The Two-Envelope Paradox
Author(s): John Broome
Source: *Analysis*, Vol. 55, No. 1 (Jan., 1995), pp. 6-11
Published by: Oxford University Press on behalf of Analysis Committee
Stable URL: http://www.jstor.org/stable/3328613
Accessed: 06-10-2015 22:02 UTC

The Two-envelope Paradox

JOHN BROOME

On the table in front of you are two envelopes. You have been told by an impeccable source that each contains a cheque, and that one cheque is twice the other. But you have no idea which envelope contains the larger cheque. You may choose one or the other envelope, and keep the cheque it contains.

Suppose that, having made a choice, you are given a chance to change your mind and switch to the other envelope instead. You might reason as follows. Let the amount of the cheque in your chosen envelope be x. Whatever x may be, there is a probability of 1/2 that the cheque in the other envelope – call it y – is $2x$ and a probability of 1/2 that it is $x/2$. So, whatever x, the expectation of y is $(1/2)2x + (1/2)(x/2)$. This comes to $5x/4$, which is more than x. So the expected value of the cheque in the other envelope is more than the value of the cheque in yours. So you should switch, whatever x may be. Consequently, you should definitely switch even though at present you do not know what x is. But this is paradoxical, because the same reasoning would lead you to switch back again if you did switch.

Frank Jackson, Peter Menzies and Graham Oppy [3] offer a solution to the paradox. They say you must have some prior probability distribution over the size of the cheques in the envelopes. Then a particular value of x will, given this prior distribution, determine probabilities for the events $y = 2x$ and $y = x/2$, conditional on x. These probabilities will not be 1/2 for every value of x. Consequently, they claim, the expectation of y, conditional on x, will not be greater than x for every value of x. It is not true that *whatever* the value of x, the expectation of y is greater. So the paradox disappears.

This is only the beginning of a solution, not the end. Jackson, Menzies and Oppy do not actually prove their claim that the expectation of y, given x, is not greater than x for every x. As it happens, it is not necessarily true. There are prior probability distributions such that, whatever the value of x, the expectation of y conditional on x is greater than x. Let us call such a distribution 'paradoxical'. Barry Nalebuff ([4], p. 187) has described one paradoxical distribution,[1] and I shall describe two more.

Let s be the value of the smaller of the two cheques. For my first example, suppose you assign prior probability 1/3 to $s = 1$, 2/9 to

[1] There are some slips in Nalebuff's calculations, but the example turns out to be correct. My first example is similar to his.

ANALYSIS 55.1, January 1995, pp. 6–11. © John Broome

$s = 2$, $4/27$ to $s = 4$, $8/81$ to $s = 8$, and so on. That is to say, your prior distribution assigns probability $2^n/3^{n+1}$ to $s = 2^n$ when n is a non-negative integer, and probability zero to any other amount.

Suppose the value of the cheque in your envelope is x. If x happens to be 1, it is certain that y is $2x$. In this case the expectation of y, given x, is definitely greater than x.

If x has some other value, what is then the probability that y is $2x$? y will be $2x$ if and only if x is the smaller of the two cheques, s. Given that your cheque is x, there are only two possibilities: either s is x, or s is $x/2$. So the probability we need is the probability that s is x, given that s is either $x/2$ or x. That is: $Prob(s = x \mid s = x \vee s = x/2)$. By the definition of conditional probability, this is:

$$\frac{Prob(s = x \;\&\; (s = x \vee s = x/2))}{Prob(s = x \vee s = x/2)} = \frac{Prob(s = x)}{Prob(s = x \vee s = x/2)}$$

$$= \frac{2^n/3^{n+1}}{2^n/3^{n+1} + 2^{n-1}/3^n} = \frac{2}{5}$$

where n is the integer such that $x = 2^n$. So the probability that y is $2x$ is $2/5$. The probability that y is $x/2$ (the only other possibility) is $3/5$. So the expectation of y, given x is:

$$E(y \mid x) = (2/5)2x + (3/5)x/2 = 11x/10$$

This is greater than x for all values of x. I have shown that, whatever the value of x, the conditional expectation of y is bigger. This distribution is paradoxical, then.

My second example is a continuous distribution, so for this one imagine the currency is continuously divisible, like atomless gold. Once again, let s be the value of the smaller cheque, and suppose the prior distribution you assign to s has a density function $f(s)$. It turns out that, when x is the value of the cheque in the envelope you have chosen, the expectation of y, the value of the other cheque, conditional on x, is:

$$E(y \mid x) = x\,\frac{2f(x) + (1/4)f(x/2)}{f(x) + (1/2)f(x/2)}$$

This formula is derived in Appendix A below.[2] Now suppose your specific density function is:

$$f(s) = 1/(s + 1)^2 \quad \text{for } s > 0$$

Substituting this function in the formula gives:

$$E(y \mid x) = x\,\frac{2(x+2)^2 + (x+1)^2}{(x+2)^2 + 2(x+1)^2}$$

[2] An incorrect version of this formula appears in Castell and Batens [2], p. 47.

Since $(x + 2)$ is bigger than $(x + 1)$, the fraction in this formula is bigger than 1. Consequently, $E(y|x)$ is bigger than x for all x. This distribution is also paradoxical, therefore.

If you have a paradoxical prior probability distribution, then for every value of x, there is a positive expectation of gain from switching envelopes. This appears to be a valid argument for switching. To emphasize the paradox, I should mention that there also appears to be a valid argument that shows you have no reason to switch.[3] Your envelope contains either s, the smaller of the two cheques, or $2s$, and these possibilities are equally likely. If your envelope contains s and you switch, you will gain s. If your envelope contains $2s$ and you switch, you will lose s. So your expected gain from switching, conditional on s, is zero, and this is true for all s. Therefore, you have no reason to switch. We have two apparently valid arguments with conflicting conclusions, then. Some people uphold the second argument and repudiate the first on the grounds that s is a constant, whereas x is a random variable. But this is not so. Once you have chosen your envelope, both s and x are fixed but unknown quantities.

Both my examples of paradoxical distributions are decent statistical distributions. There are processes for generating either: processes that will determine the value of the smaller cheque, giving each possible value just the probability that is specified in the distribution. The discrete distribution can be generated by a modified St Petersburg process. Take a coin that has 1/3 probability of falling heads and 2/3 probability of falling tails. Toss it once. If it falls heads, let s be 1. If it falls tails, toss again. If it falls heads the second time, let s be 2. If it falls tails, toss again. If it falls heads the third time, let s be 4. If it falls tails, toss again. And so on. This process assigns the required probabilities to s. My continuous distribution can be generated as follows. Pick a number t at random from a uniform distribution between 0 and 1, and let s be $t/(1-t)$. Then s has the required distribution.

Although my distributions are valid, they do have the distinctive feature of possessing no finite mean. Appendix B shows that no distribution with a finite mean (even an infinite distribution) can be paradoxical. When a distribution has no finite mean, it can be used to generate a different paradox: the St Petersburg paradox. Suppose that, for a price, you are offered the chance of playing a game in which your winnings will be an amount of money determined at random, with probabilities given by one of my distributions. How much would you be prepared to pay for this opportunity? Since the distribution has no finite mean, the expectation of your winnings exceeds any finite amount. So it seems you should be prepared to pay any

[3] This argument has been put to me by, amongst others, Douglas MacLean, Louis Marinoff and Hilmar Schneider.

amount of money, however large, to play this game. Yet in fact no one would be willing to pay very much. That is the St. Petersburg paradox.

In one respect, the two-envelope paradox is more powerful. Part of the reasoning that leads to the St Petersburg paradox is the claim that the expectation of your winnings from the game exceeds any finite amount. But strictly there is no such thing as the expectation of your winnings. Strictly, a distribution with no finite mean does not have a mean that exceeds any finite amount; it has no mean at all. So there is a weakness in the reasoning that leads to the St Petersburg paradox. But the reasoning that leads to the two-envelope paradox is free from this weakness. The only expectations it refers to are the expectations of what is in one envelope conditional on what is in the other, and these are definitely finite. True, a paradoxical distribution has no finite mean, but the reasoning that leads to the paradox does not refer to the mean of the distribution.

Still, because the two paradoxes arise from a similar phenomenon, the recognized solutions to the St Petersburg paradox can also serve as solutions to the two-envelope paradox. One solution is to point out that there is only a limited amount of money in the world, whereas the paradoxical distributions presume an unlimited supply. Another solution was originally Daniel Bernoulli's [1]. It relies on risk aversion about money. In presenting the two-envelope paradox, I assumed that, of two options, you would always prefer the one that offers you a greater expectation of money. This implies you are risk neutral about money; if you were risk averse, you would sometimes prefer a safe option to a risky one even if the safe option offers a lower expectation of money. It turns out that, if you are risk averse *enough*, neither the St Petersburg paradox nor the two-envelope paradox will arise. I do not need to go into the details.

This latter solution to the two-envelope paradox is recommended by Barry Nalebuff ([4], p. 187) It may be the best solution available, but it is unsatisfactory. Rationality does not seem to require you to be risk averse. Yet if you are not, you will still be subject to the paradox. Imagine you are the angel in charge of the happiness of some world. The cheques in the envelopes are drawn by God on the happiness bank, and the cheque you choose determines the amount of happiness there will be in your world. Suppose the possibilities for happiness are unlimited. There is surely no reason why you should be risk averse about your world's happiness. So if you believe God has determined the size of the cheques by one of the distributions I have described, you will be caught by the paradox.

University of Bristol
Bristol BS8 1TN, UK
John.Broome@bristol.ac.uk

Appendix A

Let s be the smaller of the two numbers in the envelopes, and assume your prior probabilities can be represented by a density function $f(s)$. This will be so if the distribution is continuous. This appendix derives a formula for the conditional expectation $E(y \mid x)$ in this case.

I shall first find the probability that $y = 2x$, given the value of x. There is always a zero probability that x has a particular given value. So, in order to avoid dividing by zero, I shall start by finding the probability that $y = 2x$, conditional on x lying in a small interval between z and $z + \Delta z$.

If x lies between z and $z + \Delta z$, then either s lies between z and $z + \Delta z$, or s lies between $z/2$ and $z/2 + \Delta z/2$. (I take Δz to small enough that these intervals do not overlap.) The probability that y is $2x$ is the probability that the former is the case, since in that case x must be s. So

$$Prob(y = 2x) = Prob(z \leq s < z + \Delta z \mid (z \leq s < z + \Delta z \vee z/2 \leq s < z/2 + \Delta z/2)$$
$$= \frac{Prob(z \leq s < z + \Delta z)}{Prob(z \leq s < z + \Delta z) + Prob(z/2 \leq s < z/2 + \Delta z/2)}$$

By the definition of the density function, this is approximately

$$= \frac{f(z)\Delta z}{f(z)\Delta z + (1/2)f(z/2)\Delta z} = \frac{f(z)}{f(z) + (1/2)f(z/2)}$$

and the approximation becomes precise as Δz tends to zero. When Δz tends to zero, this formula becomes the probability that $y = 2x$ conditional on x being equal to z. To put it another way:

$$Prob(y = 2x \mid x) = \frac{f(x)}{f(x) + (1/2)f(x/2)}$$

Similarly

$$Prob(y = x/2 \mid x) = \frac{(1/2)f(x/2)}{f(x) + (1/2)f(x/2)}$$

The expectation of y given x is therefore:

$$E(y \mid x) = (2x)Prob(y = 2x \mid x) + (x/2)Prob(y = x/2 \mid x)$$
$$= x\frac{2f(x) + (1/4)f(x/2)}{f(x) + (1/2)f(x/2)}$$

Appendix B

This appendix proves that, if you have a prior distribution that has a finite mean, the paradox will not arise: the expectation $E(y \mid x)$ of y conditional on x will not be bigger than x for all x.

Consider your prior bivariate distribution over x and y. It must be symmetrical, because the envelopes reveal no clue about which contains the bigger number. Suppose this distribution determines a finite mean $E(x)$ for x and a finite mean $E(y)$ for y. By symmetry, these means must be the same. For any value of x, there will be an expectation $E(y \mid x)$ of y conditional on x. I need to prove that $E(y \mid x)$ is not greater than x for all x.

To put it another way, I need to prove the difference $E(y \mid x) - x$ is not always positive. This difference is a function of x, and since x is a random variable it has an expectation:

$$E(E(y \mid x) - x)$$

The expectation of a difference is the difference in the expectations, so this is:

$$E(E(y \mid x)) - E(x).$$

Furthermore, the expectation of a conditional expectation of a variable is the unconditional expectation of the variable. This is a standard theorem that bears the main weight of my proof; unfortunately, I cannot prove it here. (See, for instance, Parthasarathy [5], p. 225.) So this is:

$$E(y) - E(x).$$

I have already said that $E(y)$ and $E(x)$ must be the same because of symmetry, so $E(y) - E(y)$ is zero. On average, then, the difference between the expectation of y conditional on x and x itself is zero. Therefore, if this difference is positive for some value of x it must be negative for some other value.[4]

References

[1] Daniel Bernoulli, 'Specimen theoriae novae de mensura sortis', *Commentarii Academiae Scientiarum Imperialis Petropolitanae* 5 (1738), translated by Louise Sommer as 'Exposition of a new theory on the measurement of risk', *Econometrica* 22 (1954) 23–36.

[2] Paul Castell and Diderik Batens, 'The two-envelope paradox: the infinite case', Analysis 54 (1994) 46–49.

[3] Frank Jackson, Peter Menzies and Graham Oppy, 'The two envelope "paradox"', *Analysis* 54 (1994) 43–45.

[4] Barry Nalebuff, 'The other person's envelope is always greener', *Journal of Economic Perspectives* 3 (1989) 171–81.

[5] K. R. Parthasarathy, *Introduction to Probability and Measure* (London: Macmillan, 1977).

[4] This paper arose out of a conversation I had with Doug MacLean on the terrace of the Parliament Building in Canberra. MacLean saw the connection between the two-envelope paradox and the St Petersburg paradox long before I did. The paper owes a lot to my conversations with several other people too, particularly with Ian Jewitt, Graham Oppy and Denis Robinson. It was mostly written while I was a visitor at the Australian National University; I am very grateful to the ANU for its hospitality.