

Critical Thinking, Summer 2004: The Two-Envelope Paradox

Tom Lored, Dept. of Astronomy, Cornell University

15 July 2004

1 The Problem

Were you able to resolve the puzzle Prof. Terzian presented to your satisfaction? If so, congratulations! This puzzling problem, known most commonly as the *Two-Envelope Paradox* or the *Exchange Paradox*, has been occupying the minds of mathematicians and philosophers at least since the 1950s, and is related to another tricky paradox, the *St. Petersburg Paradox*, that dates back to Daniel Bernoulli and the earliest days of probability theory (the 1700s). You may have encountered the problem before. Martin Gardner, a famous popularizer of mathematics and long-time author of the “Mathematical Games” column in *Scientific American*, presented the problem (without a solution!) in his 1982 book, *Aha Gotcha: Paradoxes to Puzzle and Delight*. Genius columnist Marilyn vos Savant presented the paradox (with an incorrect solution!) in a 1992 “Ask Marilyn” column in *Parade* magazine. Philosophers and mathematicians are still writing papers about both the two-envelope and St. Petersburg paradoxes.

I think it fair to say that most mathematicians would consider the two-envelope paradox to be resolved for all practical purposes. But there are some loose ends that continue to trouble some thinkers (philosophers particularly). These concern versions of the problem that could never actually arise in practice, but that seem on the face of them to be well-posed mathematically, at least to some writers. Yet it is not obvious how to resolve these theoretical versions of the paradox.

In this note I’ll try to explain the practical resolution of the paradox and convey some idea of what the loose ends are, in not-too-technical language. Let’s begin by restating Prof. Terzian’s version of the paradox:

In a game show there are two closed envelopes containing money. One contains twice as much as the other. You [randomly] choose one envelope and then the host asks you if you wish to change and prefer the other envelope. Should you change? You can take a look and know what your envelope contains.

Say that your envelope contains \$20, so the other should have either \$10 or \$40. Since each alternative is equally probable, the expected value of switching is $(1/2) \times \$10 + (1/2) \times \40 which equals \$25. Since this is more than your envelope contains, this suggests that you should switch. This reasoning works for whatever amount you find in your envelope. So it does not matter if you looked in your envelope or not.

But your envelope is as likely to contain twice as much as the other envelope, and if someone else was playing the game and had chosen the second envelope, the same arguments as above would suggest that that person should switch to *your* envelope to have a better expected value.

The problem states, “You can take a look and know what your envelope contains.” But you were not compelled to take a look. It is useful at this point to distinguish two versions of the game, depending on whether you know what’s in your envelope. If you’ve opened the envelope and learned the amount, you are confronted with the *open envelope problem* (OEP for short). If not, you must address the *closed envelope problem* (CEP).

We can use these two versions of the two-envelope game to clarify what exactly is puzzling or paradoxical about it. It appears obvious that in the CEP, there is no advantage to switching; and this conclusion is symmetric, as the second player would make the same conclusion regarding the other envelope. In the OEP, opening the envelope provides you with information not available in the closed-envelope version of the game; for this reason, perhaps we shouldn’t be too surprised that the open-envelope player reaches a different conclusion than the closed-envelope player. However, the open-envelope argument appears to hold *whatever the value found in the envelope*, therefore it doesn’t matter what value is found, and the player needn’t have opened the envelope at all. So the main puzzle is this conflict between the strategies chosen in the closed-envelope and open-envelope versions of the game. An additional puzzle is that the open envelope argument applies equally well to *both* envelopes, so that both players always believe it to be to their advantage to switch.

2 Resolving the Paradox

Resolving the problem—or *problems*—requires attention to a few different factors; this confluence of factors is part of what makes the problem tricky. I’ll provide a quick summary of the three main elements of the resolution, and then explain each of them more fully in subsequent sections. If you haven’t resolved the paradox, you might just read this summary and think more about it further on your own before reading more of the details.

We need a bit of mathematical notation before we start (indeed, part of what makes this puzzle puzzling is the absence of careful notation in the problem statement, allowing some questionable assumptions to slip by unnoticed). Let $p(\mathcal{S}|I)$ stand for the probability that statement \mathcal{S} is true given information I . By “probability” we here mean the extent to which the available information, I , implies the truth of the hypothetical statement, \mathcal{S} , where $p = 1$ corresponds to certainty that \mathcal{S} is true, and $p = 0$ corresponds to certainty that \mathcal{S} is false. Denote the amount in your envelope by A , and the amount in the other envelope by B . Then, for example, $p(A > B|I)$ is the probability that your envelope has the larger amount, given the information you have about the game (we could also have written this as $p(A = 2B|I)$, the probability that A is twice B). Throughout the rest of this note, we will take I to denote the information you have about the game *before* opening the envelope. So the closed-envelope player must make a decision based on this information, while the open-envelope player can use the amount found in his envelope as additional information.

Here are the three main ingredients for resolving the paradox:

- *Missing information:* The OEP is not a well-posed problem—there is not enough information specified in order to find a solution. If you carefully use the rules of probability theory to try to justify the quick calculation given in the statement of the problem, you will find you need further information in order to calculate the probabilities and expected values. One way to satisfy the needs of probability theory is for the host to give you enough information so that you can assign a probability that the amount in the envelope with the smaller prize, S , is equal to some value s ; i.e., you need to know $p(S = s|I)$. To make the notation a little less cumbersome, let’s call this some function, $f(s)$, so that $f(\$20)$ is the probability that there is \$20 in the envelope with the smaller prize. Once this function is specified, you can solve the

OEP (you'll find that sometimes you should switch, and sometimes not). You can also solve the CEP, and you'll find that for any legitimate $f(s)$ function, you reach the same conclusion: there is no benefit in switching if you don't know how much is in your envelope. (So in a sense the CEP is well-posed even without enough information to completely specify $f(s)$.)

- *Conditioning:* It seems obvious that $p(A > B|I) = 1/2$ and $p(B > A|I) = 1/2$; that is, it is just as likely that your envelope has the larger amount as it is that it has the smaller amount, given only the basic information describing the game. In fact, you can use probability theory to *derive* this result using any legitimate $f(s)$. But when you calculate the expected value for exchanging envelopes, these are *not* the probabilities that you must use! Rather, you must use probabilities that include the information you gained by looking in the envelope. The probability that you have the envelope with the larger prize, *given that you know there is \$20 in your envelope*, would be written $p(A > B|A = \$20, I)$. This is called a *conditional probability*, and incorporating that extra information is called conditioning. When you condition properly, you can show that sometimes the expected value in the other envelope is greater than the amount you have (so you should switch), sometimes it is less (so you should not switch), and sometimes it is equal to the amount you have (so it doesn't matter whether you switch or not).
- *Infinities:* In any real instance of the two-envelope game, the host will have a finite amount of money at her disposal. If the game is played with US currency, the most she could have is the total amount currently in circulation (and presumably she has much less than that!). Knowing that the host's resources are finite is enough to make the CEP well-posed (with the conclusion that there is no benefit to switching). It also guarantees that the OEP has a sensible solution, provided the available information lets you assign the $f(s)$ function. But there is nothing in the *mathematics* that seems to require that the game be played with finite resources. If you imagine playing with infinite resources, there are $f(s)$ probability distributions that make the problem perfectly solvable, but others (that assign relatively large probabilities to extremely large amounts) that seem to imply paradoxical reasoning like that given in the original problem description.

So, in summary, if the game is played with finite resources, it can be nicely resolved. For the closed-envelope version, calculation verifies that there is no benefit to switching, preserving the symmetry that seems evident in the problem. For the open-envelope version, knowing the amount in your envelope but not the other breaks the symmetry of the problem. One finds that the expectation value calculation in the problem statement cannot be valid; the correct calculation will lead you to sometimes prefer switching and sometimes not, depending on information not provided in the problem statement, but necessary for there to be a solution. The other player will sometimes reach the same conclusion about his envelope, but must also sometimes reach the opposite conclusion and thus not agree to switch. If the game is played with infinite resources, there are circumstances where, even given the kind of information that makes the finite-resource version solvable, one may not find a satisfactory resolution of the puzzle. And now on to the details!

3 Well-Posed Problems

Mathematics has the notion of *well-posed problems*. A problem with enough information specified to lead to a unique solution is well-posed. A problem may initially appear well-posed, but when you try working through it, you'll find you need to know information not provided in the original

problem statement in order to proceed to a single, unique solution; without that information, many solutions are possible. Such a problem is said to be *underconstrained*. In order to solve it, you must either obtain the needed information, or make reasonable assumptions (realizing that your solution depends on the truth of the assumptions). In other cases, you might work through a problem and find there is no possible solution. Such a problem is said to be *overconstrained*, or the problem statement is said to be *internally inconsistent* (some of the information given rules out all the solutions allowed by the rest of the information). In simple problems, the requirements for the problem to be well-posed are so obvious that one needn't pay too careful attention to the problem statement. But sometimes even apparently simple problems require special care to ensure they are well-posed. (I should also mention that mathematics also treats a class of problems termed *ill-posed*, but this refers to problem characteristics quite different from what we are talking about here.)

As a simple (and almost silly) example, consider this problem:

Two cars are traveling along the same road. Car *B* is traveling 10 mph faster than car *A*. After one hour, how far apart are the cars?

It may be tempting to conclude that the cars will be 10 mi apart, since $(10 \text{ mph}) \times (1 \text{ hr}) = 10 \text{ mi}$. But this isn't (necessarily) correct. It would be correct if the cars were next to each other at the start of the one hour interval, but this was not stated in the problem. In fact, if car *B* was behind car *A* at the start of the interval, it could be behind, even with, or ahead of *A*, depending on exactly how far behind it was. To solve the problem, you need additional information.

This kind of situation arises all the time in physics and astronomy. Newton's laws of motion tell us exactly how an interplanetary satellite will move in space. But the laws of motion alone aren't sufficient to determine where a particular satellite will be at a particular time unless we are given its *initial conditions*: the time of the launch, the location it was launched from, and what impulse (change in velocity, with respect to the center of mass of the solar system) it was given at launch (we are assuming here that the rocket launching the satellite burns for only a short period of time, otherwise we need further details about the rocket burn). These quantities aren't always trivial to find; for example, the position and velocity at launch will depend on where the Earth is in its orbit around the sun, where on the Earth the launch is, and how fast the Earth is rotating. Similarly, the laws of fluid mechanics tell how an element of gas orbiting a star will move. But astronomers cannot calculate the flow of gas near a star without specifying *boundary conditions* for the flow—the velocity and density where the flow originates, and where it ends. For example, a flow accreting onto a neutron star will hit the star's hard surface. That collision gets communicated up the flow (via sound waves and radiation) and affects the flow. In contrast, a flow accreting into a black hole simply disappears at the black hole's horizon. The absence of a hard surface changes the character of the entire flow, even though exactly the same physical laws “direct” both flows away from the stars. Astronomers can use this dependence of the flow properties on the boundary conditions to try to distinguish neutron stars from black holes.

The point of these examples is that well-posed problems typically require that one supplement the fundamental *laws* governing a phenomenon with *extra information*, initial conditions and boundary conditions in these examples. The same is true of probability theory. The laws of probability theory tell you how to calculate desired probabilities from other probabilities. But if you ever hope to get a number out of your calculation, “inputs” must be provided, numbers for the elementary probabilities that get manipulated. The analog to initial or boundary conditions are so-called *direct* or *prior probabilities*. For a problem in probability theory to be well-posed, it must include enough information so that numerical values can be assigned to any of the prior probabilities needed for

calculating the final probabilities or expectations of interest.

To see that additional information is needed to resolve the paradox, consider two different concrete realizations of the game. First, imagine that the game is part of a game show where the host spins a wheel to determine how much money to use as the smaller prize, with twice that amount used as the larger prize. Suppose you are shown the wheel, but that you don't see the outcome of the spin. You note that the wheel has, say, 20 amounts listed on it, in whole dollar amounts ranging from \$20 to \$2000. Now if you open your envelope and find \$20 in it, you know with certainty that the other envelope cannot have \$10 in it (since \$10 was not on the wheel). It must have \$40 in it, so you obviously decide to switch. Alternately, if you opened your envelope and found \$3000 in it, you would know not to switch; that amount is larger than the largest amount on the wheel so your envelope cannot be the one with the smaller prize.

Or suppose instead that you were playing this game as a parlor game with a friend who you knew to be hard up for money. If you found, say, \$100 in your envelope, you would probably decide to keep it, reasoning that it is improbable your friend would risk twice that much on just a game, given his meager resources. (Or perhaps you would altruistically choose to switch, reasoning that there must be just \$50 in the other envelope, and wanting your friend to keep the larger sum!)

These two scenarios indicate that knowledge of how probable it is that a certain amount is in play should affect your decision. This is born out in a more rigorous calculation; more on that as we address the second ingredient of the resolution of the paradox.

4 Conditioning

Having established that information that was missing from the problem statement should influence the player's decision, how exactly can the information be used? It appears explicitly in the calculation of the expected value in the other envelope—*when one calculates it correctly*. The expected amount in the other envelope as calculated in the problem statement is,

$$\$10 \times p(A > B|I) + \$40 \times p(A < B|I), \quad (1)$$

with the first term accounting for the case where our envelope has the larger amount ($A > B$) and the second term accounting for case where the our envelope has the smaller amount ($A < B$). The two probabilities here are each equal to $1/2$, and that gives the \$25 expectation reported in the problem statement. But this is *not correct*. The dollar amounts multiplying each term use the information that the player knows there is \$20 in his envelope, and consistency requires that the probabilities also take this information into account. The proper calculation is thus,

$$\$10 \times p(A > B|A = \$20, I) + \$40 \times p(A < B|A = \$20, I), \quad (2)$$

with the *conditional* probabilities appearing. There is a simple and powerful theorem in probability theory that one can use to calculate these probabilities called *Bayes's theorem*. But we can get the correct values with a little qualitative reasoning. The first probability—the probability that we have the greater amount, given that we have \$20—must be proportional to the probability that the *smaller* amount is \$10, which we have denoted $f(\$10)$. The second probability—the probability that we have the smaller amount given that we have \$20—must be proportional to $f(\$20)$. Since those two possibilities are the only ones (we either have the larger or smaller amount), the probabilities must sum to one. To guarantee this, we just divide the two f values by their sum, so that

$$p(A > B|A = \$20, I) = \frac{f(\$10)}{f(\$10) + f(\$20)} \quad \text{and} \quad p(A < B|A = \$20, I) = \frac{f(\$20)}{f(\$10) + f(\$20)} \quad (3)$$

You can see these two probabilities have the right proportionality and sum to one. If instead of \$20 we found some other amount, x , in the envelope, the expected amount in the other envelope would be,

$$\frac{x}{2} \times p(A > B|A = x, I) + 2x \times p(A < B|A = x, I), \quad (4)$$

with the probabilities given by,

$$p(A > B|A = x, I) = \frac{f(x/2)}{f(x/2) + f(x)} \quad \text{and} \quad p(A < B|A = x, I) = \frac{f(x)}{f(x/2) + f(x)} \quad (5)$$

Now that we have the correct formula for the expected value in the other envelope, how does it compare with the formula given in the problem? To get the equal values of $1/2$ assumed in the problem, the numerators in equation (5) must equal each other (since the denominators are the same). So we must have

$$f(x/2) = f(x), \quad (6)$$

that is, the probabilities that the smaller amount are equal to what we have (x), or half of what we have ($x/2$), must equal each other. This is certainly possible for some values of the amount, x , in our envelope. But it can't be true for *all* values. Consider the two scenarios at the end of the previous section. In the game show version of the game, if our envelope has the smallest amount on the wheel, $x = \$20$, then $f(x/2)$ is zero while $f(x)$ has some nonzero value. Similarly, if our envelope has an amount above the maximum value on the wheel, then $f(x) = 0$, but $f(x/2)$ is nonzero. In both situations, the two probabilities appearing in the expected value calculation are *not* $1/2$, as was assumed in the problem statement (in fact, one of them is 1, and the other is zero). So for the game show version of the game, it simply cannot be true that the expected value calculation gives the same result for *any* value found in our envelope, and the paradox is avoided. For the parlor game version, the poor host will have some maximum value she can gamble (e.g., her entire wealth), and presumably will be less likely to gamble large amounts than small ones. So $f(x)$ will be smaller than $f(x/2)$, and as x increases it will eventually vanish while $f(x/2)$ remains finite. Again the paradox is averted.

In addition, if you know a little probability theory, it is possible to use the equations before us to show consistency between the closed and open formulations of the game in following sense. The closed-envelope player can *imagine* that he opens the envelope and finds x in it. He would reason as described above, and decide whether to switch or not depending on whether the correct expectation calculation (conditioning on x) gave an expected amount greater or less than x . But since he is only imagining opening the envelope, he should make his final decision by averaging the expectation over the possible values of x . It is not difficult to show that such a calculation leads to an expected value in the other envelope exactly equal to to the expected amount in the first envelope for any legitimate $f(s)$ function (we'll say more in a minute about what we mean by "legitimate"). So the closed-envelope player reaches the same conclusion whether he reasons directly that there is no preference, or instead reasons the same way as the open-envelope player, but averages the calculation over all possible amounts in the envelope. This harmonizes the close-envelope and open-envelope versions of the game.

5 Infinities

What did I mean by "a legitimate $f(s)$ " above? The answer has to do with handling cases that allow infinite gambles. I can't fully address the subtle issues that arise in such cases here, but I hope to communicate some of the basic aspects of the problem.

We established previously that for the reasoning in the problem statement to hold, we must have $f(x) = f(x/2)$ for *all* possible amounts in the envelope. This cannot be true if there is a maximum amount, M , allowed in the envelope with the smaller prize. Since f is the probability for the *smaller* amount, if we happen to get the envelope with the larger amount, it can sometimes be greater than M (it can range up to $2M$), and $f(x) = 0$ for such values.

But what if there were *no limit* to the amounts in the envelopes? In this case $f(x)$ need never vanish. It turns out the reasoning in the problem statement must *still fail*. The only way it can be true is if $f(x) = f(x/2)$ for an *infinite* number of values of x . For this to be true, $f(s)$ must equal some constant, C . But since there must be an infinite number of gambles in this case, the sum of all the $f(s)$ values will be infinite for any C except $C = 0$. So there can be no game fitting the conditions of the problem statement, even admitting the possibility of infinite gambles.

The key to this argument is the requirement that the probabilities in a probability distribution must sum to one. The easiest (and most realistic) way to guarantee this is for there to be a maximum possible s . When this is true, the reasoning of the previous section always holds, and there is no paradox.

However, we can let the possible amounts extend to infinity so long as $f(s)$ falls with s rapidly enough so that the sum of all probabilities converges to one. Suppose the amounts that might be gambled are the positive integers, $1, 2, 3, \dots$. If we take the probability for the \$1 amount to be $1/2$, and the probabilities for the subsequent amounts to each be $1/2$ times the previous probability, then the probabilities look like this:

$$f(\$1) = \frac{1}{2}; \quad f(\$2) = \frac{1}{4}; \quad f(\$3) = \frac{1}{8}; \quad f(\$4) = \frac{1}{16}. \quad (7)$$

These probabilities fall quickly enough that their sum is one, even though there are infinitely many terms in the sum. In fact, more than 99.9% of the probability is in just the first 10 terms; the remaining fraction of a percent of the probability is spread over the terms from 11 to ∞ . For this distribution we can calculate the expected amount of the smaller prize,

$$\langle s \rangle = \$1 \cdot f(1) + \$2 \cdot f(2) + \$3 \cdot f(3) + \dots, \quad (8)$$

and it is highly unlikely that s will be more than several times larger than $\langle s \rangle$. For the $f(s)$ of equation (7), the expected value of s is just 2. For games described by distributions that fall quickly like this, with finite expected values, all of our conclusions above hold, even though gambles could have arbitrarily large values. The probabilities fall so quickly that, for practical purposes, it's as if there is an effective maximum amount, since the probabilities for very large amounts are vanishingly small compared to the amounts themselves.

But now imagine the game with the probabilities in equation (7), but with the possible amounts gambled given by *powers of two* rather than multiples of one. Then the game is described by this distribution:

$$f(\$1) = \frac{1}{2}; \quad f(\$2) = \frac{1}{4}; \quad f(\$4) = \frac{1}{8}; \quad f(\$8) = \frac{1}{16}. \quad (9)$$

The probabilities again sum to one, but the expected amount of the smaller prize is now,

$$\begin{aligned} \langle s \rangle &= \$1 \cdot f(1) + \$2 \cdot f(2) + \$4 \cdot f(3) + \$8 \cdot f(3) + \dots, \\ &= \$1 \times \left(\frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \dots \right). \end{aligned} \quad (10)$$

This sum is obviously infinite. For this version of the game, even though there is a 99.9% probability that the smaller prize is \$512 (that's 2^9) or less, the expected amount in play is infinite! Even

though the probabilities fall quickly, the amount gambled *grows* so quickly that one is forced to seriously consider the possibility of an arbitrarily large gamble.

For versions of the game with infinite expectations, one can create paradoxes. They can typically be resolved by following the mathematically conservative (and wise!) practice of considering a problem with an infinity to be well-posed only if it is formulated as *an infinite limit of a finite problem*. In such cases, calculations like that above apply at each stage of the limit, and things usually work out fine. But if one doesn't specify the infinite game as a well-defined limit of a finite game, there may be no unique resolution to the paradox.

So to answer the question that began this section, a “legitimate” version of the game (defined by an $f(s)$ function) is one where the $f(s)$ distribution sums to one, and where the expected amount in play is finite. These conditions are met for any $f(s)$ when there is a maximum amount that might possibly be in play, but are only sometimes met when an arbitrarily large amount could be in play.

6 Losse Ends

There is a significant literature on the two-envelope problem, with papers being written on it to this day. Perhaps surprisingly, it was only within the last decade or so that the paradox was clearly resolved for games with finite expectations. The list below includes some papers that explain the resolution with more mathematical detail than I've provided here.

Recent papers on the paradox address two loose ends in the above discussion. The first concerns the assignment of the distribution $f(s)$ describing the game. In the game show version of the game, it is easy to see how one would assign this function (assuming the host is playing fairly!). If the 20 possible amounts are listed along the edge of the wheel in arcs of equal length, then one would set $f(s) = 1/20$ for each amount s appearing on the wheel. If the arcs are of unequal length (e.g., if the lower amounts are in big arcs and the larger ones in small arcs, as is typically the case in arcade games), then the probabilities will just be proportional to the arc lengths.

But what of other versions of the game, such as the parlor version? How should I assign $f(s)$ based on what I know about my friend's wealth? More generally, how does one convert vague knowledge into a specific prior probability assignment? This is an important open problem in probability theory, important not only for the two-envelope problem but also for practical use of probability theory in analyzing scientific data. One can view such analyses as quantifying what we know about a phenomenon upon learning of some new data about it. This corresponds to calculating a conditional probability (conditioning on the new data), and the result depends on the (unconditioned) prior probability—what you know after getting the data depends on what you knew before you had it. Fortunately in most scientific settings, the data are so informative that the final conclusions are largely independent of the prior probabilities (we learn so much that whatever we knew before is swamped by the new information). But this is not always true, particularly in forefront areas of science, so scientists must sometimes think carefully about prior probabilities. The two-envelope problem is an example of a problem where the conclusions depend sensitively on the prior information.

The second focus of recent writing on this problem is on the role of infinities in the problem (and in probability theory and decision theory more generally). The mathematics becomes more subtle in such cases, and some recent literature studies the extent to which the problem can be considered well-posed when infinite gambles are permitted.

Several of the papers below treat these issues in greater length than is possible here. Some are available on the web as noted. Many provide references to additional literature on this fascinating

“paradox.”

References

- [1] C. J. Albers, B. P. Kooi, & W. Schaafsma (2003) “Trying to Resolve the Two-Envelope Problem,” *Synthese*. Argues that the problem has no satisfactory solution because of the problem of assigning prior probabilities. <http://www.math.rug.nl/casper/>
- [2] N. M. Blachman (1996) “Comment on Christensen and Utts, Bayesian resolution of the ‘Exchange Paradox’,” *The American Statistician*, **50**, 98–99.
- [3] S. J. Brams & D. M. Kilgour (1995) “The box problem: To switch or not to switch,” *Mathematics Magazine*, **68**, 27–34. A very clear presentation of the algebra for various general cases.
- [4] D. J. Chalmers (1994) “The two-envelope paradox: A complete analysis?” A philosopher weighs in on loose ends in analyses presented in the philosophical literature. jamaica.u.arizon.edu/chalmers/papers/envelope.html
- [5] D. J. Chalmers (2002) “The St. Petersburg two-envelope paradox,” *Analysis*, **62**, 155–157. Discussion of the infinite-expectation version of the paradox described above. jamaica.u.arizon.edu/chalmers/papers/stpete.html
- [6] R. Christensen & J. Utts (1992) “Bayesian resolution of the ‘Exchange paradox’,” *The American Statistician*, **46**, 274–276. Be sure to also see Blachman (1996) for a correction.
- [7] E. Linzer (1994) “The two-envelope paradox,” *The American Mathematical Monthly*, **10**, 417–419. A treatment in terms of simple algebra along the lines of this note, but more brief.
- [8] J. D. Norton (1998) “When the sum of our expectations fails us: The exchange paradox,” *Pacific Phil. Quarterly*, **79**, 34–58. Focuses on infinite versions of the game. http://www.pitt.edu/jdnorton/papers/Exchange_paradox.pdf
- [9] G. Priset & G. Restall (2003) “Envelopes and indifference.” An accessible treatment focusing on taking care in assigning the probabilities needed. <http://consequently.org/papers/envelopes.pdf>
- [10] T. Ridgway (1993) “Comment on Christensen and Utts, Bayesian resolution of the ‘Exchange Paradox’,” *The American Statistician*, **47**, 311. Quotes vos Savant’s “Ask Marilyn” column and points out her error.