

The Classical and Maximin Versions of the Two-Envelope Paradox

BRUCE LANGTRY
DEPARTMENT OF PHILOSOPHY
UNIVERSITY OF MELBOURNE
brucenl@unimelb.edu.au

Received by Greg Restall

Published August 2, 2004

<http://www.philosophy.unimelb.edu.au/ajl/2004>

© 2004 Bruce Langtry

Abstract: The *Two-Envelope Paradox* is classically presented as a problem in decision theory that turns on the use of probabilities in calculating expected utilities. I formulate a *Maximin Version* of the paradox, one that is decision-theoretic but omits considerations of probability. I investigate the source of the error in this new argument, and apply the insights thereby gained to the analysis of the classical version.

I TWO VERSIONS OF THE PARADOX

The *Two-Envelope Paradox* is a decision-theoretic problem that is widely taken to turn on considerations of probability.¹ Here is the Classical Version, as found in many recent contributions to the recent literature:

You were invited to select one of two sealed envelopes. You were informed that both contain money, one twice as much as the other. This is all that you know about them: you do not know, for example, anything about the process by which each envelope came to hold the amount of money that it holds. You have chosen one envelope—call it ‘Alpha.’ You are now offered two options: you may keep Alpha, or else swap it for the other envelope, Beta. What is it in your financial self-interest to do? Surely you have no reason to prefer one option to the other. But there is an argument, the Classical Argument, for the view that you should swap:

¹Raymond Smullyan, *Satan, Cantor and Infinity*, Oxford University Press 1992, pp. 189–192, offers what he calls ‘a version of the paradox’ whose conclusion and supporting arguments do not concern what the agent should do. Smullyan does not claim to provide materials that could be used in the construction of a paradox in decision theory, let alone explain how to perform the construction.

- (C-1) Let x be the amount of money in Alpha.
- (C-2) Then Beta contains either $2x$ or $x/2$.
- (C-3) These two possibilities are equally likely.
- (C-4) The expected value of keeping Alpha is $(0.5 \times x) + (0.5 \times x) = x$, and the expected value of swapping to Beta is $(0.5 \times 2x) + (0.5 \times x/2) = 1.25x$.
- (C-5) Therefore you should swap.

The problem is to locate the error in this argument. That there is an error somewhere is guaranteed not only by the implausibility of the conclusion, but also by the fact that similar reasoning for the opposite conclusion can be entertained, with the first step ‘Let y be the amount of money in Beta.’

All discussions of the Classical Argument known to me seek to identify the error as one bound up with probabilities. I will now set forth a new version of the paradox, the MAXIMIN Version. It is manifestly decision-theoretic, and closely related in structure to the Classical Version, except that it omits considerations of probability.

You were invited to select one of two sealed envelopes. You were informed that both contain money, one twice as much as the other, and this is all that you know about them. You have chosen one envelope, Alpha. You may now keep Alpha, or else swap it for the other envelope, Beta. You believe that you are not in a position to assign probabilities to the various possible outcomes of your choices—e. g., you cannot safely use the principle of indifference to conclude that there is a 0.5 probability that Alpha contains twice as much money as the Beta. Being risk-averse, you decide to employ the maximin rule.

Here are three parallel arguments for three mutually incompatible conclusions about what you should do. First, THE x -ARGUMENT: Let x be the unseen amount of money in Alpha. Then the amount of money in Beta is either $2x$ or $x/2$. We draw up a payoff table as follows:

	Beta has $2x$	Beta has $x/2$
SWAP	$2x$	$x/2$
KEEP	x	x

The action with the best worst possible outcome is KEEP, so you should perform this action.

Second, THE y -ARGUMENT: Let y be the unseen amount of money in Beta. Then Alpha contains either $2y$ or $y/2$.

	Alpha has $2y$	Alpha has $y/2$
SWAP	y	y
KEEP	$2y$	$y/2$

The action with the best worst possible outcome is SWAP, so you should perform this action.

Third, THE z-ARGUMENT. Let z be the total amount of money that has been distributed between the envelopes, so that one of the envelopes contains $2z/3$ and the other contains $z/3$.

	Alpha has $2z/3$	Alpha has $z/3$
SWAP	$z/3$	$2z/3$
KEEP	$2z/3$	$z/3$

The worst possible outcome of SWAP is identical with the worst possible outcome of KEEP, so you should be indifferent between the actions.

If we describe the Classical and Maximin Versions as two versions of *the* Two-Envelope Paradox, rather than two different paradoxes, we suggest more than that they have some structural similarities: we suggest that each has at least one resolution that overlaps in vital and especially instructive ways with at least one resolution of the other. This is indeed the conjecture that I want to test in what follows. If it is correct, then there is some error in the Classical Argument other than probabilistic ones.²

In Section 2 of this paper I investigate the Maximin Version, seeking to explain both why the x -argument and the y -argument are fallacious, and why the z -argument reaches what is intuitively the correct conclusion. In Section 4 I use some insights of Section 2 to resolve the Classical Version. Sections 3 and 5 concern side-issues.

2 AMBIGUITY AND ERROR IN MAXIMIN VERSION REASONING

Consider the first payoff table, presented above in connection with the x -argument. The table is supposed to be listing, for each of the two relevant actions, all the epistemically possible financial outcomes—that is, all the physically possible payoffs compatible with the evidence that there are two sealed envelopes, one containing twice as much money as the other, and that you initially chose one but now have the option of swapping it for the other. The payoff table does not in fact do this.

To establish this, let us ponder how we are to interpret the symbol ‘ x ’ that appears in various places in the payoff table. It was stipulated that x is the

²This conclusion is not tied especially closely to the maximin rule. Someone who believes that calculations involving probabilities are here unsafe might argue: After I have made my choice and the envelopes have been opened, I will be glad or disappointed in proportion to the difference between the dollar amount I in fact obtained and the dollar amount I would have obtained if I had chosen otherwise. If I were to swap to Beta, and it turned out that Alpha contained more money, then my disappointment would be measured by $x/2$, whereas if I were to keep Alpha and it turned out that Beta contained more money, then my disappointment would be measured by x . My policy is to act in such a way as to minimise my possible disappointment. Therefore I should swap envelopes.

unseen amount of money in Alpha. But this stipulation itself requires interpretation.

Consider the assertion ‘The Prime Minister used to be Keating but now is Howard.’ The phrase ‘the Prime Minister’ is not being used to pick out an individual, with the result that the sentence says of this individual that he used to be identical with Keating but is now identical with Howard. Rather the phrase ‘the Prime Minister’ is being used to state a condition, and what is asserted is that this condition was once fulfilled by Keating but is now fulfilled by Howard.

Suppose that the singular term ‘Jack the Ripper’ was originally introduced by someone’s saying, ‘Let us call the person who murdered these women “Jack the Ripper”.’ Is it epistemically certain for us, given that these women were murdered by one person, that Jack the Ripper is a murderer? It is, if we are using the schema ‘Jack the Ripper is #’ as an abbreviation of the schema ‘The person who murdered these women is #.’ In contrast, the sentence ‘None of Jack the Ripper’s friends know that he is a murderer’ should be interpreted not as ascribing to the friends ignorance of the meaning of the name ‘Jack the Ripper,’ but rather as saying ‘Take whichever person murdered these women: it is true *of him* that none of his friends know that he is a murderer.’ And since you and I do not know the identity of Jack the Ripper, there is no-one of whom we can truly say ‘It is epistemically certain that he murdered these women.’ If we are using the expression ‘Jack the Ripper’ to pick out a specific (but unknown) individual, then the sentence ‘It is epistemically certain for us, given that these women were murdered by one person, that Jack the Ripper is a murderer,’ expresses a false *de re* modal proposition. Indeed, the following expresses a *de re* modal truth: ‘Consider Jack the Ripper’s father: it is epistemically possible that *he* is the real murderer of these women.’

So what about the definite description ‘The amount of money in Alpha’ and the symbol ‘ x ’? Well, the most obvious way to think about the payoff table is this:

Consider this specific amount of money, x . It is now safely in Alpha, and will be neither increased nor decreased by your decision. If you KEEP, you will get this amount of money. If you SWAP, you will get either double or half this amount of money. It is clear that KEEP has the best worst possible outcome.

Plainly, in this line of thought ‘ x ’ is being used *de re*: it is intended to pick out a specific sum—how much is at present unknown but it is either \$10 or ... Given this understanding of the role of ‘ x ,’ there arises a decisive objection to the argument based on the payoff table displayed above. It is epistemically possible (since physically possible, and consistent with the evidence) that Alpha contains a sum of money other than the foregoing specific amount. For example, the following expresses a *de re* modal truth: ‘Consider x minus one cent: it is epistemically possible that *this amount* is in Alpha.’

Therefore, assuming for the sake of argument that we can properly draw up a payoff table expressed in terms of ‘ x ,’ then the table should list, as pos-

sible outcomes of your performing KEEP, your obtaining x minus one cent, x plus one cent, x minus 2 cents, x plus 2 cents, and so on. For each of these, the payoff table should list two corresponding possible outcomes of your performing SWAP, namely your obtaining twice or half the specified outcome of your performing KEEP. (This does not entail that the table should incorporate the assumption that there are twice as many possible outcomes of SWAP as of KEEP. That assumption would reflect double counting,—e. g., neglecting the fact that \$2 is both twice \$1 and half of \$4; and of course such an assumption would be inconsistent with our intuitive conviction that, however this or that table might be drawn up, KEEP and SWAP are symmetrical with respect to possible outcomes.)

We should now draw up a revised table, centred on x . How many columns does the complete table contain? Although for obvious reasons the amount of money in Alpha must be less than, say, $\$10^{100}$, there is no maximum amount of money that Alpha might contain. So the complete table contains indefinitely many columns. Since we are concerned with whether there is a best worst possible outcome, we can concentrate on the left hand side, which lists possible outcomes of KEEP that are below x . Here is a fragment:

	Alpha: $\$x - 0.1$ Beta: $\$(x - 0.1)/2$	Alpha: $\$x - 0.1$ Beta: $\$2x - 0.2$	Alpha: $\$x$ Beta: $\$x/2$	Alpha: $\$x$ Beta: $\$2x$
SWAP	$\$(x - 0.1)/2$	$\$2x - 0.2$	$\$x/2$	$\$2x$
KEEP	$\$x - 0.1$	$\$x - 0.1$	$\$x$	$\$x$

Because the full table has indefinitely many payoff columns to the left of those I have shown, it is plain that the full table does not register a worst possible outcome either of SWAP or KEEP. Indeed, there is no best worst possible outcome specifiable in terms of ‘ x .’ Hence the argument for KEEP fails.

It might be said that there is a worst possible outcome of KEEP, namely that you get 1 cent. This remark has a number of things wrong with it, the most important of which is that it does nothing to rehabilitate the argument that we have been considering for KEEP, since obviously if we count *Your getting 1 cent* as a possible outcome of KEEP then we will have to count it as a possible outcome of SWAP too.

For the last few paragraphs I have been working with the most obvious interpretation of ‘ x ,’ namely that it picks out that specific sum of money that is in fact contained in Alpha, and that the payoff table records *de re* modal truths. But we should look briefly at the alternative interpretation, that ‘ x ’ instead specifies a condition that various different sums of money might fulfil.

	Beta has $2x$	Beta has $x/2$
SWAP	$2x$	$x/2$
KEEP	x	x

Under the alternative reading, the presence of ‘ x ’ in both cells in the ‘KEEP’ row fails to embody the truth that Alpha contains some definite amount of money that you will get if you perform KEEP. Instead, ‘ x ’ merely indicates that if you perform KEEP you will get some sum that will fulfil the condition *is the amount of money in Alpha*. Similarly with the expressions ‘ $2x$ ’ and ‘ $x/2$ ’. Under the alternative interpretation of the table, we are not entitled to assume that the expression ‘ $2x$ ’ in the first row refers to a greater dollar amount than the expression ‘ $x/2$ ’ in the first row. Therefore we cannot read off from the table the statement that KEEP has the best worst possible outcome. Plainly this interpretation can fairly be labelled ‘deviant,’ in that the table so interpreted does not reflect or facilitate the intended argument for performing KEEP.

We now have a diagnosis of where people’s reasoning goes wrong in the x -ARGUMENT and the y -ARGUMENT. Ambiguity in the symbols ‘ x ’ and ‘ y ’ obscures problems with the construction of the payoff tables, and so causes us to make fallacious inferences from them.

Let us now briefly consider the z -ARGUMENT, the one that reaches what is intuitively the correct conclusion. The role of ‘ z ,’ I shall assume, is to pick out that specific sum of money that is in fact distributed between the envelopes. It is physically possible, and consistent with the evidence, that a sum of money other than z was distributed between the envelopes. The payoff table given in connection with the z -ARGUMENT is therefore incomplete. But fact does not undermine the crucial lemma of the z -ARGUMENT, that there is no best worst possible outcome. Indeed, the properly revised table will enable the formulation of a modified z -ARGUMENT, along obvious lines, which supports that conclusion.

3 TWO VARIATIONS ON THE x -ARGUMENT

Before we leave the Maximin Version, we should look at two interesting variations on the theme. The first is a line of reasoning designed to avoid the use of singular terms to denote outcomes:

Your decision will not cause money to be transferred into or out of either of the envelopes. If you perform KEEP then you will be wealthier than you were before you picked up Alpha. If you perform SWAP, then you will also be wealthier than you were before you picked up Alpha: you will be either more wealthy or less wealthy than you would be if you performed KEEP, but you do not know which. If you perform SWAP then you might end up less wealthy than you would be if you performed KEEP, but if you perform KEEP then you cannot end up less wealthy than you would be if you performed KEEP. Therefore, of the two actions KEEP and SWAP, KEEP has the best worst possible outcome.

The conclusion is implausible, and a parallel argument for the opposite conclusion can be formulated simply by interchanging the words ‘KEEP’ and ‘SWAP.’ So there must be a fallacy in the inference from the premises to the conclusion. True: if you perform SWAP then you *might* end up less wealthy than

you *would* be if you performed KEEP. But although this entitles us to assert that SWAP has a possible outcome that is worse than at least one possible outcome of KEEP, it does not entitle us to assert that SWAP has a possible outcome that is worse than *all* possible outcomes of KEEP.

But surely, it might be objected, there is no possible outcome of KEEP in which you end up less wealthy than you will be if you perform KEEP, whereas there is a possible outcome of SWAP in which you end up less wealthy than you would be if you performed KEEP. Well it cannot be the case that both you perform KEEP and you have the description *ends up less wealthy than you will be if you perform KEEP* true of you. Nevertheless there is a possible outcome of your performing KEEP in which you end up less wealthy than you will in fact be if you perform KEEP. (Suppose that Alice has mistakenly believed that you are a basketball player. On meeting you for the first time she says, 'I assumed that you were taller than you in fact are.' One interpretation of this sentence ascribes to her an assumption that is logically impossible; there is a more charitable interpretation that does not.)

The second variation on the X-ARGUMENT urges you to reason in this way:

- (X-1) Suppose, contrary to fact, that you opened Alpha and discovered how much money it contained. Then *whatever this amount was*, you would identify KEEP as the action that had the best worst possible outcome. (Suppose, for example, that you found that there was in fact \$10 in Alpha. Then you would identify your getting \$10 as the worst possible outcome of KEEP, and would identify your getting \$5 and your getting \$20 as the two possible outcomes of SWAP. You would therefore identify KEEP as the action that had the best worst possible outcome. And so on.)
- (X-2) Therefore in the actual situation, in which you do not know how much money Alpha contains, you should perform KEEP.

The invalidity of the argument is obvious from the fact that a parallel argument for the opposite conclusion can be formulated simply by replacing 'Alpha' with 'Beta' and 'SWAP' with 'KEEP'. But where, exactly, does the argument go wrong? Agreed: after you have discovered how much money Alpha contains, SWAP has a worse epistemically possible outcome than any epistemically possible outcome of KEEP. Moreover this truth is independent of the specific amount that Alpha in fact contains.³ But it is dependent on the fact that *you have discovered* how much money Alpha contains: the information about Alpha, whatever the dollar amount involved, restricts the range of epistemically possible outcomes of KEEP more narrowly than it restricts the range of epistemically possible outcomes of SWAP. If you had instead discovered how much money Beta contains, the resulting ranges of epistemically possible outcomes would have been different. Therefore you should not slide from the fact that it

³Let us not press the point that there is an exception: the case in which you discover that Alpha contains one cent.

does not matter how much money Alpha contains to the thesis it does not matter whether you know how much money Alpha contains. And you should not conclude that in the actual case, in which you do not know how much money Alpha contains, SWAP has a worse possible outcome than any possible outcome of KEEP and so you should perform KEEP.⁴

4 AMBIGUITY AND ERROR IN THE CLASSICAL ARGUMENT.

Here again is the Classical Argument to the conclusion that you should swap envelopes.

- (C-1) Let x be the amount of money in Alpha.
- (C-2) Then Beta contains either $2x$ or $x/2$.
- (C-3) These two possibilities are equally likely.
- (C-4) The expected value of KEEP is $(0.5 \times x) + (0.5 \times x) = x$, and the expected value of SWAP is $(0.5 \times 2x) + (0.5 \times x/2) = 1.25x$.
- (C-5) Therefore you should swap.

There is an ambiguity in the symbol ‘ x ’ that appears in the argument 1.–5. Either ‘ x ’ is being used *de re*, to pick out a specific sum whose size we have no rational way of estimating, or else ‘ x ’ is being used to identify a condition that various specific sums can fulfil.

Suppose that the former is the case. Then (C-4) is false. The expected value of KEEP cannot be the definite but unknown quantity denoted by ‘ x ’, and the expected value of SWAP cannot be the definite but unknown quantity denoted by ‘ $1.25x$ ’, since (given only the information on hand) KEEP and SWAP *have* no expected values.

Here are four ways of bringing out the point. Firstly, consider the argument:

- (L-1) This table is roughly circular.
- (L-2) Let ‘ L ’ denote the average length in centimetres of the table’s diameters
- (L-3) Therefore our best estimate of the circumference of the table is πL .

⁴Suppose that we were to write in to the specification of the problem the existence of an actual spectator who knows how much money Alpha contains but does not know how much money Beta contains, and who has no causal interaction with you, the agent. Whatever amount of money is in Alpha, SWAP has, relative to the epistemic state of this spectator, a worse possible outcome than any possible outcome of KEEP. You, we might suppose, know this. Does it follow that SWAP has, relative to *your* epistemic state, a worse possible outcome than any possible outcome of KEEP? Surely not.

If 'L' is being used to refer to the specific number, but we have no idea what number that is, and no other relevant information, then we do not have a best estimate of the circumference of the table. This is consistent with a truth implied by the first premise, that when we do discover the length of the diameter and the circumference, we will find that the latter is roughly π times the former. Furthermore, suppose that there is a second table, also roughly circular, whose average diameter is twice the length of the first table's. We can sensibly assert that our best estimate of the ratio of the two circumferences is 2, but not that the ratio of our best estimates of the circumferences is 2.

Here is the second consideration. Suppose you introduce the expression 'Jack the Ripper' by declaring that you will be using it *de re*, to pick out the murderer of these women. Nevertheless you do not know who Jack the Ripper was, and you do not know of Jack the Ripper that he murdered anyone. Similarly, if you the agent are using the expression 'x' *de re*, to refer to the specific sum of money that is in fact in Alpha, it remains the case that you do not know how much money is in Alpha, and you do not know of the specific sum of money that happens to be in Alpha that it is equal to the expected value of *KEEP*. Hence you do not know that the expected value of *KEEP* is equal to x. Hence you are not in a position to accept (C-4) But if you, an ideally rational agent, are not in a position to accept (C-4), then how can (C-4) be true?⁵

Thirdly, if 'x' is tied to the specific amount of money that happens to be in Alpha, then (C-4) entails truths like the following:

- If $x = \$10$ then the expected value of *KEEP* is \$10, and the expected value of *SWAP* is \$12.50.
- If $x = \$11$ then the expected value of *KEEP* is \$11, and the expected value of *SWAP* is \$13.75.

But (C-4), so understood, *cannot* be true. Suppose that while prospective buyers of tickets in a certain lottery know the probability that any given ticket will win, they have no clue to how much the prize is. If it turns out that the prize was \$10,000, this does not mean that the expected value of buying a ticket was 100 times what it would have been if the prize money had been \$100. The contrary supposition largely detaches the concept of expected value from its point, its established roles in evaluation and choice. Similarly, the expected value of *KEEP* cannot depend on how much money Alpha in fact contains (unknown to the agent). So if 'x' is being used *de re*, to pick out this specific amount of money, the expected value of *KEEP* cannot be identified with x. Indeed, given the information on hand, it cannot properly be identified in any other way, either. Now the concept of expected value is such that if an action possesses an expected value for an agent at a time then that expected value can be identified by

⁵An extended presentation and defence of this argument obviously would require delving deeply into the theory of *de re* belief.

the (ideally rational) agent at that time, given the information then available to the agent. Therefore *KEEP* lacks an expected value.

Fourthly, if 'x' is being used to pick out the specific sum that is in fact in Alpha, then *KEEP* and *SWAP* have no expected values because some or all of the possible outcomes do not possess determinate probabilities. To calculate the expected value of an action you consider each possible outcome and multiply its value (here measured by units of currency) by the probability of obtaining it given that you perform the action in question; you then add together the products thereby obtained. Now as I argued in Section 2 of this paper, if 'x' is being used to pick out the specific sum that is in fact in Alpha, it is epistemically possible that Alpha contains a sum of money other than the foregoing specific amount. For example, the following expresses a *de re* modal truth: 'Consider x plus \$1000: it is epistemically possible that this amount is in Alpha.' Although there is no doubt that Alpha contains a finite amount of money, there is no maximum amount of money that Alpha might contain. So there are indefinitely many possible outcomes of performing *KEEP* and *SWAP*. Hence the probability that you will obtain a given outcome if you perform, say, *SWAP* is indefinite. Specifically, the probability that you will get will get $2x$ if you perform *SWAP* is not 0.5 but indefinite. Hence no calculation of expected value is possible.⁶

Let us turn now to the other candidate interpretation of 'x,' according to which 'x' identifies a condition that various specific sums might fulfil—viz. the condition *is the amount of money in Alpha*. Once again, *KEEP* and *SWAP* do not have expected values. There are a great many relevant possible outcomes of *KEEP*—e. g., that you get \$1, that you get \$1000, ... All we know about their values is that, whatever they are, they fulfil the condition *is measured by whatever quantity of money is in Alpha*. All we know about the values of the great many relevant possible outcomes of *SWAP* is that some of them fulfil the condition *is measured by twice whatever quantity of money is in Alpha* while others of them fulfil the condition *is measured by half of whatever quantity of money is in Alpha*. We also know that these two conditions are equally likely to be fulfilled. But these facts do not enable us to multiply relevant values by relevant probabilities and thereby confer a coherent meaning on the assertion, made in the context specified above, 'The expected value of *KEEP* is equal to whatever quantity of money is in Alpha.' The matter turns on the 'logic' of *expected value*. (There is no problem about the meaning of the statement 'The payoff of *KEEP* will be equal to whatever quantity measures the money in Alpha.')

It is perhaps worth making a brief comment on the idea that 'x' should be regarded as a variable. When we say that the area of a circle is πD , where D is the diameter, the letter 'D' is often understood as a variable ranging over

⁶Nevertheless, leaving aside general doubts about the principle of indifference, we can sensibly say that if the agent knows that Alpha contains \$10, and that Beta contains either \$20 or \$5, and nothing more is known, then the epistemic probability of getting \$20 if you perform *SWAP* is definite—in fact, is 0.5. And so on.

lengths, and in that case what we mean by what we say is this: There are many possible numerical values that the surface area of a circular table top can have, and (given an appropriate choice of units of measurement) they all fulfil the condition *is equal to π multiplied by whatever number fulfils the condition measures the length of the table's diameter*. Similarly if we treat 'x' as a variable ranging over possible dollar amounts in Alpha when we say 'The value of KEEP is x' what we are taking ourselves to be saying is that the value of KEEP fulfils the condition *is measured by whatever quantity of money is in Alpha*. This being so, the suggestion that 'x' should be regarded as a variable amounts to basically a variant of the idea, considered in the previous paragraph, that 'x' identifies a condition that various specific sums might fulfil.⁷

The argument so far presented in this Section was constructed before I had read Terence Horgan's fine paper on the Classical Version.⁸ Our accounts are similar in that both focus on ambiguities in expressions such as 'x' and 'the amount of money in Alpha.' Horgan distinguishes between *standard* and *non-standard* expected utility, in terms of whether or not the calculation involves the specifications of possible outcomes, utilities and probabilities made using *canonical* or *non-canonical* referring expressions. (A referring expression is canonical for an agent only if its referent is epistemically determinate for the agent, given the agent's total available evidence.) Horgan distinguishes several kinds of non-standard expected utility—x-based, y-based, and z-based—for the actions KEEP and SWAP; the differences are a matter of *what* non-canonical referring expression is used to specify the states of the world and the payoffs.⁹ Horgan regards the classical argument up to and including step (C-4) as sound,

⁷Someone might try to evade the foregoing arguments by employing the idea of conditional expected probability, where to calculate the expected value of an action conditional upon the obtaining of a state of affairs S, you consider each possible outcome and multiply its value (measured by units of currency) by the probability of obtaining it given that both you perform the action in question and S obtains; you then add together the products thereby obtained. Accordingly, let us try rewriting step (C-4) of the Classical Argument as (C-4') *The conditional expected value of KEEP—i. e. the expected value relative to the assumption that Alpha contains x—is $(0.5 \times x) + (0.5 \times x) = x$, and the conditional expected value of SWAP is $(0.5 \times 2x) + (0.5 \times x/2) = 1.25x$.*

Surely if 'x' is being used to pick out the specific sum of money that happens to be in Alpha, then (C-4') entails that if $x = \$10$ then $EV(\text{SWAP} \mid \text{Alpha contains } x)$ is \$10, and that if $x = \$11$ then $EV(\text{SWAP} \mid \text{Alpha contains } x)$ is \$13.75, and so on. For reasons similar to those given in the main text, the function EV, when put to use in (C-4'), does not amount to a concept of conditional *expected value*, and even if it did then KEEP and SWAP would have no expected values relative to 'Alpha contains x' because some or all of the possible outcomes would not possess determinate probabilities.

⁸'The Two-Envelope Paradox, Nonstandard Expected Utility, and the Intensionality of Probability,' *Noûs* 34 (2000), 578–603.

⁹Horgan's pp. 584f, 592 make it clear that x-based nonstandard expected utility is calculated by letting 'x' go proxy for the expression 'the actual amount of money in Alpha,' where, as the word 'actual' is a rigidifying operator and so 'x' is a rigid designator. So talk of x-based non-standard expected utility involves treating 'x' as picking out the specific dollar amount at Alpha in fact contains, rather than as identifying a condition that various specific dollar amounts can fulfil.

provided only that we realise that the 'expected utility' spoken of in (c-4) is x -based and non-standard. Horgan locates the error in the classical argument as lying in the inference from (c-4) to (c-5). For there is no correct rule of rational decision theory requiring the agent to maximise non-standard expected utility. What a rational agent should do is rather maximise standard expected utility.

I, however, deny the existence of x -based non-standard expected utility, at least if it is supposed to be *expected utility identified in a non-standard way*, or a *variety of expected utility*, or a *notion of expected utility*.¹⁰

Teachers might distinguish several kinds of overall grade, based on different ways of calculating a student's overall grade from the numerical marks given for many pieces of work. They might, for example, take the average, or alternatively ignore the lowest mark and take the average of the others. What makes the outcomes of these calculations deserve the description 'overall grade'? The ways the teachers and other people are in principle prepared to put the outcomes to use: forming an impression of the student's abilities, deciding which of several competing students should be awarded some prize, and so on. Even if the teachers always use one method, they regard it as intelligible that one might use the others for the relevant purposes, and regard the choice of methods as requiring serious deliberation. But what makes it the case that the results of Horgan's calculations deserve the description 'species of *expected utility*? The x -based and y -based calculations do not issue in any output, specified in such terms as ' x ' and ' $1.25y$,' that you can regard as offering serious help in ranking available actions, or events in terms of the degree to which you should welcome them. (Of course if you are in a position to specify the output in terms of numerals then you know either how much money is in Alpha or how much is in Beta, and you face a decision problem quite different from that posed in the Classical Version.) The point here goes beyond Horgan's remark that what the basic normative principle of decision theory prescribes is choosing the action with maximum *standard* expected utility.¹¹ Suppose, for example, that a coin is to be tossed and if heads result then you will get the contents of Alpha, whereas if tails result you will get the contents of Beta; a calculation of ' x -based non-standard expected utility' for the events *heads* and *tails* will not tell you which event to *hope for*.

Furthermore, Horgan commits himself to the statements that if Alpha, unknown to the agent, contains \$16 dollars then the x -based non-standard expected utility of KEEP is 16, while if Alpha, unknown to the agent, contains \$32 then the x -based non-standard expected utility of KEEP is 32.¹²

¹⁰We have extended the concept of number so that complex numbers can be regarded as a kind of number, but it remains the case that although the expressions 'plastic duck' and 'Christian Scientist' have legitimate uses, plastic ducks are not a variety of duck, and the notion of a Christian Scientist is not a notion of scientist.

¹¹ibid., p. 593.

¹²ibid., p. 592.

This is rather like saying that if you do not know anything about the amount of a lottery prize, and the prize is in fact \$10,000, and if you are an ideally rational agent, then there is some kind of expected utility such that the action of buying a ticket is in fact twice what it would have been had the prize been \$5,000. No. As I have already pointed out, such an assertion largely detaches the concept of expected value from its point, its established roles in evaluation and choice. If you have no idea how much the prize is, then the action of buying a ticket has no expected utility of any kind.

The fact that in the two-envelopes case Horgan's numbers 16 and 32 are the results of calculations bearing some structural similarities to calculations of standard expected utility does not settle whether the x -based non-standard utility being talked about here is a kind of expected utility. It is not.

5 EXPECTED UTILITY AFTER OPENING ONE ENVELOPE

It is worth pausing to consider the following variation on the Classical Argument:

- (O-1) Suppose, contrary to fact, that you opened Alpha and discovered how much money it contained. Then *whatever this amount was*, you would identify SWAP as the action that had the higher expected utility. (Suppose, for example, that you found that there was in fact \$10 in Alpha. Then KEEP would have an expected value of \$10 while SWAP had an expected value of \$12.50. And so on.¹³)
- (O-2) Therefore in the actual situation, in which you do not know how much money Alpha contains, you should perform SWAP.

Let us, for the sake of discussion, leave aside general doubts about the principle of indifference. After you have discovered how much money Alpha contains, SWAP has a higher expected value than KEEP. This truth is independent of the specific amount that Alpha in fact contains. But it is dependent on the fact that *you have discovered* how much money Alpha contains: for if you had not opened Alpha, and had instead been told how much money Beta contains, the expected values of SWAP and KEEP would have been different. Therefore you should not slide from the truth that it does not matter how much money Alpha contains to the thesis it does not matter whether you know how much money Alpha contains. And you should not conclude that in the actual situation, in which you do not know how much money Alpha contains, SWAP has a higher expected value than KEEP and so you should swap envelopes.

¹³The idea behind the words 'and so on' might be expressed in the schema: *If the agent knows that Alpha contains \$x, and that Beta contains either \$2x or \$x/2, and that is the agent's total relevant information, then the expected value of SWAP is \$1.25x.* (The 'if ... then' structure of the schema hides the fact that the expected value function incorporates epistemic probabilities that are relative to the agent's information about Alpha's contents.)

6 CONCLUSION

In both the Classical and Maximin versions of the paradox, you should be indifferent between KEEP and SWAP because you do not have a good reason for preferring one action to the other. (I have identified fallacies involved with what at first appeared to be good reasons.) More specifically, as far as your beliefs, desires, and values are concerned, you are symmetrically situated with regard to the envelopes, except for the fact that one of them is already in your hand while other is not, and *this* fact does not provide a good reason for preferring one action to the other.¹⁴

¹⁴Unpublished work by Graham Priest and Greg Restall first drew my attention to the Two-Envelope Paradox. I thank them, and also James Dreier, Alan Hajék, Allen Hazen, Lloyd Humberstone, Graham Oppy and two anonymous referees, for stimulating and helpful discussion.

The *Australasian Journal of Logic* (ISSN 1448-5052) disseminates articles that significantly advance the study of logic, in its mathematical, philosophical or computational guises. The scope of the journal includes all areas of logic, both pure and applied to topics in philosophy, mathematics, computation, linguistics and the other sciences.

Articles appearing in the journal have been carefully and critically refereed under the responsibility of members of the Editorial Board. Only papers judged to be both significant and excellent are accepted for publication.

The journal is freely available at the journal website at

<http://www.philosophy.unimelb.edu.au/ajl/>.

All issues of the journal are archived electronically at the journal website.

SUBSCRIPTIONS Individuals may subscribe to the journal by sending an email, including a full name, an institutional affiliation and an email address to the managing editor at ajl-editors@unimelb.edu.au. Subscribers will receive email abstracts of accepted papers to an address of their choice. For institutional subscription, please email the managing editor at ajl-editors@unimelb.edu.au.

Complete published papers may be downloaded at the journal's website at <http://www.philosophy.unimelb.edu.au/ajl/>. The journal currently publishes in pdf format.

SUBMISSION The journal accepts submissions of papers electronically. To submit an article for publication, send the \LaTeX source of a submission to a member of the editorial board. For a current list of the editorial board, consult the website.

The copyright of each article remains with the author or authors of that article.