind of circular reasoning we would reject in other contexts?

It is worth thinking about how we might modify Seems → Belief in response to cases of this kind. Here's one suggestion:

Restricted Seems → Belief
If it seems to you that P is true, and you have no argument against P, and the seeming is not caused by a belief you should not have, you should believe P.

This rule restricts the seemings you should trust to the ones that are not caused by beliefs you should not have. This would block the result that Jill should believe that Jack is angry.

The problem, though, is that it is hard to know how one could employ this rule. After all, the problem with the cases under discussion is that one can't tell from the inside when a seeming is caused by one of one's beliefs.

This rule restricts the seemings you should trust to the ones that are not caused by beliefs you should not have. This would block the result that Jill should believe that Jack is angry.

The problem, though, is that it is hard to know how one could employ this rule. After all, the problem with the cases under discussion is that one can't tell from the inside when a seeming is caused by one of one's beliefs.

Here's an idea. Perhaps you should not trust seemings when **you have good reason to think that** the seeming is based on an unjustified belief, and hence good reason to think that the seeming is unreliable:

Restricted Seems → Belief 2.0
If it seems to you that P is true, and you have no argument against P, and you have no good reason to think that the seeming is unreliable, you should believe P.

What does this say about the case of Jack and Jill? It says that, if Jill is not aware that her belief played a role in Jack seeming angry, she should form the belief that he is angry. (After all, she had no way of knowing that the seeming was unreliable.)

Restricted Seems → Belief 2.0

If it seems to you that P is true, and you have no argument against P, and you have no good reason to think that the seeming is unreliable, you should believe P.

What does this say about the case of Jack and Jill? It says that, if Jill is not aware that her belief played a role in Jack seeming angry, she should form the belief that he is angry. (After all, she had no way of knowing that the seeming was unreliable.)

But suppose that Jill is told that beliefs about people can play a big role in determining how you perceive their emotions. This information would give her reason to think that Jack's seeming angry to her is unreliable — and in that case she should not reinforce her belief that Jack is angry.

In general, it seems like the best course of action for those who accept Seems → Belief is to restrict it in some way, and to educate themselves about the various situations in which background beliefs (or other mental states) are most likely to affect how things seem to them.

Let's look at one last idea about basic beliefs. This is the idea that we should sometimes have a belief for **practical** reasons. Our focus will be a famous argument from Blaise Pascal which makes use of this idea.

Pascal was a 17th century French philosopher, theologian, and mathematician; he made foundational contributions to, among other areas, the early development of the theory of probability.

Pascal was one of the first thinkers to systematically investigate the question of how we should make decisions under situations of uncertainty, where we don't know all of the relevant facts about the world, or the outcomes of our actions. He was (with his contemporary Fermat) the first to formulate the idea of **expected utility**.

Pascal was one of the first thinkers to systematically investigate the question of how we should make decisions under situations of uncertainty, where we don't know all of the relevant facts about the world, or the outcomes of our actions. He was (with his contemporary Fermat) the first to formulate the idea of **expected utility**.

The expected utility of an action can be calculated by looking at the various possible outcomes of the action, and assigning each a value — measuring how good the outcome is — and a probability — measuring how likely the outcome is. Because you know that one of the outcomes is going to happen (but not more than one) the probabilities should sum to 1. To get the expected utility, you multiply each outcome's value by its probability, and add them all up.

So consider a bet in which a fair coin is flipped. Suppose that you get $5 if it comes up heads, and lose $3 if it comes up tails. Then the expected utility is:

$$(0.5 * \$5) + (0.5 * \$3) = \$1$$

Because this is a positive expected utility, you should take the bet if offered (and if you have nothing better to do).

Many have thought that expected utility considerations should guide our actions. Perhaps we should all act so as to maximize the expected utility of our actions. (We'll come back to this idea in a few weeks.)

Pascal had the thought that forming beliefs is just one sort of action we perform. So, if in general it makes sense for our actions to be guided by expected utility, why not also form beliefs on the basis of expected utility?

This suggests the following rule of belief:

Expected Utility → Belief
If believing P has a higher expected utility than not believing P, you should believe P.

This rule of belief led Pascal to a famous argument for belief in God.

This rule of belief led Pascal to a famous argument for belief in God.

> "It would be unwise of you, since you are obliged to play, not to risk your life to win three lives at a game in which there is an equal chance of winning and losing. But here there is an infinity of happy life to be won ... and what you are staking is finite. ... And thus, since you are obliged to play, you must be renouncing reason if you hoard your life rather than risk it for an infinite gain, which is just as likely to occur as a loss..."

Our question is: how might Pascal argue that believing in God has higher expected utility than nonbelief?

First, he emphasizes that "there is an equal chance of gain and loss" — an equal chance that God exists, and that God does not exist. This means that we should assign each a probability of 1/2.

Second, he says that in this case the amount to be won is infinite. We can represent this by saying that the utility of belief in God if God exists is ∞.

Let's suppose, plausibly, that if we believe in God, and God does not exist, this involves some loss of utility. This loss will be finite — let's symbolize it by the word "loss".

One might represent these assumptions as follows:

Courses of action	God exists	God does not exist	Expected utility
believe	∞	loss	∞
don't believe	0	0	0
	Probability = 0.5	Probability = 0.5	

Courses of action	God exists	God does not exist	Expected utility
believe	∞	loss	∞
don't believe	0	0	0
	Probability = 0.5	Probability = 0.5	

So it looks as though the expected utility of believing in God is infinite, whereas the expected utility of nonbelief is 0. If the rule of expected utility is correct, it follows that it is rational to believe in God - and it is not a very close call.

Let's look at a few objections to the idea that the above chart accurately represents our choice of whether or not to believe in God.

Courses of action	God exists	God does not exist	Expected utility
believe	∞	loss	∞
don't believe	0	0	0
	Probability = 0.5	Probability = 0.5	

Objection 1: the probability that God exists is not 1/2, but some much smaller number -- say, 1/100.

Courses of action	God exists	God does not exist	Expected utility
believe	∞	loss	∞
don't believe	0	0	0
	Probability = 0.01	Probability = 0.99	

Objection 1: the probability that God exists is not 1/2, but some much smaller number -- say, 1/100.

This is a real strength of Pascal's argument: it does not depend on any assumptions about the probability that God exists other than the assumption that it is nonzero. In other words, he is only assuming that we don't know for sure that God does not exist, which seems to many people - including many atheists - to be a reasonable assumption.

Courses of action	God exists	God does not exist	Expected utility
believe	∞	loss	∞
don't believe	0	0	0
	Probability = m	Probability = n	

Objection 2: Pascal is assuming that, if God exists, there is a 100% chance that believers will get infinite reward.

To accommodate this possibility, we would have to add another column to our chart, to represent the two possibilities imagined. Let's call these possibilities "Rewarding God" and "No reward God", and let's suppose that each has a nonzero probability of being true.

Courses of action	Rewarding God exists	No reward God exists	God does not exist	Expected utility
believe	∞	0	loss	∞
don't believe	0	0	0	0
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Objection 2: Pascal is assuming that, if God exists, there is a 100% chance that believers will get infinite reward.

As this chart makes clear, adding this complication has no effect on the result. Pascal needn't assume that God will certainly reward all believers; he need only assume that there is a nonzero chance that God will reward all believers.

Courses of action	Rewarding God exists	No reward God exists	God does not exist	Expected utility
believe	∞	0	loss	∞
don't believe	0	0	0	0
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Objection 3: God might give eternal reward to believers and nonbelievers alike.

Let's call the hypothesis that God will give eternal reward to all "Generous God."

Courses of action	Rewarding God exists	Generous God exists	God does not exist	Expected utility
believe	∞	∞	loss	∞
don't believe	0	∞	0	∞
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Objection 3: God might give eternal reward to believers and nonbelievers alike.

Setting aside the possibility of No reward God, which we have seen to be irrelevant, taking account of the possibility of Generous God has a striking effect on the expected utilities of belief and nonbelief.

Objection 3: God might give eternal reward to believers and nonbelievers alike.

Setting aside the possibility of No reward God, which we have seen to be irrelevant, taking account of the possibility of Generous God has a striking effect on the expected utilities of belief and nonbelief.

Now, it appears, belief and nonbelief have the same infinite expected utility, which undercuts Pascal's argument for the rationality of belief in God.

However, Pascal seems to have a reasonable reply to this objection. It seems that the objection turns on the fact that any probability times an infinite utility will yield an infinite expected value. And that means that any two actions which have some chance of bring about an infinite reward will have the same expected utility.

But this is extremely counterintuitive. Suppose we think of a pair of lotteries, EASY and HARD. Each lottery has an infinite payoff, but EASY has a 1/3 chance of winning, whereas HARD has a 1/1,000,000 chance of winning. What is the expected utility of EASY vs. HARD? Which would you be more rational to buy a ticket for?

But this is extremely counterintuitive. Suppose we think of a pair of lotteries, EASY and HARD. Each lottery has an infinite payoff, but EASY has a 1/3 chance of winning, whereas HARD has a 1/1,000,000 chance of winning. What is the expected utility of EASY vs. HARD? Which would you be more rational to buy a ticket for?

How might we modify our rule of expected utility to explain this case? Would this help Pascal respond to the case of Generous God?

A natural suggestion is to say something like this: if two actions each have infinite expected utility, then (supposing that neither action has a very high chance of leading to a very bad outcome) it is rational to go with the action that has the higher probability of leading to the infinite reward. This sort of supplement to the rule of expected utility explains why it is smarter to buy a ticket in EASY than in HARD; and it also helps Pascal solve the problem of Generous God, since the believer receives an infinite reward if either Generous God or Rewarding God exists, whereas the nonbeliever only gets a reward in the first of these cases.

Courses of action	Rewarding God exists	Generous God exists	God does not exist	Expected utility
believe	∞	∞	loss	∞
don't believe	0	∞	0	∞
	Pr. = m	Pr. = n	Pr. = 1-m-n	

If we adopt this modified rule — which says that in cases where two outcomes each have an infinite expected utility, one should choose the action more likely to lead to one of these outcomes —then this argues for belief in the case of Generous God, so long as m≠0.

Courses of action	Rewarding God exists	Generous God exists	God does not exist	Expected utility
believe	∞	∞	loss	∞
don't believe	0	∞	0	∞
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Objection 4: God might give eternal reward to just those who do not believe.

It is conceivable that God would do the opposite of rewarding belief, and instead would reward only disbelief. Call this hypothesis 'Anti-Wager God.'

Courses of action	Rewarding God exists	Anti-Wager God exists	God does not exist	Expected utility
believe	∞	0	loss	∞
don't believe	0	∞	0	∞
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Objection 4: God might give eternal reward to just those who do not believe.

It is conceivable that God would do the opposite of rewarding belief, and instead would reward only disbelief. Call this hypothesis 'Anti-Wager God.'

Courses of action	Rewarding God exists	Anti-Wager God exists	God does not exist	Expected utility
believe	∞	0	loss	∞
don't believe	0	∞	0	∞
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Objection 4: God might give eternal reward to just those who do not believe.

It is no longer obvious that belief has a higher chance of reward than nonbelief: we need an argument that Rewarding God is more likely to exist than Anti-Wager God. This shows that Pascal's argument can't be completely free of commitments to the probabilities of certain theological claims.

Courses of action	Rewarding God exists	Anti-Wager God exists	God does not exist	Expected utility
believe	∞	0	loss	∞
don't believe	0	∞	0	∞
	Pr. = m	Pr. = n	Pr. = 1-m-n	

Note also that this scenario is analogous to the hypothesis that God rewards only the adherents of certain specific religions, only one of which can be believed.

So far we have focused on objections which try to show that expected utility calculations do not deliver the result that it is rational to believe that God exists.

I want now to consider three quite different lines of reply to Pascal's argument, which do not involve trying to find a flaw in his calculations.

In cases with infinite utility, the rule of expected utility fails.

It is impossible to form beliefs on the basis of expected utility calculations.

It is irrational to form beliefs on the basis of expected utility calculations.

In cases with infinite utility, the rule of expected utility fails.

Consider the following bet:

The St. Petersburg

I am going to flip a fair coin until it comes up heads. If the first time it comes up heads is on the 1st toss, I will give you $2. If the first time it comes up heads is on the second toss, I will give you $4. If the first time it comes up heads is on the 3rd toss, I will give you $8. And in general, if the first time the coin comes up heads is on the nth toss, I will give you $$2^n$$.

The St. Petersburg

I am going to flip a fair coin until it comes up heads. If the first time it comes up heads is on the 1st toss, I will give you \$2. If the first time it comes up heads is on the second toss, I will give you \$4. If the first time it comes up heads is on the 3rd toss, I will give you \$8. And in general, if the first time the coin comes up heads is on the nth toss, I will give you $\$2^n$.

Would you pay \$2 to take this bet? How about \$4?

Suppose now I raise the price to \$10,000. Should you be willing to pay that amount to play the game once?

What is the expected utility of playing the game?

The St. Petersburg

I am going to flip a fair coin until it comes up heads. If the first time it comes up heads is on the 1st toss, I will give you \$2. If the first time it comes up heads is on the second toss, I will give you \$4. If the first time it comes up heads is on the 3rd toss, I will give you \$8. And in general, if the first time the coin comes up heads is on the nth toss, I will give you $\$2^n$.

What is the expected utility of playing the game?

We can think about this using the following table:

Outcome	First heads is on toss #1	First heads is on toss #2	First heads is on toss #3	First heads is on toss #4	First heads is on toss #5	.....
Probability	\$2	\$4	\$8	\$16	\$32	.....
Payoff	1/2	1/4	1/8	1/16	1/32	.....

Outcome	First heads is on toss #1	First heads is on toss #2	First heads is on toss #3	First heads is on toss #4	First heads is on toss #5	.....
Probability	$2	$4	$8	$16	$32	.....
Payoff	1/2	1/4	1/8	1/16	1/32	.....

The expected utility of playing = the sum of probability x payoff for each of the infinitely many possible outcomes. So, the expected utility of playing equals the sum of the infinite series

1+1+1+1+1+ 1+1+1+1+1+ 1+1+1+1+1+1+1+1+1+1+.....

But it follows from this result, plus the rule of expected utility, that you would be rational to pay any finite amount of money to have the chance to play this game once. But this seems clearly mistaken. What is going on here?

Does this show that the rule of expected utility can lead us astray? If so, in what sorts of cases does this happen? Does this result depend essentially on their being infinitely many possible outcomes?

But it follows from this result, plus the rule of expected utility, that you would be rational to pay any finite amount of money to have the chance to play this game once. But this seems clearly mistaken. What is going on here?

Does this show that the rule of expected utility can lead us astray? If so, in what sorts of cases does this happen? Does this result depend essentially on their being infinitely many possible outcomes?

Suppose that we set an upper bound of 100 coin flips on the game, so that if you get to the 100th flip you get $\$2^{100}$ (a very large number) no matter how the coin comes up. Then the expected utility of playing will be \$100. Would you pay \$99 to play this game?

Most would say not. One possibility is that this is explained by a combination of risk aversion and decreasing marginal utility. Could these also play a role in the evaluation of Pascal's wager?

It is impossible to form beliefs on the basis of expected utility calculations.

Suppose that I offer you $5 to raise your arm. Could you do it?

But now suppose I offered you $5 to believe that you are not now sitting down. Can you do that (without standing up)?

Cases like this suggest that it is impossible to form beliefs on the basis of expected utility calculations.

It is impossible to form beliefs on the basis of expected utility calculations.

What does he have in mind here?

Pascal considered this objection, and gave the following response:

> “I am so made that I cannot believe. What do you want me to do then?”
>
> “At least get it into your head that, if you are unable to believe, it is because of your passions, since reason tells you to believe and yet you cannot do so. Concentrate then not on convincing yourself by multiplying proofs of God’s existence, but by diminishing your passions.”

It is irrational to form beliefs on the basis of expected utility calculations.

Let's now turn to our last line of objection to Pascal.

Pascal's argument, as we have reconstructed it, relies on the following principle.

Expected Utility → Belief
If believing P has a higher expected utility than not believing P, you should believe P.

This principle seems plausible. But so does this one:

Low Probability → No Belief
If you think that P has a very low probability of being true, you should not believe P.

Expected Utility → Belief
If believing P has a higher expected utility than not believing P, you should believe P.

Low Probability → No Belief
If you think that P has a very low probability of being true, you should not believe P.

Pascal's reasoning shows that these rules can come into conflict, because sometimes believing something which you think has a very low probability of being true can have a higher expected utility than not believing it.

One important question for those who find Pascal's argument convincing is: how could this second principle be false?