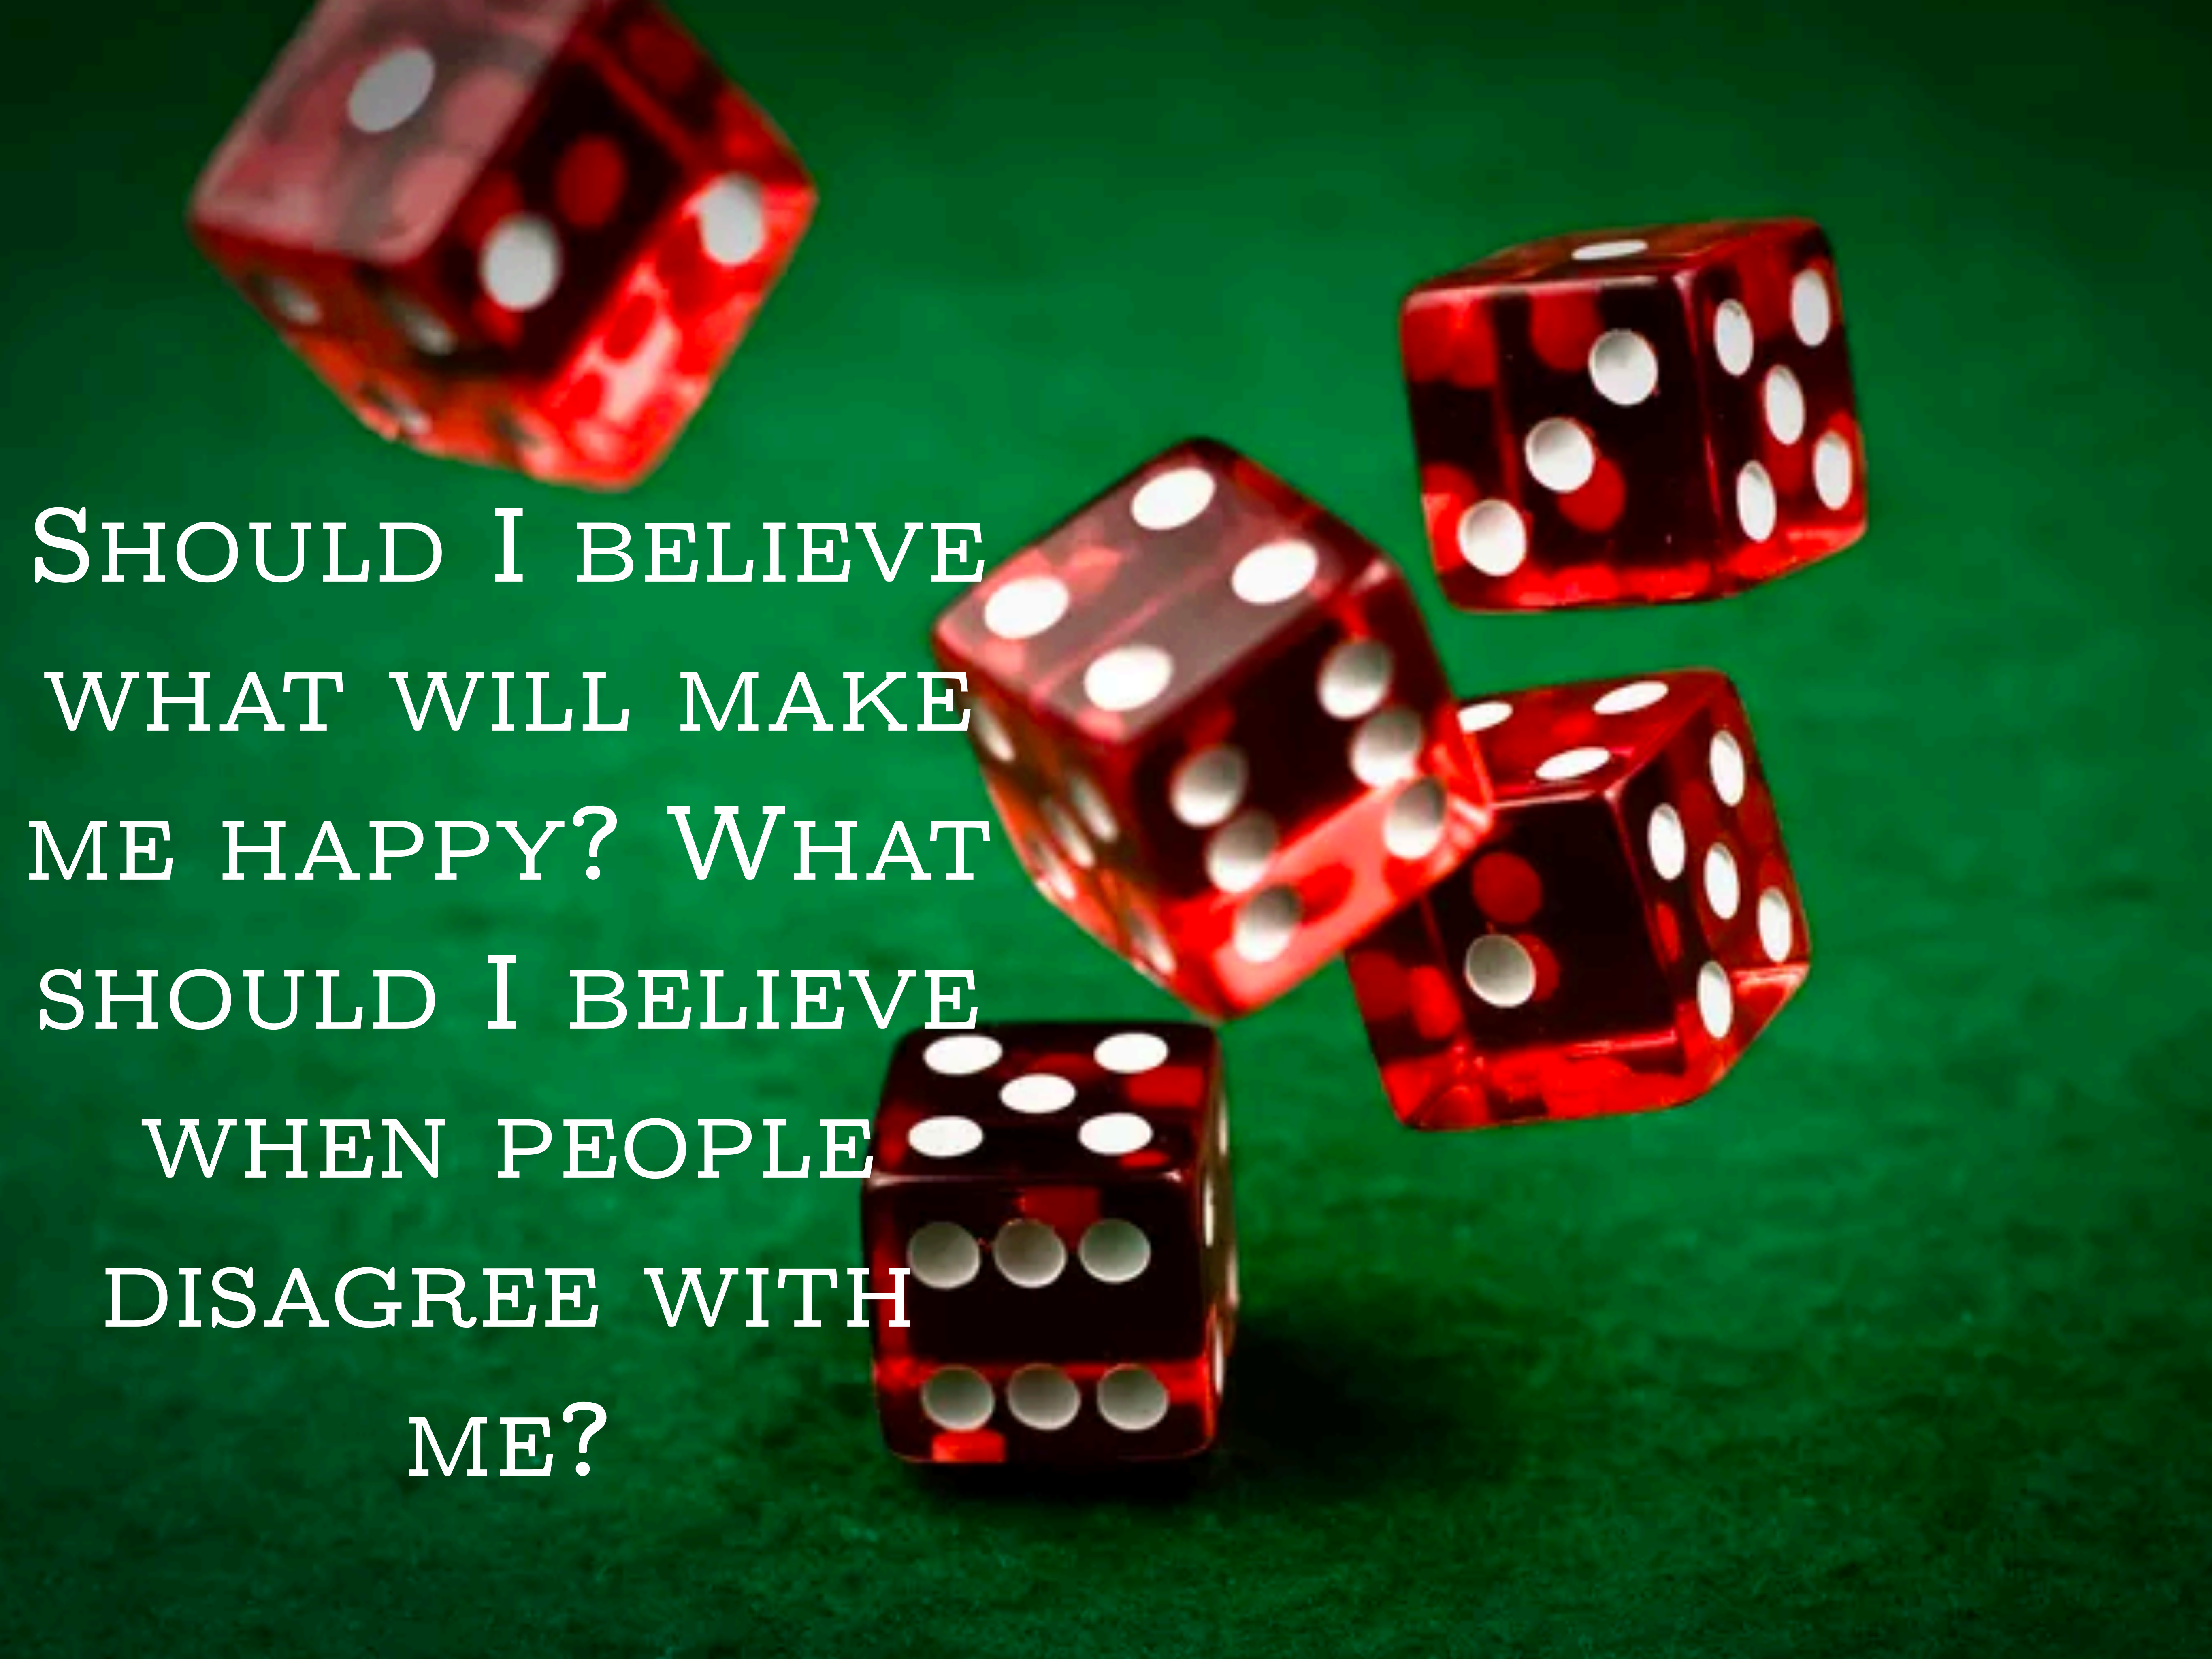
Should I believe what will make me happy? What should I believe when people disagree with me?

Last class we introduced the conservative idea that you should stick with your existing beliefs unless you encounter a good argument against them.

Our first topic today is a challenge to this idea. It is an attempt to show that certain situations can provide us good reason to abandon our beliefs despite their not providing any argument (at least directly) for their falsity. These are situations of **disagreement**.

Let's start with an example which seems to show that we should modify our views in response to certain kinds of disagreement.

Let's start with an example which seems to show that we should modify our views in response to disagreement of this kind.

The horse race

Imagine that you are at a horse track with a friend. Two horses, A and B, are competing for the lead down the stretch. At the finish, it is extremely close, but it looks to you that horse A won. You are highly confident that you are correct.

Your friend then turns to you and says "I can't believe that B won."

Should you now be less confident in your initial judgement?

Splitting the bill

You are in a restaurant with some friends, and the bill comes. You've agreed to split the bill equally. You think that everyone owes $19.

Your friend says, "OK, everybody should chip in $18."

Should you now be less confident that everyone owes $19?

These are simple cases of disagreement. Many people have the intuition that, in cases like these, disagreement should lead us to revise our beliefs.

Of course, we are making some background assumptions about these cases. For example, in the case of the horse race we are assuming that the other person has as good of eyesight as you do. And in the splitting the bill case we are assuming that your friends are not terrible at math.

Whenever someone has as many of the qualities which are relevant to correctly answering some question as you do, we will say that that person is your **epistemic peer**.

Intuitively, the idea is that your epistemic peer is just as likely to give the correct answer to some question as you are. So, if an epistemic peer disagrees with you, you should reduce your confidence in your own belief.

Exactly how much should you reduce your confidence?

Exactly how much should you reduce your confidence?

The natural answer to this question starts with the fact that, in ordinary life, we don't just believe or disbelieve things; we also take them to have a certain probability of being true. The probability that you take P to have is called your credence in P. Credence can be expressed as a percentage, or as a number between 0 and 1 (1 means that you are sure that P is true, 0 that you are sure that P is false).

Probability-splitting
If you assign P credence N, and come across an epistemic peer who assigns P credence M, then you should change your credence in P to the average of N and M.

Suppose that both you and your friend have credence of 0.9 in your initial views about the winner of the horse race. This rule says that, on learning of your disagreement, you should both adjust your credence to 0.5.

Probability-splitting
If you assign P credence N, and come across an epistemic peer who assigns P credence M, then you should change your credence in P to the average of N and M.

This rule also seems to say the right thing about cases in which you meet a large number of epistemic peers who disagree with you.

The poll

I put an argument on the screen, and conduct a poll, asking you to say whether the argument is valid or invalid. You confidently answer "Valid." When the poll results show up, you find to your surprise that you are the only student who answered this way.

We can think of this as a case in which you have many simultaneous disagreements. Supposing for simplicity that everyone initially has credence 1 in her answer and that everyone in the class is your epistemic peer, the Probability Splitting Rule would suggest that you should lower your credence in your initial answer to 0.5, then to 0.25, then to 0.125, then to .... a small number.

Here's a problem case for the probability splitting rule:

An argument for astrology?

Astrology is the view that we can predict the events in ordinary people's lives by the time of their birth and the relative locations of the stars and planets. I have the view that astrology is completely unscientific; there's just no evidence to show that it works. But 45% of Americans (62% between the ages of 18 and 24!) think that astrology is either "scientific" or "sort of scientific." So, following the advice of The Equal Weight View, I significantly increase my credence in the scientific status of astrology.

Other, similar examples are easy to come by. 20% of Americans think Obama was born in Kenya; 15-20% endorse core parts of the QAnon conspiracy theory; 30% think global warming is a hoax; etc. Should any of these facts lead me to revise my views on these topics?

A reply: remember that we need to restrict the relevant cases of disagreement to disagreement between epistemic peers. And part of being an epistemic peer is having access to the same evidence.

Let's suppose for the moment that Probability-Splitting gives the correct response to cases of peer disagreement. Does this have any practical consequences?

Consider any religious, moral, or political view you have. There would seem to be plenty of people who have the same evidence as you, have thought about the issues as much as you, and are as smart as you, who have a view opposite to yours.

This suggests an argument with massive consequences for what you believe about these domains.

This suggests an argument with massive consequences for what you believe about these domains.

THE DISAGREEMENT-AGNOSTICISM ARGUMENT

```
1. For virtually every moral, political, or religious
   view you have, you have at least roughly as many
   epistemic peers who disagree with you as you have
   epistemic peers who agree with you.
2. If you assign P credence N, and come across an
   epistemic peer who assigns P credence M, then you
   should change your credence in P to the average of N
   and M. (Probability-Splitting)
-------------------------------------------
C. You should not have credence > 0.5 in virtually any
   moral, political, or religious claim.
```

Is it reasonable to deny premise (1)?

Let's consider instead the possibility that we might reject premise (2), Probability-Splitting.

Probability-splitting
If you assign P credence N, and come across an epistemic peer who assigns P credence M, then you should change your credence in P to the average of N and M.

Is this plausible? Let's look at two arguments against this rule of belief.

The first is that the principle is in a certain way self-refuting. There are plenty of people who have thought about disagreement as much as you have who think that the probability-splitting rule is false.

What, given that, does the probability-splitting rule tell you to think about itself?

So there is a sense in which, given actual beliefs of your epistemic peers, this rule of belief is unstable: it recommends against itself. (In this respect, it is similar to foundationalism, and some of the other negative rules of belief we have discussed.)

Probability-splitting
If you assign P credence N, and come across an epistemic peer who assigns P credence M, then you should change your credence in P to the average of N and M.

The second argument is simpler. The main point is that this rule makes the facts about what we ought to believe oddly hostage to the beliefs of others.

It is for that reason a somewhat conservative rule of belief: it argues in favor of thinking what other people think.

Would this make it impossible to be a self-aware radical and to be rational in your beliefs?

The opposite of Probability-Splitting is the view that I should give no weight at all to the fact that people disagree with me. The mere fact of disagreement, after all, is no argument against the truth of any of my beliefs.

The central challenge for this position are the cases mentioned at the outset, like the horse race and check-splitting.

Probability-splitting
If you assign P credence N, and come across an epistemic peer who assigns P credence M, then you should change your credence in P to the average of N and M.

The opposite of Probability-Splitting is the view that I should give no weight at all to the fact that people disagree with me. The mere fact of disagreement, after all, is no argument against the truth of any of my beliefs.

The central challenge for this position are the cases mentioned at the outset, like the horse race and check-splitting.

If you are inclined to think that you should at least become less confident in those cases, but don't think that we should go all the way to Probability-Splitting, the question is how we might stake out a principled middle ground between these two opposing views. Should we reduce confidence, but by less than Probability-Splitting recommends? Or is it the case that we should reduce confidence in some kinds of disagreements, but not others?

Let's turn to our second topic for the day. This is the idea that we should sometimes have a belief for **practical** reasons. If you think about it, forming a belief is just one thing among others that we do. So why wouldn't forming beliefs be governed by the same principles that govern our actions more generally?

This thought is the basis of a very famous argument due to Blaise Pascal. 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 $m \neq 0$.

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.

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."

What does he have in mind here?

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?