The Paradox of the Two Envelopes

Chloé de Canson

1. Introduction

There are two envelopes in front of you. You are told that they each contain some amount of money, and that one contains an amount twice as large as the other, but you do not know which. You pick an envelope, and you are given the opportunity to change. Call the chosen envelope M, for "mine", and the non-chosen one O, for "other". Let us call this the initial situation. Now, let x be the amount in M. There are two possibilities: either O contains $x/2$, or it contains $2x$. They are epistemically equally likely. You represent this decision problem in the following matrix **1**.

	O contains $x/2$	O contains $2x$
Stick	x	x
Change	$x/2$	$2x$

Matrix 1

The expected utilities of sticking and changing are, respectively:

$$E[S] = x/2 + x/2 = x$$

$$E[C] = 1/2(x/2) + 1/2(2x) = 5x/4$$

So, it appears that it is rational for you to change to the other envelope. Three arguments for the absurdity of this conclusion are often sketched in the literature. Firstly, given that both envelopes are assumed to be indistinguishable, one cannot be rationally preferred to the other. Secondly, the same argument can be given for changing back, after having changed once, and so *ad infinitum*. Thirdly, an analogous argument can be given for rationality of sticking, starting with y as the amount in O. So, something must have gone wrong in the above argument. But what? This is the *two-envelope paradox*.

In this paper, my main argument is that the above presentation of the two-envelope paradox switches between two distinct decision problems, such that only one is adequately described by matrix **1**, and only the other makes it absurd to change envelopes. The structure of the paper is as follows. I begin by examining how to spell out precisely the conditions of the initial situation (§2). I then show that matrix **1** does not satisfy one of these conditions (§3), but that a number of decision matrices which yield equal expected utilities for sticking and changing do (§4). This gives me enough background to diagnose the source of the paradox: a conflation between two very similar yet distinct decision problems, one of which is adequately represented by matrix **1** and for which it is rational to change, the other one, not adequately represented by **1**, for which it is not rational to change (§5). I end this paper by considering several alternative resolutions of the paradox present in the literature, and building on some of the results established in the paper to show that they are erroneous: some concerning boundedness and infinity (§6), and others regarding designation and intensionality (§7).

Before beginning, I want to outline the basic formalism I will be employing throughout the paper. Every decision problem will be represented in the following way. First, I will specify a finite set of possible states $\Omega = \{w_1, w_2, ..., W_n\}$.[1] These are the states that the agent considers possible. A probability function p, representing a rational agent's epistemic credences in each of these states, will be defined on an algebra $\mathcal{A}$ over this set. Actions will be represented as random variables, which are functions $X : \Omega \to \mathbb{R}$ from the set of possible states of the world to the real numbers. All the decision problems I will represent will have two possible actions: sticking (S) and changing (C). Accordingly, there will be a random variable S representing the action of sticking, which assigns to every possible state the amount of money that the agent gets if she decides to stick with envelope M. Similarly, there will be a random variable C which assigns to every state the amount that the agent gets if she chooses to change to envelope O. So, it will be possible to represent every decision problem in a matrix of the form:

	w_1	w_2	$\cdots$	w_n
S	x_{S1}	x_{S2}	$\cdots$	x_{Sn}
C	x_{C1}	x_{C2}	$\cdots$	x_{Cn}

Sample Matrix

All the sets of states I will consider will be finite, and all the random variables that I will mention will be discrete, so their expectation values (interpreted as the expected utilities of performing the action they represent) will be calculated as $E[X] = \sum_i (w_i p_i)$. Lastly, I will sometimes mention

[1] There is a finite amount of money in the world, so there are a finite number of possible states with different amounts in M and O. See §6.

a set P which corresponds to the set of values the agent deems possible in either envelope. This should give us enough background to tackle the two-envelope paradox.

2. The Initial Situation

In the two envelope situation, we are initially told three things: (1) that each envelope contains some amount of money, (2) that one contains an amount twice as large as the other, and (3) that it is epistemically equally likely for each of the envelopes to contain more than the other. In this section, I try to spell out these assumptions precisely, and I offer some comments on how they are related to one another.

Let's start with (1), the claim that both envelopes contain some amount of money. Even though this is not made explicit in the formulation of the problem, I think it is a crucial premise that the agent lacks some information about the amount in each envelope. The first, quite obvious, type of information the agent lacks, is how much money there is in each envelope. We will discuss this later. For now, simply suppose that there is some set $P \subseteq \mathbb{R}$, which corresponds to all the amounts that the agent considers to be possibly contained in one of the envelopes. The second, more interesting, type of ignorance that the agent has, and this is really crucial, is that for any amount of money x, her evidence does not let her say whether that amount is in one envelope or in the other. In other words, if the agent considers it possible for M to contain some amount x, then she must also consider it possible for O to contain x. If she thinks it possible to get x by changing, she must also think it possible to get x by sticking, and vice versa. We spell this out formally as follows.[2]

$$\begin{gathered}\text{There exists a } w \in \Omega \text{ such that } S(w) = x \text{ iff there exists a } w' \in \Omega \\ \text{such that } C(w') = x \text{ for all } x \in P\end{gathered} \tag{1}$$

Now, on to (2). It says that, in every possible state of the world, there is twice as much in one envelope than in another. In other words, in every state, the maximum that you get by either sticking or changing is equal to twice the minimum that you get by either sticking or changing. So, formally:

$$\begin{gathered}\text{There exists a constant random variable } R : \Omega \to \mathbb{R} \text{ such that} \\ R(w) = \frac{min[S(w), C(w)]}{max[S(w), C(w)]} = 1/2 \text{ for all } w \in \Omega\end{gathered} \tag{2}$$

[2] Another way to formulate this is $S^{-1}(B) \in \mathcal{A}$ iff $C^{-1}(B) \in \mathcal{A}$ for all $B \subseteq P$.

Finally, we get to (3), which says that the agent considers it equally likely for each envelope to contain more than the other. So, if an agent picks an envelope at random, it is equally likely that she gets more by sticking, than she gets by changing. This can be expressed formally as:

$$p[S > C] = p[C > S] \tag{3}$$

This concludes our analysis of the initial situation. There are three important assumptions made: (1) that every amount possibly in M is also possibly in O, (2) that one envelope contains twice as much as the other in each possible state of the world, and (3) that it is equally likely for each envelope to contain more than the other.

3. Diagnosing the Problem

Recall that the two-envelope paradox arises when the initial situation is modelled as matrix **1**, where the expected utility of changing is greater than the expected utility of sticking. I will call such a matrix a ***paradoxical matrix***. Solving the paradox should then consist in explaining what went wrong; why matrix **1** is not suited to represent the problem at hand. There are a number of proposals in the literature, some of which I discuss later in the paper, and many of which propose quite complicated diagnoses. By contrast, I contend that the solution to the paradox is actually very simple: matrix **1** violates one of the assumptions of the inital situation.

There is an initial problem with matrix **1**; namely, that it is not clear what the possible states are. What is contained in Ω? Which options does the agent believe to be possible? Suppose, first, that the agent considers two equiprobable states. In both states, M contains x, but O contains $x/2$ in one and $2x$ in the other. In other words, take the matrix at face value. Then, clearly, it is consistent with constraints (2) and (3). Indeed, in each possible state of the world w, the amount the agent gets for sticking and the amounts she gets for changing are such that once is twice as large as the other. So, the model validates (2). Furthermore, since the two states are equiprobable, and that M contains more than O in one, and less in the other, (3) also comes out true. However, (1) is violated: there is a state w in which $S(w) = x$, but there is no state w' in which $C(w') = x$. Furthermore, there exist states w, w' for which $C(w) = x/2$ and $C(w') = 2x$, but no states w for which $S(w) = x/2$ or $S(w) = 2x$. So, if the agent believes that there are two possible states, those described by matrix **1**, then she will have misrepresented the initial situation. Matrix **1** cannot correctly represent both the initial situation and the states that the agent considers possible.

There are two ways to go from here. The first is to argue that constraint (1) is not a good constraint on representations of the initial situation. I discuss this in §4. The second is to claim that matrix **1** is not to be taken at face value: it is something like a summary. The agent does not actually consider that there are only two possible states, those in the decision matrix, but the states she actually considers are "conveniently" classified under these two categories. It is important to note that, unless she knows at least how many possible states there are, and what the values of changing and sticking are for each of these states, she will not be in a position to draw a decision matrix, and thus will not be able to calculate expected utility. So, let us charitably assume that the agent knows, somehow, that the amounts in the envelopes can vary between 1 and 8. Then, matrix **1** above can be seen as a coarse-graining of matrix **2** below: the agent actually considers four states, two in which M contains 2, and two in which it contains 4. She assigns credence $1/4$ to each of these states.[3]

	w_1	w_2	w_3	w_4
S	2	2	4	4
C	1	4	2	8

Matrix 2

Note that there is no possible state in which it contains 1 or 8, because, if matrix **2** is going to be a fine-graining of matrix **1**, it will have to be the case that, for every value of M, there is an epistemically possible state where O contains half that value, and another where it contains twice. This blocks the two extremal values of our interval.

Now, matrix **2** is distinct from matrix **1**, and the expectation values for sticking and swtiching remain paradoxical:

$$E[S] = 1/4(2 + 2 + 4 + 4) = 12$$

$$E[C] = 1/4(1 + 4 + 2 + 8) = 15$$

This will only be successful at restoring the paradox if matrix **2** respects constraint (1). But again, it does not. Counterexample: C takes value 1 at w_1, but there exists no state at which S takes value 1. Now, we see that this is a devastating argument against the suggestion that matrix **1** is a summary or coarse-graining of the agent's actual representation of the decision problem. Matrix

[3] In fact, this is not required by the setup. All that is required is that it is as likely for M to contain more than O, than it is for O to contain more than M. All that is needed to satisfy this is that $p(w_1) + p(w_3) = p(w_2) + p(w_4)$. But, assuming a uniform distribution over all the states is simpler in this case and nothing rests on it, so I will go along with that.

1 has as inbuilt that, for every possible value of M, it is possible for it to be twice the value of O, and half the value of M. But if the possible amounts present in the envelopes are bounded, then this assumption will not hold.[4] This leaves rejecting that (1) is an assumption about the initial situation as the only possible objection to my diagnosis. I turn to discussing this now.

4. The Plausibility of (1)

Here is a reminder of what (1) claims: if the agents thinks that some amount x can be in M, then she also thinks that it can be in O. Another: calling P the set of possible amounts in at least one of the two envelopes, the range of S and the range of C are identical, and they are both equal to P. It seems to me unclear how to convince the sceptical reader that (1) is indeed assumed. Short of a compelling positive argument, my strategy will be to give some reasons not to be suspicious of it. I will proceed in a twofold manner: first, I will present some non-paradoxical matrices and show that they all satisfy it (§4), and second, I will show that it is not a very strong requirement at all (§6).

The literature is full of matrices that are claimed to represent the initial situation, but do not yield the paradoxical inequality between the expectation values of S and C. I present three and show that they all satisfy (1). To be sure, this is not meant as a definitive argument for (1), far from it, but it is meant to appease lingering suspicions. Note, in passing, that all these also satisfy (2) and (3), in case you were also suspicious of those.

The first matrix is akin to one proposed by Horgan (2000), which he calls the urn case. Suppose the agent knows that the three values that may be in either envelope are 1, 2, 4 and 8. Then, she can model the situation in the following way:

	w_0	w_1	w_2	w_3	w_4	w_5
S	1	2	2	4	4	8
C	2	1	4	2	8	4

Matrix 3

Here, if the agent assigns the same credence, $1/6$, to each of the six states, the matrix clearly satisfies the three identified constraints. Furthermore, the expected utilities of sticking and changing are equal:

$$E[S] = 1/6(1 + 2 + 2 + 4 + 4 + 8) = 3.5$$

[4] This consideration about the boundedness of the possible values of x is what led Jackson, Menzies and Oppy (1995) to diagnose the problem differently than I have. I discuss this in §6.

$$E[C] = 1/6(2 + 1 + 4 + 2 + 8 + 4) = 3.5$$

You might notice that the non-paradoxical matrix **3** is identical to the paradoxial matrix **2**, but for the addition of states w_0 and w_5. Those two states were precisely the states whose inclusion into the matrix was forbidden by the constraint that matrix **2** be a coarse-graining of matrix **1**. Furthermore, the inclusion of these two extra states enabled matrix **3** to satisfy constraint (1). This reinforces my conclusion that the failure to satisfy (1) was the reason behind the paradoxicality of matrix **1**.

This is a natural point to reflect on what I have been meaning by "coarse graining". Suppose that the agent indeed knows that the envelopes contain either 1, 2, 4 or 8, and that she is in the decision problem of the two-envelope paradox. The six states w_0 to w_5 can be grouped in a natural way in two categories, those in which M contains more than O, that is, those for which $S(w) > C(w)$; and those in which the opposite holds. So, the agent might coarse-grain her initial set of possible states Ω, to a new set $\mathbb{\Omega}$ which contains two coarse-grained states: $\mathbb{w}_1$, corresponding to the situation where M contains more than O, and $\mathbb{w}_2$, corresponding to the situation where O contains more than M. More formally, there is a many-to-one function $f : \Omega \to \mathbb{\Omega}$, such that:

$$f(w) = \begin{cases} \mathbb{w}_1 & \text{if } S(w) > C(w), \\ \mathbb{w}_2 & \text{if } S(w) < C(w). \end{cases}$$

What would the values of $\mathbb{S}$ and $\mathbb{C}$ be for these coarse-grained states? One suggestion, in the case where we know the initial amounts that M and O may contain, such as in matrix **3**, is to make an averaging matrix such as matrix **4**. For any $\mathbb{w} \in \mathbb{\Omega}$, which is the image under f of some states $w_i, ..., w_k \in \Omega$, $\mathbb{S}(\mathbb{w}) = E[S(w_i, ..., w_k)]$. Note that this definition is such that, if the fine-grained matrix satisfies (1)-(3), so does the coarse-grained matrix, provided that $p(\mathbb{w}) = \sum_{n=i}^{j} p(n)$ for all $\mathbb{w} \in \mathbb{\Omega}$.

	$\mathbb{w}_1$	$\mathbb{w}_2$
$\mathbb{S}$	$14/3$	$7/3$
$\mathbb{C}$	$7/3$	$14/3$

Matrix 4

So far, we have been assuming that the agent knows what the possible values of S and C are. In other words, we have been assuming that the agent knows which amounts M and O may contain. But this assumption is not part of the initial situation. Is that where the problem comes from? I argue that it is not. The initial situation may be represented non-paradoxically even when the agent has no idea about what the fine-grained possible states are, so long as she knows that the coarse-grained utility matrix (using the coarse-grained possible states $\mathbb{w}_1$ and $\mathbb{w}_2$) respect constraints (1), (2) and (3). Here is an example, from Jackson et al. (1994). Call x the smallest value in the two envelopes, whatever it is. Now, there are two states of the world, one in which M contains x and O contains $2x$, and another in which M contains $2x$ and O contains x.

	M contains x O contains $2x$	M contains $2x$ O contains x
Stick	x	$2x$
Change	$2x$	x

Matrix 5

Suppose that each of these gets assigned credence $1/2$. Then, the expected values of sticking and changing are equal:

$$E[S] = 1/2(x) + 1/2(2x) = 2/3(x)$$

$$E[C] = 1/2(2x) + 1/2(x) = 2/3(x)$$

Furthermore, the three constraints identified in §2 are satisfied: (1) every amount possibly in M is also possibly in O, (2) one envelope contains twice as much as the other in each possible state of the world, and (3) it is equally likely for each envelope to contain more than the other. Note that, starting from matrix **5** and then fine-graining (maybe, because the agent then learns what the possible values in M and O are) will always yield a fine-grained utility matrix that satisfies (1), (2) and (3). For example, if she learns that the possible values are 1, 2, 4, and 8, the fine-graining of matrix **5** will be matrix **3**.

Horgan (2000) is worried about using non-rigid designators in state descriptions, as in matrix **5**, where “x” refers to a different amount in each state (more on this in §7). So, he proposes to use rigid designators instead. He names the heighest actual amount in either M or O Heidi, and the lowest Lois. The following matrix results:

	O contains $1/2$(Heidi)	O contains 2(Lois)
Stick	Heidi	Lois
Change	$1/2$(Heidi)	2(Lois)

Matrix 6

Horgan then claims that we know that Heidi $=$ 2 Lois, and that both states are equally likely.

From these two claims, it follows that the respective expected utilities of sticking and changing are equal:

$$E[S] = {}^{1}\!/\!{}_{2}(\text{Heidi}) + {}^{1}\!/\!{}_{2}(\text{Lois})$$

$$E[C] = {}^{1}\!/\!{}_{2}(\text{Lois}) + {}^{1}\!/\!{}_{2}(\text{Heidi})$$

Once again, the matrix respects (1), (2) and (3), and is non-paradoxical. This has been the case for all the matrices presented so far in this section. It has turned out that there are a number of ways we can represent the initial situation, and that all those that do not respect (1) are paradoxical, and all those that do are not. This should be an indication that (1) is really a crucial implicit assumption of the initial situation, and that failing to satisfy it, as in matrix **1**, is the source of the problem.

5. The Situation of Matrix 1

So far, I have been arguing that matrix **1**, because it does not satisfy (1), does not represent the decision problem of the initial situation. But then, what does it represent? I believe that it represents an alternative situation that has been discussed since Cargile (1992). The situation is the following: you are given an envelope, call it M, containing some amount of money x. Opposite you is another envelope, O. A fair coin is flipped and, if it lands heads, $2x$ is placed in the other envelope, and if it lands tails, ${}^{1}\!/\!{}_{2}$ is placed in O. As Cargile says, in this case, it is obviously rational to switch to the other envelope. Indeed, the expected utility of changing is higher than that of sticking for any x. To convince yourself, imagine the following situation. You are given 2, and a fair coin is about to be flipped. If it lands heads, you get an extra 2, and if it lands tails, you owe your friend 1. Surely, you should take that bet.

The coin flip situation can clearly be modelled by matrix **1**. The puzzle then, is: what is the salient difference between the initial situation of the two-envelope paradox, and the coin flip situation? And, relatedly, why are we tempted to accept that the initial situation is modelled correctly, when it is modelled in the same way as the coin flip situation? My contention is that the salient distinction between the two cases has nothing to do with causation, counterfactuals or temporality, as Katz and Olin (2007) propose. Rather, more like what Clark and Shackel (2000) suggest, it has to do with what may evasively be described as point of view.

Take the coin flip case. Whether the coin has been flipped and the money has been placed before or after you are allowed to enter the room with the two envelopes, the situation is the same to you. So, the two cases might be perfectly equivalent from these perpectives. The difference is that, in the coin flip case, you are given information about O from the point of view of having M, but, in the initial situation, you are given information about O and M from a third perspective. So, in some sense, you have more information in the coin flip case, because the number of epistemically possible states is more restricted: you are only uncertain about the value of O given the value of M. But, in the initial situation, you are not in a position to assume the value of M and ask what the value of O might be, because, by doing so, and given the restriction (2) on the relative values of M and O, you would be excluding some possible values for O. Let us illustrate this. You are considering the set of possible values $x_1, x_2, ... \in P$. You assume that M contains some $x \in P$ and you write down what that entails for the possible values of O. If this is when you draw your decision matrix, after having assumed that x is the amount in M, you are not modelling the possible states in which x is the amount in O. That is, you are modelling the situation as if you had more information than you actually have.

So, where, precisely, does the reasoning in the two-envelope paradox go wrong? It is, I contend, when you say "Let x be the value in M". When you do that, you are insidiously changing the decision problem from one where you are considering the two envelopes and their relations to one another (so, in that situation, "one contains twice as much as the other" means that every value that is allowed for one is allowed for the other – that is, (1) holds in that situation); to one where you are considering one envelope and *its* relation to another (and in that situation, (1) no longer holds).[5]To really make vivid how insidious this move is, consider the following. Suppose that you are told that the possible amounts in the envelopes are 1, 2, and 4. Suppose then that the amount in M is 2. Then, O contains either 1 or 4. So you should switch to O in order to maximise expected utility. I hope that you are outraged at the move of having fixed M to be 2. What if O contains 2 and M contains 4?!

What are we to do then, with the accusations of paradoxicality we described in §1? There were three reasons for claiming that one should not switch in the initial situation. We have shown that the initial situation, when properly modelled, does not actually yield that conclusion. But the coin flip situation does. Is that a problem? Let's go through them in turn.

a. Of the two envelopes, it is clear that none is preferable. Well, that is simply not true in the coin flip case. I would definitely pay 2 for a fair bet on either 1 or 4.

b. Some analogous argument could be given for switching back. One would simply have to

[5] Katz and Olin (2007) note that the two states "M contains x and O contains $x/2$" and "M contains x and O contains $2x$" are not jointly exhaustive, because "O contains x" is consistent with the initial situation and is inconsistent with either state. This is correct. However, they think that the reason why a state in which O contains x has been excluded from the utility matrix is that, by saying "let x be the amount in M," we introduced a rigid designator. I think this is incorrect—see §7.

call the amount in the new envelope say, y, and run the same argument. But, if one did that, after having already ran the argument with x, this would be akin to suppressing the information about the amount in O given that in M that one had in the first case. And, if one did preserve that information, rationality would dictate not to change back, as per matrix **7**.

	O contains $x/2$	O contains $2x$
Stick	$x/2$	$2x$
Change	x	x

Matrix 7

c. The same argument could be run to argue for the rationality of sticking. Such an argument would require considering the amount in M given the amount in O. But this would be a different decision problem; it would be one where you are given O, and the coin is flipped to determine the amount in M. So, that's not a problem for the coin flip case that starts with M.

So, what's really going on in the two-envelope paradox is the following. There is a first decision problem, the initial situation. Any representation of it should satisfy (1), (2) and (3). Furthermore, in this situation, it would not be rational to switch envelopes. There is also a second decision problem, the coin flip case. Any representation of that will not satisfy the constraints—and (1) in particular. Furthermore, it would be rational, in that situation, to switch envelopes once. The reasoning described in §1 goes wrong when, by stipulating that x is the amount in M, the decision problem is changed. So, the initial situation is described by a decision matrix that really represents the coin flip situation. This is why we get the paradoxical conclusion that we ought to switch envelopes. Before ending this paper, I examine in §6 and §7 some responses to the two-envelope paradox that can be found in the literature. I use the tools that we have been deploying throughout this paper to argue that these proposed solutions are misguided.

6. Boundedness and Infinity

Jackson et al. (1994) propose to solve the paradox by claiming that what goes wrong in the reasoning of §1 is that the probability assignments are incorrect. They argue that the reason we have for assigning credence $1/2$ to both possibilities in matrix **1** is that we accept that "if x is the amount of money in some particular envelope, it is equally likely that $2x$ or $x/2$ is the amount in the other envelope," and they reject this claim. Let us call this claim, as Horgan (2000) does, the Asymmetrical Symmetry Condition, and let us make it precise in the following way:

$$\text{For all } x \in P, \text{ there are two distinct states } w \in \Omega \text{ such that}$$
$$S(w) = S(w') = x, C(w) = 2x \text{ and } C(w') = x/2; \text{ and } p[w] = p[w'] \tag{ASC}$$

This says that, for every amount x that the agent considers possible, there are two possible states w and w' such that M contains x in both; and O contains $2x$ in one and $x/2$ in the other. Furthermore, these states are equiprobable. To reject this precisification, one could claim that it should not hold for every x in P; but for every x in some subset of P which would correspond to the amounts the agent considers to be possibly in M. But of course, that would require the agent to have more information than she has: she would have to be aware of some amounts that she can find in O but not in M, which is ruled out by the setup of the problem.

Once the (ASC) is made precise, it is obvious that, as Jackson et. al. claim, it is false. Indeed, provided that there is an upper and lower bound on P (i.e. that there is a highest and lowest amount of money), there is some $x \in P$ such that $2x \notin P$, and some such that $1/2 \notin P$. This entails that, there are some $x \in P$ such that there is only one state $w \in \Omega$ such that $S(w) = x$. Because they show that the (ASC) is false by appealing to the boundedness of P, and because they think that the justification for the credence assignments in matrix **1** is the (ASC), Jackson et. al. conclude that the problem with matrix **1** has to do with probability assignments and comes from the boundedness of P.

This, I think, is misguided. Indeed, matrices **2** and **3** are both bounded top and bottom, and one is paradoxical whereas the other is not.[6] So, boundedness must be independent from pradoxicality. Furthermore, a weakened form of the (ASC), call it the Conditional Asymmetrical Symmetry Condition, is both true and sufficient grounds for the probability assingments in matrix **1**. So, probabilities cannot be the culprit either.

$$\text{For any } x \in P, \text{ if any two distinct possible states } w \text{ and } w' \text{ are such that } S(w) = S(w') = x,$$
$$\text{then } C(w) = 2x \text{ and } C(w') = x/2; \text{ and } p[w] = p[w'] \tag{CASC}$$

[6] This is also a reason to reject Broome's (1995) analysis of the paradox, namely that "no distribution with a finite mean can be paradoxical" (p. 8). This is plainly false: matrix **2** together with the uniform distribution is clearly paradoxical, and yet the distribution has a finite mean. There are a number of other papers discussing infinity and the two-envelope paradox. These include Clark and Shackel (2000), Meacham and Weisberg (2003), Arntzenius et al. (2004).

The (CASC) clearly justifies the probability assignments of matrix **1**. Furthermore, not only is it true, but it is entailed by (2) and (3) from §2. This indicates that what Jackson et. al. have shown to be going wrong in matrix **1** must be what is needed to strengthen the (CASC) to the (ASC)—the assumption that, for every $x \in P$, there are two distinct states w and w' are such that $S(w) = S(w') = x$. And that is a condition about what the states are, not what their probabilities are—in accordance with what I have been arguing.

7. Designation and Intensionality

I want to finish by exploring some arguments surrounding the two-envelope paradox that mention designation. Katz and Olin (2007) and Horgan (2000) both present arguments that rely on the premise that "x" must be a rigid designator. Katz and Olin write: "if we were not using "[x]" rigidly, we could not make sense of the comparisons of utility across different possible states of the world, in which case our calculations of expected utility would not make sense" (p. 909). In a similar vein, Horgan writes: "the argument's expected utility calculations are therefore bogus, since "x" lacks a single constant referent througout the course of the calculations" (p. 584).

Although warnings about using non-rigid designators appear compelling at first, it is quite hard to pin down precisely what is wrong with them. I want to argue that, actually, they are nothing to be afraid of. First, let me start with an example. I have invited my sister, Agatha, and my friend, Camilla, over for tea. I leave them for a second to pour the water into the teapot and, when I come back, my chocolate bar is gone. I have the choice between making a remark and remaining silent. If Agatha took the chocolate bar and I make a remark, she will give me one of hers and I will have chocolate after all. Furthermore, I will feel no embarrassment. However, if Camilla took the chocolate, she would also give me one of her chocolates, but I would be mortified for having said something. I represent the problem in the following decision matrix:

	Agatha is the culprit	Camilla is the culprit
Say something	Get chocolate	Get chocolate and be mortified
Remain silent	Get no chocolate	Get no chocolate

Matrix 8

In this case, it seems quite clear that I could (at least in principle) assign utilites to the possible outcomes, probabilities to the possible states, and make an expected utility calculation. But "the culprit" is a non-rigid designator: it refers to Agatha in one state, and to Camilla in the other.

How does this consideration translate to the two-envelope case? Consider the following matrix **9**. The two possible states are: the one where the amount in M is the smallest, and the one where amount in M is the largest. Given that each state is equiprobable and that, in each state, the largest amount is twice as large as the lowest amount, an expected utility calculation seems perfectly kosher in this case.

	The amount in M is the smallest	The amount in M is the largest
Stick	The smallest amount	The largest amount
Change	The largest amount	The smallest amount

Matrix 9

Compare matrix **9** above to matrix **10** below. In matrix **10**, let "x" be the total amount in both envelopes. So, there are two possible states, one where M contains $x/3$ and O contains $2x/3$, and one where M contains $2x/3$ and O contains $x/3$. Here again, there seems to be no reason to be sceptical of the possibility of a perfectly good expected utility calculation for changing and sticking.

	M contains $x/3$ O contains $2x/3$	M contains $2x/3$ O contains $x/3$
Stick	$x/3$	$2x/3$
Change	$2x/3$	$x/3$

Matrix 10

Matrices **9** and **10** both model the initial situation. However, "the amount in M" refers to two different entities in the two different states of matrix **9**, and "x" refers to the same entity in the two different states of matrix **10**. That is, the states have been picked out with a non-rigid designator in one case, and with a rigid one in the other. This seems to me to indicate that using non-rigid designators across states to pick out those states does not automatically render the expected utility calculations "bogus" or "nonsensical", even in the case of the two-envelope paradox.

Does it, however, indicate that there is nothing to be worried about? Is using a term whose referent varies across states, just as good as using term whose referent does not, for the purpose of our utility calculation? Let us go back to the formalism. We have a set of states Ω, on which are defined two random variables S and C. The expected utility of the act of say changing, is just the expectation value of the random variable C. In order to calculate this expectation value, all we need is to know is, for each state $w \in \Omega$, its probability and the value $C(w) \in \mathbb{R}$. So, we should ask: does non-rigid designation of states pose an adiditional challenge to assigning credences and utilities to these states?

Credences, just like beliefs, are intensional. So, I can have a higher credence in "the inventor of the post-it note is a genius" than in "my neighbour is a genius", even if "the inventor of the post-it note" and "my neighbour" corefer. Similarly, on many accounts of utility, that is also intensional. For any chocolate, I might assign a higher utility to eating it under the description "a piece of dark chocolate with a touch of sea salt" than under the description "the last chocolate in the box". But, if these create any worries for expected utility calculations,[7] they do so whether rigidly or non-rigidly referring terms are involved. And, if intensionality cannot be responsible for non-rigid designators' being worse than rigid ones in state descriptions, I am unsure what could be.

I speculate that, even though considerations of intensionality actually concern matrices with both rigidly and non-rigidly designated states, they are what underpins this suspicion of non-rigid designators. In matrix **3**, it seems perfectly clear and transparent which values S and C assign to the states. But in matrix, say, **9** it is not. Before I examine this line of thought further, I want to make it clear again that this worry is independent from the rigidity of the terms in the state descriptions. Matrix **6**, which employs rigid designators, is not transparent in the way that **3** seems to be; and the following matrix **11**, which does use a non-rigidly denoting term, is.

	The amount in M is 1	The amount in M is 2
Stick	1	2
Change	2	1

Matrix 11

The difference with what I will call the mysterious matrices and the non-mysterious ones might seem evident at first: in the latter case, I know what the amounts in the envelope are, but in the former case, I don't. Horgan, in the context of an argument I admit I don't fully understand, has branded numerals as "canonical" ways of representing numeral quantities, and expressions such as "x" or "the amount in M" are "non-canonical" ways of specifying them. Maybe, a non-myseterious matrix is one that only uses canonical terms in its outcome descriptions. However, this line of argumentation presupposes an account of what terms or types of terms count as canonical.

Here are some reasons to doubt that the mysterious/non-mysterious distinction has any role to play in analysing the paradox. First, suppose that you are told that there is a set P of possible

[7] Maybe I should cite Chalmers on the objects of credence here.

amounts in M and O. You are not told what elements the set has, or even how many it has. But you are told that the expected utility of sticking, is equal to the expected utility of changing (this is a matrix representing the initial situation, so it must), and is equal to 3.5. So, you are in the situation of matrix **4**. Now, here are two possibilities that you might consider for P. Maybe $P = \{1, 2, 4, 8\}$ (this is matrix **3**). Or, maybe, $P = \{7/3, 14/3\}$. You don't know. Is this a mysterious or a non-mysterious matrix? After all, you are using numerals only. But you don't know what *really* is in the envelope. So which is it?

Secondly, what could make a term canonical as opposed to non-canonical? It could not be that the former are rigid designators, whereas the latter are non-rigid. Indeed, "Heidi" is rigid (matrix **6**), "the smallest amount" is rigid (matrix **9**), but they don't tell us what *really* is in the envelope. Maybe canonical terms are those such that we always know what they refer to. But I know what "Heidi" refers to: it refers to the largest amount in either M or O! Maybe then, they are those such that, if you gave me another canonical coreferring term, I would always know that the two corefer. But that seems too strong: if "17" and "seventeen" are canonical terms referring to 17, then presumably, so are "XVII" and "dix-sept"—but I can very well fail to know how to read roman numerals or speak French.

I have presented you with a phletora of matrices that adequately represent the problem. Some of them use rigid designators, some not. Some of them are "mysterious", some not. Some, you can't even tell. Some use "canonical" terms, some use "non-canonical" terms, but it is not obvious what differentiates them. The one thing that all these non-paradoxical matrices do have in common, though, is that they all satisfy the conditions (1), (2) and (3) that I proposed in §2. Although this is no definitive argument, it seems to me to be a compelling reason to accept that mode of designation is not relevant to the two-envelope paradox in particular. Of course, this is not to say that designation does not pose any challenges to the analysis of credences within decision theory and formal epistemology, but I do not think that it poses a special problem to the two-envelope paradox.

References

Arntzenius, F., A. Elga, and J. Hawthorne (2004). Bayesianism, infinite decisions, and binding. *Mind 113*(450), 251–283.

Broome, J. (1995). The Two-Envelope Paradox. *Analysis 55*(1), 6–11.

Cargile, J. (1992). On a Problem About Probability and Decision. *Analysis 52*(4), 211–216.

Clark, M. and N. Shackel (2000). The Two-Envelope Paradox. *Mind 109*(435), 415–442.

Horgan, T. (2000). The Two-Envelope Paradox, Nonstandard Expected Utility, and the Intensionality of Probability. *Noûs 34*(4), 578–603.

Jackson, F., P. Menzies, and G. Oppy (1994). The Two Envelope "Paradox". *Analysis 54*(1), 43–45.

Katz, B. D. and D. Olin (2007). A Tale of Two Envelopes. *Mind 116*(464), 903–926.

Meacham, C. and J. Weisberg (2003). Clark and Shackel on the Two-Envelope Paradox. *Mind 112*(448), 685–689.