deciding to believe

In some cases, it is pretty easy to say what someone should believe. And we have some plausible positive rules of belief which reflect this:

Experience → Belief
If your sense experience tells you that P, and you have no reason to think that your sense experience is misleading, believe P.

Self–Evident → Belief
If P is self-evident, believe P.

Proof → Belief
If you can prove P, believe P.

But none of these rules tell us what to do in hard cases: cases in which we have no experience, or proof, that P is true, and no experience, or proof, that not-P is true. What should we do in these cases?

But none of these rules tell us what to do in hard cases: cases in which we have no experience, or proof, that P is true, and no experience, or proof, that not-P is true. What should we do in these cases?

This is the question taken up by William James in our first reading today. His view is that nothing about the nature of reason or rationality tells us what to do in these cases:

> There are two ways of looking at our duty in the matter of opinion— ways entirely different, and yet ways about whose difference the theory of knowledge seems hitherto to have shown very little concern. *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.

There are two ways of looking at our duty in the matter of opinion—ways entirely different, and yet ways about whose difference the theory of knowledge seems hitherto to have shown very little concern. *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.

James is thinking that the aim of reason is to do two things:

Believe true things.

Avoid believing false things.

But in hard cases, James points out, we face a choice about which of these commandments to follow.

Believe true things. Avoid believing false things.

But in hard cases, James points out, we face a choice about which of these commandments to follow.

Suppose, to take one of the cases he discusses, that I am considering the question of whether God exists. And suppose that for me this choice is what he calls a live one — it would be possible for me to believe that God exists, and possible for me to withhold belief.

Then if I follow the first rule, I will believe that God exists, since that maximizes my chances of believing more true things. If I follow the second rule, I will withhold belief, since that maximizes my chances of avoiding false beliefs.

Which rule should I follow? James' point is that in cases like this one, nothing about the nature of reason or rationality tells me which rule to follow. I know that these are the two rules which should guide me; but nothing tells me which rule to follow in this particular case.

Which rule should I follow? James' point is that in cases like this one, nothing about the nature of reason or rationality tells me which rule to follow. I know that these are the two rules which should guide me; but nothing tells me which rule to follow in this particular case.

Now of course that doesn't mean that reason never tells me which rule to follow. If I have excellent evidence that some claim is true, then the chances of me gaining a truth by believing are much higher than the chances of me believing a falsehood by so doing — so it seems clear that in this case I should follow the first rule. And the opposite for cases in which I have excellent evidence against a claim.

But in some cases I am in neither position. And in these cases, James says, the right question is not what reason tells us to do, but rather the practical question of what sort of person I want to be.

But in some cases I am in neither position. And in these cases, James says, the right question is not what reason tells us to do, but rather the practical question of what sort of person I want to be.

One attitude is the one James ascribes to Clifford. On this view, the right thing to do in every hard case is to withhold belief: we should always privilege the rule which says that we should avoid false belief.

James says in reply:

> Clifford's exhortation has to my ears a thoroughly fantastic sound. It is like a general informing his soldiers that it is better to keep out of battle forever than to risk a single wound. Not so are victories either over enemies or over nature gained. Our errors are surely not such awfully solemn things. In a world where we are so certain to incur them in spite of all our caution, a certain lightness of heart seems healthier than this excessive nervousness on their behalf

Suppose you agree with James that in the kinds of hard cases we have discussed we should not always withhold belief. It is equally true that we should not always believe. For consider the case in which it seems equally possible to me that God exists and that God does not exist. Surely in this case I should not believe both claims.

So when should I believe, and when should I withhold?

Our second reading for today gives an answer to this question. Perhaps we should treat forming a belief like any other action. Perhaps, in that case, we should just form beliefs which will bring about the best results.

Blaise 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 thought that one such decision was the decision whether or not to believe in God:

> Let us then examine this point, and say, "God is, or He is not." But to which side shall we incline? Reason can decide nothing here. There is an infinite chasm which separates us. A game is being played at the extremity of this infinite distance where heads or tails will turn up. What will you wager? According to reason, you can do neither the one thing nor the other; according to reason, you can defend neither of the propositions.

Pascal thought that God so far exceeds our comprehension that we have no way of using our reason to decide whether or not God exists.

But, Pascal thinks, this does not remove the necessity of choosing whether or not to believe in God.

But, Pascal thinks, this does not remove the necessity of choosing whether or not to believe in God.

> —Yes; but you must wager. It is not optional. You are embarked. Which will you choose then; let us see. Since you must choose, let us see which interests you least. You have two things to lose, the true and the good; and two things to stake, your reason and your will, your knowledge and your happiness; and your nature has two things to shun, error and misery. Your reason is no more shocked in choosing one rather than the other, since you must of necessity choose. This is one point settled. But your happiness? Let us weigh the gain and the loss in wagering that God is. Let us estimate these two chances. If you gain, you gain all; if you lose, you lose nothing. Wager them without hesitation that He is.—"That is very fine. Yes, I must wager; but I may perhaps wager too much."—Let us see.

—Yes; but you must wager. It is not optional. You are embarked. Which will you choose then; let us see. Since you must choose, let us see which interests you least. You have two things to lose, the true and the good; and two things to stake, your reason and your will, your knowledge and your happiness; and your nature has two things to shun, error and misery. Your reason is no more shocked in choosing one rather than the other, since you must of necessity choose. This is one point settled. But your happiness? Let us weigh the gain and the loss in wagering that God is. Let us estimate these two chances. If you gain, you gain all; if you lose, you lose nothing. Wager them without hesitation that He is.—"That is very fine. Yes, I must wager; but I may perhaps wager too much."—Let us see.

Pascal is here drawing an analogy between the choice whether or not to believe in God and the choice whether or not to make a bet.

Betting, after all, is another case in which we make decisions under uncertainty.

The study of how it is rational to act under certain kinds of uncertainty is now known as "decision theory." We can use some concepts from decision theory to get a bit more precise about how Pascal's argument here is supposed to work.

We are facing a decision in which we have only two options: belief or nonbelief. And there is one unknown factor which will determine the outcome of our choice: whether or not God exists. So, pairing each possible choice with each possible outcome, there are four possibilities. Our question should be: when faced with a decision like this, what should guide our choice?

We can get clearer on this question by considering a simple bet:

I offer you the chance of choosing heads or tails on a fair coin flip, with the following payoffs: if you choose heads, and the coin comes up heads, you win $5; if you choose heads, and the coin comes up tails, you lose $1. If you choose tails, then if the coin comes up heads, you get $2, and if it comes up tails, you lose $1.

As in Pascal's case, we have a decision with two options - heads or tails - and one relevant unknown - the way the coin will flip will turn out.

I offer you the chance of choosing heads or tails on a fair coin flip, with the following payoffs: if you choose heads, and the coin comes up heads, you win $5; if you choose heads, and the coin comes up tails, you lose $1. If you choose tails, then if the coin comes up heads, you get $2, and if it comes up tails, you lose $1.

We can represent this choice as follows:

Courses of action	Outcome 1: The coin comes up heads	Outcome 2: The coin comes up tails
Choose heads	win $5	lose $1
Choose tails	win $2	lose $1

Courses of action	Outcome 1: The coin comes up heads	Outcome 2: The coin comes up tails
Choose heads	win $5	lose $1
Choose tails	win $2	lose $1

Obviously, given this choice, you should choose heads. One way to put the reason for this is as follows: there is one possibility on which you are better off having chosen heads, and no possibility on which you are worse off choosing heads. This is to say that choosing heads dominates choosing tails.

Courses of action	Outcome 1: The coin comes up heads	Outcome 2: The coin comes up tails
Choose heads	win $5	lose $1
Choose tails	win $2	lose $1

This suggests the following rule for rational decision making:

The rule of dominance
If you are choosing between A and B, and A dominates B, you should choose A

Some passages in Pascal's argument suggest that he had this sort of rule in mind.

—Yes; but you must wager. It is not optional. You are embarked. Which will you choose then; let us see. Since you must choose, let us see which interests you least. You have two things to lose, the true and the good; and two things to stake, your reason and your will, your knowledge and your happiness; and your nature has two things to shun, error and misery. Your reason is no more shocked in choosing one rather than the other, since you must of necessity choose. This is one point settled. But your happiness? Let us weigh the gain and the loss in wagering that God is. Let us estimate these two chances. If you gain, you gain all; if you lose, you lose nothing. Wager them without hesitation that He is.—"That is very fine. Yes, I must wager; but I may perhaps wager too much."—Let us see.

The rule of dominance
If you are choosing between A and B, and A dominates B, you should choose A

Some passages in Pascal's argument suggest that he had this sort of rule in mind.

Pascal's claim that if you lose, you lose nothing is some indication that he thought that belief dominated non-belief; the thought would be that in the case where God exists (i.e., where you win), you are better off believing, and that in the case where God does not exist (i.e., where you lose) you are no worse off.

If Pascal is right about this, then we might represent our decision about whether or not to believe in God as follows:

Courses of action	God exists	God does not exist
believe	win everything	win nothing, lose nothing
don’t believe	win nothing, lose nothing	win nothing, lose nothing

If this is the correct representation of our choice whether or not to believe, then belief dominates non-belief. Since the rule of dominance seems very plausible, this would be a very powerful argument that we rationally ought to believe in God.

Courses of action	God exists	God does not exist
believe	win everything	win nothing, lose nothing
don’t believe	win nothing, lose nothing	win nothing, lose nothing

But is this the correct representation of the choice?

In the case of false unbelief, wouldn’t one be undertaking religious obligations which one might have avoided? And isn’t having a false belief something bad in itself?

Suppose, then, that we make a small change in our representation of the choice.

Courses of action	God exists	God does not exist
believe	win everything	lose something
don’t believe	win nothing, lose nothing	win nothing, lose nothing

Suppose, then, that we make a small change in our representation of the choice.

Now does the rule of dominance tell us to believe in God?

However, there is some indication that this is not the best interpretation of Pascal’s argument anyway.

However, there is some indication that this is not the best interpretation of Pascal's argument anyway.

"…since there is an equal chance of gain and loss, if you stood to win only two lives for one you could still wager, but supposing you stood to win three?

…it would be unwise of you, since you are obliged to play, not to risk your life in order 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, one chance of winning against a finite number of chances of losing, and what you are staking is finite. That leaves no choice; wherever there is infinity, and where there are not infinite chances of losing against that of winning, there is no room for hesitation, you must give everything. 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, just as likely to occur as a loss amounting to nothing."

Here Pascal is thinking of bets where you might win or lose something by playing, but where what you win is greater than what you lose. But in bets of this sort, dominance reasoning will often not tell us whether or not to take the bet, since it will not be the case that taking the bet will never leave you worse off than not taking it.

Let's consider how we might reason about decisions of this sort, where it is not the case that one option dominates the others, and so where the rule of dominance does not tell us what to do.

Let's consider how we might reason about decisions of this sort, where it is not the case that one option dominates the others, and so where the rule of dominance does not tell us what to do.

I'm about to flip a coin, and offer you the following bet: if the coin comes up heads, then I will give you $5; if it comes up tails, you will owe me $3. You know that it is a fair coin. Should you take the bet?

Courses of action	Heads	Tails
take the bet	$5	-$3
don't take the bet	$0	$0

Courses of action	Heads	Tails
take the bet	$5	-$3
don't take the bet	$0	$0

Here neither course of action dominates the other; but it still seems that you should clearly take the bet. Why?

There is a ½ probability that the coin will come up heads, and a ½ probability that it will come up tails. In the first case I win $5, and in the second case I lose $3. So, in the long run, I'll win $5 about half the time, and lose $3 about half the time. So, in the long run, I should expect the amount that I win per coin flip to be the average of these two amounts — a win of $1.

Here neither course of action dominates the other; but it still seems that you should clearly take the bet. Why?

There is a ½ probability that the coin will come up heads, and a ½ probability that it will come up tails. In the first case I win $5, and in the second case I lose $3. So, in the long run, I'll win $5 about half the time, and lose $3 about half the time. So, in the long run, I should expect the amount that I win per coin flip to be the average of these two amounts — a win of $1.

We can express this by saying that the expected utility of taking the bet is $1. It seems that one should take this bet because the expected utility of doing so is greater than the expected utility of not taking the bet.

To calculate the expected utility of an action, we assign each outcome of the action a certain probability, thought of as a number between 0 and 1, and a certain value (in the above case, the relevant value is just the money won). In the case of each possible outcome, we then multiply its probability by its value; the expected utility of the action will then be the sum of these results.

Let’s see how this looks by returning to our simple bet.

Courses of action	Heads	Tails	Expected utility
take the bet	$5	-$3	.5 * $5 + .5 * (-$3) = $1
don’t take the bet	$0	$0	.5 * $0 + .5 * $0 = $0
	Probability = 0.5	Probability = 0.5	

The higher expected utility of taking the bet seems to explain why this would be the right move.

The higher expected utility of taking the bet seems to explain why this would be the right move.

Reflection on this sort of example seems to make the following principle about rational action seem quite plausible:

The rule of expected utility
It is always rational to pursue the course of action with the highest 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, as the example of the bet illustrates, tells us what we should do in certain situations about which the rule of dominance is silent. So, even if we think that belief does not dominate non-belief, perhaps we can use the rule of expected utility to reconstruct Pascal's argument.

Let's return to the passage discussed above.

> "…since there is an equal chance of gain and loss, if you stood to win only two lives for one you could still wager, but supposing you stood to win three?
>
> …it would be unwise of you, since you are obliged to play, not to risk your life in order 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, one chance of winning against a finite number of chances of losing, and what you are staking is finite. That leaves no choice; wherever there is infinity, and where there are not infinite chances of losing against that of winning, there is no room for hesitation, you must give everything. 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, just as likely to occur as a loss amounting to nothing."

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 concede the point made above in connection with dominance reasoning: 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	.....

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+......

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:

> "...is there really no way of seeing what the cards are? ...I am being forced to wager and I am not free; I am begin held fast and I am so made that I cannot believe. What do you want me to do then?"
>
> "That is true, but at least get it into your head that, if you are unable to believe, it is because of your passions, since reason impels 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.

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