

CAN TWO ENVELOPES SHAKE THE FOUNDATIONS OF DECISION-THEORY?*

**Olav Gjelsvik,
University of Oslo.**

The aim of this paper is to diagnose the so-called two envelopes paradox. Many writers have claimed that there is something genuinely paradoxical in the situation with the two envelopes, and some writers are now developing non-standards theories of expected utility. I claim that there is no paradox for expected utility theory as I understand that theory, and that contrary claims are confused. Expected utility theory, or theory of decision under uncertainty, was first developed by Frank Ramsey, and in its essence it only relies on the basic laws of probability. It is furthermore a standard prescription to think of utilities as bounded, and not as infinite. (Making this prescription matters for the mathematics, but the prescription has also strong intuitive support.¹ Another consideration is that we will never face unbounded utilities in any choice.) I shall follow this prescription throughout this paper, but I shall make paranthetical remarks about the infinite case where that prescription is lifted. The two envelope paradox seems to me not to have anything essentially to do with infinity. Perhaps the two-envelope set-up has brought writers in contact with the problems concerning infinite utilities, but those problems should be approached straight on and not in a roundabout way. My main concern is whether there is a problem here for standard expected utility-theory with bounded utilities. Throughout the paper I assume that we are fully risk-neutral.

The choices

* I am much indebted to Aanund Hylland and Ole Jørgen Skog for penetrating discussions, also in writing, about the two envelope problem and the contents of this paper. I am also indebted to John Broome, Tim Williamson, Wlodek Rabinowics, Søren Halldén, and audiences at the Universities of Lund and Oslo, and Oriel College, University of Oxford.

¹ About these points, see Lindley, D.V. **Making Decisions**, chapter 1.

It is important to see the exact logical and informational structure of the problem. We must keep apart the situation with unopened envelopes, where you have no knowledge of the actual amounts in any envelope, and the situation I shall call opened envelope, where you know the actual amount of one of the envelopes, namely the opened envelope. I shall first turn to unopened envelopes.

A: Unopened envelopes

You are given the choice between two envelopes and are told that they contain one cheque each. One envelope contains a cheque with twice the amount of the cheque in the other envelope, but you do not know which envelope contains the larger amount. For the sake of clarity I shall assume that the situation has come about in two independent steps: First, the amounts which go into the envelopes have somehow been decided upon. Secondly, it has been decided by a fair coin which envelope contains the larger amount. Think of this as a one-off situation, the kind of situation which is basic in expected utility-theory.

One clear intuition is that there is nothing which differentiates in the choice between the two envelopes, A and B. We have no way of knowing what the selected amount is, and we have no way of knowing which of the two envelopes contains the larger cheque, and which contains the smaller. Since a fair coin has been used in deciding which envelope contains the larger cheque, it is equally probable that A contains the larger amount and that it contains the smaller amount. These two possibilities are the relevant possibilities in the case of A. The same goes for envelope B. The expected utility of choosing A and of choosing B must ought therefore to be the same. If the smaller amount is represented by the letter "z", then the expected utility of choosing A (or B) can simply be represented as $(0.5 z + 0.5 2z) = 1,5 z$. To repeat: Relative to this way of structuring the problem there are two relevant equiprobable states of both A and B. These states are linked in this way: If A contains the larger cheque, B contains the smaller, and vice versa.

Of course there is a well-known piece of reasoning in support of the conclusion that when you have chosen one unopened envelope it is always

rational to swap to the other. Grass is always greener on the other side, or so it seems to be. This is the paradoxical reasoning. The reasoning goes like this:

The envelope you have selected, A, contains a cheque with a certain amount written on it. Let us call this amount it contains x . Let us simply say that the expected utility of not swapping is x .

Either envelope B contains a cheque with the amount $2x$, or it contains an envelope with the amount $x/2$. The one situation is as likely as the other. The expected utility of swapping is then $1,25 x$! ($.5 \times 2x$ plus $.5 \times x/2$). Of course you should swap!

Of course, we can carry out a similar reasoning in favour of choosing A. Let us call the amount on the cheque in B y . The expected value of A is $1,25 y$ by the same type reasoning. This is an argument for swapping back.

This is paradoxical. " $x=1,25y$ ", and " $y=1,25 x$ " cannot both be true. What has happened here?

The short answer is that there is, in this situation, no expected utility like x or y . I shall develop this answer in answer in a cautious way. There is a way in which we do not keep track of the original epistemic sensitivities when we generate paradox. It is epistemically relevant for the original decision problem that irrespective of the total amount of money involved, both A and B may be in one of two possible states; both may contain the larger or the smaller amount. B contains the larger if A contains the smaller and vice versa. If we calculate expected utility with **the amount A contains as the unit**, and we simply assign this value to A, then we risk that we no longer pay proper attention to the two relevant epistemic possibilities in the case of A. Remember that the two states of A have links to the two states of B. (The case of B is parallel.)

We have to keep clearly apart two different cases: The case where both A and B can be in two different states (but always "opposite of each other"), **and the case where A stays in one state, and B can be in two possible states, either twice as big as A or half the size of A**. When envelopes are unopened, our case is the former case: There are two relevant states of A, linked to two relevant states of B, and not just one of A and two of B.

Let clarify this further, and focus on the two descriptions

"The amount A contains".

"The smaller amount"

Both descriptions can be read with wide scope or with narrow scope if we quantify into modal contexts. A decision matrix is nothing but an overview of the relevant epistemic possibilities or the epistemically relevant possible worlds. To do expected utility theory we have to assign probabilities to each such epistemically possible and relevant outcome (utility).

Consider the two possible decision matrices. In the first matrix, I, the variable, the unit for calculating expected utility, is introduced and fixed as a variable by the description "The smaller amount". In the second case, matrix II, the variable is introduced and fixed by the description "The amount in A".

I.	States of nature:	B contains $2z$	B contains z
Strategy	Stay with A	Get z	Get $2z$
	Swap	Get $2z$	Get z

II	States of nature	B contains $2x$	B contains $1/2 x$
Strategy	Stay with A	Get x	Get x
	Swap	Get $2x$	Get $1/2x$

Look first at the description "the amount in A" and the corresponding decision matrix (Matrix II.). The description **has** to be read with narrow scope to respect properly the epistemic situation we are in when both A and B each can be in these two different ("opposite") states respectively, (containing either the larger or the smaller amount, (and always the "opposite" of the other envelope)). But it **has** to be read with wide scope in order to be seen as fixing a unit for calculating expected utility throughout all worlds relevant for the decision-problem. We cannot otherwise coherently matrix (like for instance II). But we cannot have it both ways when it comes to scope. (We could also go through this reasoning with "The amount in B".)

The justification for saying that a description **has** to be read with wide scope to supply a unit for calculating expected utility is this: When representing a decision problem, we must fix the unit, to be represented by a variable, by which we calculate expected utility. This fixing must make it possible to let this variable (in an instantiation, for instance with a monetary sum) be seen as standing for one and the same particular value through all the (epistemically) possible worlds relevant for calculating expected utility. By reading "The amount in A" with wide scope, we capture as a unit the amount A actually contains. If we then proceed to represent A as having this very content throughout the decision-problem, **and that is what we do when use the reasoning above which leads to paradox**, then we lose track of the two relevant possible states of A, the one state where it contains the larger, and the other where it contains the smaller amount. Then there is no longer a symmetry between the choices, we are structuring the situation like the case where there is an amount put in A, and either half or the double amount is put in B. **That is, however, not our situation.** If we on the other hand give the description "The amount in A" narrow scope to capture the two possible states of A, and try to think of the description as introducing a unit for calculating expected utility, then we are not able to calculate expected utility properly: We are no longer calculating by a unit representable by a variable we can see as having the same value in an instantiation (with a monetary value) throughout the relevant possible worlds.

There is clear contrast to the case of the description "The smallest amount". If this description, with a wide scope reading, is used to fix the unit by which we calculate , as we did in matrix I, there is no problem in representing the

original decision problem. This is a problem where A either contains z or $2z$, and B contains $2z$ or z . We can instantiate this variable, give it the value "a", and there would be no problem in representing z as having this value throughout the relevant epistemic possibilities.

Concluding diagnosis: We have various options for how to represent and structure the decision problem we are facing. This structuring is partly done by using descriptions (quantifier phrases) to introduce variables. We can in this case introduce a variable by various descriptions, by the description "smallest amount", (or "the larger amount"), and by "the amount in A" and "the amount in B". The first way of structuring this problem, by the description "the smaller amount", is a representation which captures the original decision problem., and it introduces a variable which can be instantiated throughout the relevant possible worlds. The calculation which supports the paradoxical conclusion introduces a unit for calculating expected utility by means of the description "The amount A contains". (or "...B contains") This introduction of a unit for calculating expected utility is illegitimate in the case of the unopened envelopes. In our example it appears to make us blind to the two different relevant states of A; to the fact that A contains the larger amount when B contains the smaller, and contains the smaller amount when B contains the larger. Openness towards the relevance of both these states of A in the unopened case, on the other hand, leaves us without a properly fixed variable for calculating expected utility if we try to fix it by "the amount in A"..

You should remain indifferent between A and B as long as they are unopened, or so I conclude. This result generalizes from this one-off case to cases of repeated choice.

B: Opened envelope.

But, we may ask, what if I open one of them, namely A, and find for instance GBP100 in it? The thing to note first, is that we receive crucial information when we learn the amount in A. Before the opening of A, there are very many possibilities for what the amount in A may be. After the opening, the number of possibilities is reduced to one. If we knew how many possibilities there were to begin with, we could in favourable cases calculate how many bits of information the opening of envelope A carries

with it (by applying Shannon and Weavers mathematical theory of information.).

How does this reception of information matter for the decision-problem? One reply is this: We receive a lot of information upon opening, but we do not receive information as to whether A contains the smaller or the larger amount. Since this bit of information is not received, we have to structure the decision problem by using these descriptions. We should simply regard the information we have received as not relevant for the decision problem at hand. That means that the basic situation is unchanged, and that we still should remain indifferent. I shall put aside this for the moment, and discuss it later.

The other reply is this: The information we receive about the amount in A changes the structure of the decision-problem. After this information is received, there is one and only one (epistemically) relevant state of A. Secondly, since A contains GBR 100, we know for certain that B either contains GBP 50 or GBP 200.

Should I on this line accept an offer of swapping to B? We should say: That depends. That depends on what we think the probability is that B contains GBP 50, and on what we think the probability is that B contains GBP 200. Expected utility theory will calculate the expected value of each option on the basis of such probabilities. Since the answer from expected utility theory depends on these probabilities, a fundamental prior question is whether we can at all assign probabilities to these possible states. If we cannot at all assign probabilities, we cannot use expected utility-theory when making the choice about swapping or not. Of course we can nevertheless use well-known decision-strategies. Maximin is an example of such a strategy. If we use maximin, we will decide not to swap.

The point to be clear about, is that we need to be able to assign probabilities in order to apply expected utility theory. And only when we can apply expected utility-theory, does this theory at all face the possibility of paradox. Of course, when we actually assign probabilities, our assignments may be empirically wrong. That is, however, no conceptual problem.

There are two further points to be made here. We work with the simplification that utilities might be directly represented by money. Since

utility is bounded, we then make the assumption that money is finite. We could work with infinite sums of money, and let only utilities, and not money be bounded. But the simplification is unproblematic. Also, in order to apply expected utility-theory to this decision-problem, we must somehow have an a priori probability density-function, describing the distribution of probability for a continuous random variable like the content of the envelope with the smallest amount (above called z). When we know the actual amount of the opened envelope A , we can then relative to the density-function work out how probable it is that the other envelope contains half the amount of A and how probable it is that it contains twice the amount of A .

An important property of a probability density-function is best seen when viewed as a graphic representation: the total area under the density curve must be 1; such a function represents how the total probability of 1 is distributed over the range of possible values for the random variable z . The probability for z to be between two values, a and b , is represented by the area under the density curve between a and b . The density is either positive or 0. (With a continuous random variable, it strictly speaking only makes sense to speak of the probability that z lies in an interval.) In case z (the smallest number) is picked randomly from a certain finite range, as it is in this case, and we know the highest possible value of z , we will definitely not want to swap in case the opened envelope contains more than the highest possible value of z . In case we do not know the highest possible value of z , we might have a hunch or make a guess at that, and make our choices about swapping on that basis. As long as there is a finite highest value for the amount in the envelope, the density-function will reach 0, and we will never know before opening of an envelope whether we will want to swap after opening.

Several writers on the two-envelope problem have discussed the infinite case, and some have pointed to parallels between the two envelope case and the St. Petersburg paradox. But we put aside infinite expectations, as is standardly done. In the normal case of a practical decision-maker facing the two envelopes problem, the value of the density-function will be 0 for very high values of z . It has been said in discussions of the two-envelope problem that we should not bother with infinite values of z because there is only a finite amount of money in the world anyway. It is true that there is only a finite amount of money in the world. This observation about the amount of money in the world has its role in justifying to ourselves that we can put

very high values of z aside as having no positive probability. The conceptual point is that we should treat utilities as bounded.

(In case it is supposed that z ranges from 1 to infinity, and there are infinite utilities, there are possible density-functions which are such that relative to them, it will always be rational to swap to the unopened envelope according to expected utility theory, as shown by Nalebuff, Broome and others. What this means, is that relative to such density-functions you will be able to know beforehand that you will swap. That might create a danger of regenerating the paradox for those special cases: You know that if A is opened you will prefer B, and if B is opened, you will prefer A. That, however, is not really paradoxical. What we have then is not a situation where opposing preferences are both rational for a person at a time, as it was above. We have an indifferent person who knows that in the light of some information she/he will prefer B, and in the light of a different piece of information there will be a preference for A. The preferences one then forms is relative to some information received, and clearly stable relative to that information. Note also that what is required for this particular type of situation to arise is that a) the amounts in the envelopes is known to grow to infinity, (or that utilities are infinite/unbounded) b) that a special type density-function is held to hold, c) one receives information about what is in one envelope, but not in the other, and one's information is partial in this sense. The situation we are facing here is perhaps odd and puzzling, but not paradoxical. It illustrates the caution with which we have to treat the case of unbounded utilities.)

Has the situation really changed essentially upon opening?

There is one final issue we have to deal with. A consideration recently put forward by Clark and Shakel (in **Mind**) against the view that the situation changes essentially upon opening is this: Imagine that you and I in this situation, the two envelope situation, get to look into one envelope each. (We are then both in open envelope situation, but look into "opposite" envelopes.) You have seen what is in envelope A and prefer B. I have seen what is in B and I prefer A. A third party offers us the option of paying him a 10% tax on our perceived gain (not on the actual gain of course) for each swapping. We pay and swap repeatedly. We would still both prefer to swap and pay this tax on each swapping. We, you and I, would both lose money in the long run, Clark and Shakel claim. We would both pay money to a third party, and there would be no more money around than what is in the

envelopes. Conclusion: It is rational to disregard to information we receive upon opening an envelope. ² Therefore, Clark and Shackel argues, the situation has not changed essentially.

This reasoning is not acceptable for the standard case with finite expectations. We must in fact tread carefully here, as long as we think of utilities as just given by monetary amount, and we think of the probability-estimates as correct. (It is not news that false beliefs might lead to losing a bet, and the same goes for wrong probability-estimates.) In the case of opened envelope, where the expected value on swapping is positive, then you might clearly gain by swapping in that case. The same might go for two or three or many repeated such openings. Eventually we would not pay any tax - that all relates to the density-function. No good reason has been given for thinking that we would necessarily be worse off by always following expected utility in this type of case as long as we restrict ourselves to finite expectations, as I do, and correct probability-estimates. If we allow infinite expectations in this situation, we get a situation I will not go into: we might face the task of ranking various infinite sums against each other.

It has also been argued by Clark and Shaket that "If all those cases where you have 2 in your envelope are picked out, then the average gain for them (in swapping) would be positive. In considering the average gain for a given value in your envelope we are not considering a truly representative example, one for which we are as likely to have the larger sum in our envelope as the smaller".³

Consider the claim that us having the larger sum in our envelope must be as likely (as probable) as us having the smaller. Why should it be thought to be so after the opening of the envelope? The whole point about the density-function would be to provide an answer to exactly how likely that is; we can work out the answer from the given density-function and the information we receive. What we actually know is that relative to the type of density-function Clark and Shaket are discussing, (they are thinking about the infinite case when it always rational to swap), it is never equiprobable after opening that the amount we find there is as likely to be the smaller as the larger amount. (If it were, we would soon come across cases where it is not

² By Clark and Shackel (2000) p. 429.

³By Clark and Shackel (2000) p. 430.

rational to swap.) The answer in the cases considered is therefore that it is never equally likely that the unopened envelope contains the smaller or the larger amount, and it is precisely the case that the likelihood of the amount in the opened envelope being the smaller or the larger must relate to the probability density-function one relies upon. The premise about equiprobability seems necessary to reach the conclusion that nothing has changed essentially, but that premise must be rejected, and our conclusion stands.

Conclusions in the opened case

The opening of an envelope gives us information which changes the structure of decision-problem. We might in real life hold that we are not able to assign probabilities in the opened case. This might be so when we have no information at all about how the amounts in the envelopes have been picked. In that case we would not make the choice about swapping according to expected utility-theory, but in some other way.

If we can assign probabilities, and that is an assumption, then we would make our choice on the basis on some probability-density function. In the normal case, we would assign probabilities on the basis of various types of a priori knowledge, our knowledge of the people involved, etc. (If it is assumed that utility is represented by amount of money, that the amount of money in the envelopes can be infinite, and that a specific type of density-function is held to cover the case. We would then be able to know before opening that we would prefer to swap after opening, otherwise we would not know that.)

We do not regenerate paradox. (This holds even in the case with infinite utilities which I am not really considering: Prior to receiving the information that we would receive such information in the case of one of the envelopes, we would be indifferent between the envelopes. We would in the light of the information received form a preference on the basis of our beliefs about the case. We would change our preference in case we were to be told that after all the other envelope would in fact be opened first. That is not paradoxical, even if it might be considered odd.)

Conclusion about case of opened envelope: Either we apply expected utility theory to the problem at hand or we do not. If we do not, there is after all

no paradox to be generated from expected utility-theory . If we do apply this theory, we do have stable solutions. Given a specific type of probability density-function, and an infinite case (infinite utilities), we would know beforehand that we would prefer to swap after opening. Still, we would either be indifferent, before receiving information, or hold a stable preference after receiving information.

C: Conclusion.

Paradox is lost, and not regained. Until an envelope among A or B is opened, we should remain indifferent between A and B. After an envelope has been opened, and we know the amount in it, and if we hold that we can apply expected utility-theory to the problem, we will find it rational to swap to the other envelope (B), in case we hold it sufficiently probable that B contains the larger amount. It is easy to calculate what is sufficient here: We must hold it more than 1/3 probable that B contains the larger envelope.⁴ We do not, in finite or standard cases, **know** before opening that we will swap. But even if we do know that we will swap, in some special non-standard cases with infinite utilities, that does not generate paradox. Expected utility theory is completely unaffected by the two-envelope paradox.

Literature

- Arntzenius, F. and D. McCarthy. 1997. "The two envelope paradox and infinite expectations." **Analysis** 57, 42-50.
- Broome, J. 1995 "The two-envelope paradox." **Analysis** 55, 211-16.
- Castell, P and D. Batens. 1994 "The Two-Envelope paradox: The Infinite case" **Analysis** 54, 46-59.
- Clark, M. and N. Shackel. The Two-Envelope Paradox. **Mind** 2000, pp. 415-442.

⁴If we call the amount in opened A "x". Then, in order to be indifferent between B and A:

$$B=(x/2)p \text{ plus } (2x)(1-p)=x=A$$

$$p=2/3$$

This means that to be indifferent between A and B after an opening of A we must hold it twice as likely that B contains $x/2$ as that it contains $2x$. Whether we do that depends on the density-function.

- Jackson, F., P. Menzies and G. Oppy. 1994 "The two-Envelope 'paradox'" **Analysis** 54, 46-59
- Lindley, D.V. **Making Decisions**. 2nd edition, London, John Wiley & Sons, 1985.
- McGrew, T., D. Shier and H. Silverstein. 1997 "The two-envelope paradox resolved", **Analysis** 57, 28-33.
- Nalebuff, B. "The other person's envelope is always greener." 1989 **Journal of Economic Perspectives** 3, 171-81.
- Rawling, P. 1997 "Perspectives on a Pair of Envelopes" **Theory and Decision**, 43, 253-277.
- Scott, A.D., and M. Scott. 1997. "What's in the two envelope paradox?" **Analysis** 57, 34-41.