Repeated St Petersburg two-envelope trials and expected value

In David Chalmers 'The St Petersburg two-envelope paradox', two envelopes are filled with money according to the following procedure: 'a coin is flipped until it comes up heads, and if it comes up heads on the nth trial, 2^n is put into the envelope' (2002: 155). Chalmers then considers and rejects reasoning which attempts to show that if given envelope A, it is in one's interest to switch to B, *before* opening envelope A. But *after* opening envelope A, it does *seem* that it is in one's interest to switch envelopes. As Chalmers writes:

> When a distribution over finite amounts has an infinite expected value, then … any specific result will be disappointing. It will then always be in one's interest to do things over, if given the opportunity. (2002: 156)

In a similar vein Arntzenius and McCarthy write:

> We would prefer to abandon our winnings for another shot at the gamble, no matter how much we won. Bizarre, but not inconsistent. The bizarreness of this case, we think, derives from the fact that we know in advance that no matter what we get, it will be less than the expected value of the gamble. (1997: 48-9)

In this paper we argue that it is not in one's interest to switch envelopes, even after opening one envelope. Put more cautiously, we argue that there is no advantage in switching envelopes over repeated trials (where a trial is described below). We suggest that a paradox arises if it is assumed that infinitely many trials have been completed. We also present a case where (repeatedly) switching to an unopened envelope with a higher expected value (than the value of one's current, opened envelope) leads to a worse overall outcome over any finite number of trials.

Imagine that infinitely many envelopes (numbered 1, 2, 3, 4…) are filled according to the St Petersburg procedure described above. Two agents then select strategies. Each will open envelope 1, and then have the option of switching to envelope 2; next each opens 3, and has the option of switching to 4, etc. In general, trial n consists of envelopes 2n-1 and 2n; 2n-1 is opened, and the agent then has the option of switching to 2n. One strategy is the always-switch strategy; another is the always-stay strategy. After infinitely many trials, the always-stay strategy results in an agent's having all of the odd envelopes; the always-switch strategy results in the evens. But there can be no reason why the evens should be preferable to the odds, that is, switching is not, in general, beneficial. We would expect each strategy to outperform the other half of the time.[1]

The same holds true if only finitely many trials are run. The evens, in general, cannot be preferable to the odds. Computer simulations reinforce this conclusion; we found that the always-switch strategy outperformed the always-stay strategy about half of the time. That is, switching, even after opening the odd envelope, is not beneficial. Notice that actually opening the odd numbered envelopes is playing no role in either strategy. That is, one strategy (always-stay) gets the odds, the other (always-switch) gets the evens, regardless of what is initially seen in any odd envelope. Though it is not beneficial to always-switch as compared with always-

[1] Though, especially in the case where finitely many trials are run, some percentage would result in ties.

staying, it is beneficial to switch *selectively*, that is, to switch depending on the amount that the odd envelope contains. In this way, opening the odd envelopes plays a role. For example, the strategy of switching if the odd envelope contains $2 is beneficial because the even envelope cannot contain less.[2] Then, depending on the number of trials that are run, switching out of higher values than $2 can also be beneficial. As Feller argued and Easwaran recently discussed:

> [Feller] suggests that a sequence of n plays of the St Petersburg game should have value n log n … Thus, it is suggested that the longer a sequence of plays of the St Petersburg game, the more one should be willing to pay for each (one should be willing to pay log n for each of n total plays). (Easwaran 2008: 638)

We ran computer simulations many times of many different numbers of trials, but in this paper let us focus on and discuss runs of 1,048,576 (=2^{20}) trials. According to Feller and Easwaran, one should place a value of $20 per envelope for this many plays of the St Petersburg game. As a slight refinement, we believe that n St Petersburg games should be valued at $(2 + log n) per play. Over many such runs, the average of the 1,048,576 trials did hover around $22 (for both the always-stay and always-switch strategies), though the average would often exceed $22 by a great amount, and would rarely go below $18. For example, on three runs on different underlying datasets, the always-stay strategy averaged $22.78 per envelope, the always-switch strategy averaged $21.73, and the switch if the odd envelope contained $2 averaged $41.21 (note that with a selective switching strategy, more than 1,048,576 envelopes are being considered).

[2] We use the '$' to ease reading, that is, to disambiguate envelope numbers from the value an envelope contains.

We give these examples because we felt that these results were typical, which is a subjective evaluation. Note that it does not make sense to repeat the runs and give averages (of averages), because that would be equivalent to a single longer run and so should have a higher value. To put this point another way consider the story of the king and the peasants. A seller is selling 1,048,576 St Petersburg games. Note that each game can be sold multiple times. The king buys all 1,048,576 games, and so pays $22 per game. Each of 1,024 peasants buys 1,024 games, and so each pays $12 per game. The king has paid $10 more per game than the peasants, for the same 1,048,576 underlying games; sometimes it is not good to be the king.[3] Note that it may be difficult for the seller and a peasant to come to an agreement on price. The seller knows that 1,048,576 games are to be run, and so values them at $22 each, whereas a peasant, even if he knows this fact, is only buying 1,024 games and so only wishes to pay $12.[4]

Let us now address the question: How is it that, in general,[5] the always-stay strategy can do as well as the always-switch strategy? What emerges from considering a finite number of trials is that the always-switch strategy benefits when switching out of low values (of odd, opened envelopes), hurts when switching out of high values, and the two balance each other out. This is how the always-switch strategy ends up being the same as the always-stay strategy. For example, of 1,048,576 odd envelopes, roughly half should contain $2. Then, the always-switch strategy means that these odd envelopes containing $2 are switched out of roughly 524,288 times. It is reasonable to value this many new St Petersburg envelopes at $21 each. Thus, the

[3] At least as compared with *collusive* peasants, that is, peasants who get together and average their winnings.
[4] We did look at 1,024 averages, each of 1,024 St Petersburg games (each average corresponds to a peasant's average value per envelope). The median of the 1,024 averages hovered around 12, whereas the average of the averages was often near 22 (which corresponds to the king's average value per envelope). As far as we could tell, most averages were near 12, but some were strongly atypical, meaning simply that they were very far above 12 (e.g., 517). A few peasants got very lucky.
[5] That is, obviously on any given run one strategy can, and generally will, outperform the other. But each exceeds the other half of the time.

gain from switching out of $2 is roughly 524,288 * $19 = $9,961,472, where the $19 is $21 - $2. Repeated computer simulations confirmed these numbers. But switching out of high values hurts. For example, there should be roughly 2,048 odd envelopes containing $512. Switching 2,048 times has a $13 per envelope value. The loss then from switching out of odd envelopes containing $512 is roughly 2,048 * $499 = $1,021,952, where the $499 is $512 - $13. The following table gives the relevant data for 1,048,576 trials (where the data discussed above is in bold):

roughly this many envelopes	contain this in value ($)	value of each switch ($)	added value per switch ($)	added value of all switches ($)
524,288	**2**	**21**	**19**	**9,961,472**
262,144	4	20	16	4,194,304
131,072	8	19	11	1,441,792
65,536	16	18	2	131,072
32,768	32	17	-15	-491,520
16,384	64	16	-48	-786,432
8,192	128	15	-113	-925,696
4,096	256	14	-242	-991,232
2,048	**512**	**13**	**-499**	**-1,021,952**
1,024	1,024	12	-1,012	-1,036,288
512	2,048	11	-2,037	-1,042,944
256	4,096	10	-4,086	-1,046,016
128	8,192	9	-8,183	-1,047,424
64	16,384	8	-16,376	-1,048,064
32	32,768	7	-32,761	-1,048,352
16	65,536	6	-65,530	-1,048,480
8	131,072	5	-131,067	-1,048,536
4	262,144	4	-262,140	-1,048,560
2	524,288	3	-524,285	-1,048,570
1	1,048,576	2	-1,048,574	-1,048,574
				Total = 0

The data presented above do outline how the always-stay strategy can manage to be as good as the always-switch strategy, though the story is more complicated. For example, the data in the table do not consider values of more than $1,048,576, though e.g., there is a 39% chance that one or more envelopes (out of 1,048,576 envelopes) will contain $2,097,152 (= 2^{21}). Interestingly, the table also indicates that an envelope should be valued at $20,[6] though we believe that envelopes should be valued at $22.[7]

To this point we have argued that, on an intuitive level, it cannot benefit an agent to always switch to the even envelopes, even after opening the odds. The even envelopes cannot be preferable to the same number of odds. Selectively switching is advantageous if the right strategy is used, where the correct strategy will depend on Feller's result discussed above. We have also indicated the reason why the always-switch strategy is not advantageous, namely that though switching out of low values is beneficial, switching out of high values is detrimental, and the two cancel each other out.

We now show that it is possible to always opt for a better expected value (EV), and yet do worse over any finite number of trials. Imagine the following strategies. One strategy is always-stay. The other is a modified-switch strategy, namely switch to the even envelope with the following condition: if the even envelope contains more than \$$2^{\text{ODD_ENV} + 1}$, then $0 is given, where ODD_ENV is the odd envelope's value. For example, if the odd envelope contains $8, then if the even envelope contains more than $2^{8+1} = 2^9$ = $512, then $0 is given. Note two facts. 1) This modified-switch strategy is worse than the always-switch strategy, in that sometimes $0 is given, instead of the even envelope's value. 2) Switching is recommended based on EV, in that the EV of switching is ODD_ENV + 1, as compared with a certain value of ODD_ENV. That is, the expected gain from switching is $1.

---

[6] Multiply the first two columns by row, add those values, and divide by 1,048,575.

[7] Or more generally, n games should be valued at 2 + log n. We base this conclusion on our computer simulations (we estimate that we have run one billion St Petersburg games), and a consideration of low numbers of games, e.g., 2 games should be valued at $3 each (and not $1).

It follows that the modified-switch strategy is worse than the always-stay strategy, as

always-stay = always-switch > modified-switch, and so

always-stay > modified-switch

Modified-switch is worse than always-stay[8] even though modified-switch *is recommended* over always-stay based on a consideration of EV. Computer simulations supported the conclusion that always-stay outperforms modified-switch. On many runs of 1,046,578 trials the always-stay strategy hovered around $22 per envelope, whereas the modified-switch strategy hovered around $6 per envelope. Always-stay outperformed modified-switch every time.[9] Expected value is not always all that it is cracked up to be.

Above we argued that the always-switch strategy cannot be better than the always-stay strategy, over finitely many or infinitely many trials; the evens cannot be preferable to the odds. Let us conclude with reasoning that attempts to highlight the paradoxical nature of this result by now arguing that *it is* in the agent's interest to always switch over infinitely many trials. Note that when there are *finitely* many trials, the high values have not been hit enough times to make switching out of those odd envelopes worthwhile. But as soon as the game is played *infinitely* many times, then it is almost certain that every value will be hit infinitely many times. Then switching out of every value is beneficial. That is, we can ask: When switching out of odd envelopes containing $2, what resulted? The answer is benefit, as 1/2 of the infinite is infinite, and so we have switched out of $2 enough times to benefit. When switching out of $4, what

[8] This, of course, is a probabilistic conclusion. Modified-switch could outperform always-stay.
[9] Over approximately 100 comparisons.

resulted? Benefit, by the same reasoning. In general, for any n, when the initial odd envelope contained $\$2^n$, there is benefit in switching out of those (infinitely many) envelopes. We can also ask: When is the always-switch strategy harmful? That is, present an n, such that switching out of $\$2^n$ is harmful. There are no such n. Then switching is benefiting a great deal and never harming. Thus switching is beneficial. That is, in the infinite case the always-switch strategy outperforms the always-stay strategy. Of course, by the initial reasoning (the evens cannot be preferable to the odds) switching is not beneficial. We therefore believe that the infinitely repeated St Petersburg two-envelope case is paradoxical.[10]

It is commonly believed that when a finite value is received in a game that has an infinite EV, it is in one's interest to redo the game (note that switching to an unknown but already determined value is equivalent to redoing). We have argued against this belief, at least in the repeated St Petersburg two-envelope case. We have also shown an example where repeatedly switching to a higher EV leads to a worse outcome over any finite number of trials. Finally, over infinitely many trials of the repeated St Petersburg two-envelope game, the always-switch strategy is paradoxically both better than and the same as the always-stay strategy.

References

Arntzenius, F. and McCarthy, D. 1997. The two envelope paradox and infinite expectations. *Analysis* 57: 42-50.

Chalmers, D.J. 2002. The St Petersburg two-envelope paradox. *Analysis* 62: 155-57.

Easwaran, K. 2008. Strong and weak expectations. *Mind* 117: 633-41.

[10] Our way out of the paradox is to suggest that it is impossible to complete infinitely many tasks. But a further consideration of this complex and much discussed topic is beyond the scope of this paper.

Feller, W. 1968. *An introduction to probability theory and its applications*. 3rd Edition, Volume 1. New York: John Wiley & Sons.