Problem
Consider two envelopes, one containing precisely twice as much money as the other one. The two envelopes are shuffled and, without preference one way or the other, one of the envelopes is handed to you. Without looking inside it, you are given the option of swapping envelopes. Do you swap?

Consider the following approach: Let's call the amount of money inside the envelope given to you, x. Then the other envelope can either contain $x/2$ or $2x$. Since the two envelopes are randomly shuffled, there is no reason to prefer one of these over the other and thus the probability of these two cases are 50/50. If we do not swap, we are guaranteed to earn x. If, however, we swap envelopes, our expected final value would be $0.5 \cdot x/2 + 0.5 \cdot 2x = 1.25x$. This is larger than x and thus we should switch.

This argument seems silly though. We have no information on either of the envelopes, why in the world should we switch? Furthermore, suppose that we do, in fact, switch. Faced with this new envelope, doesn't the same argument above suggest that you ought to switch back? Wouldn't this continue forever?

What's wrong with this argument?

Discussion
The discussions usually just give an alternative solution that they claim is correct and stop there. This solution goes as follows: Call the amount in the smaller envelope x. Then the two envelopes will contain x and $2x$ respectively. You have an equal chance of having either of these envelopes in your hand. Suppose that you have the smaller envelope. If you switch you would gain $+x$. However, if instead you have the larger envelope already, this same strategy would make you lose that amount: $-x$. The expected return you would get by switching would then be $+x \cdot \frac{1}{2} - x \cdot \frac{1}{2} = 0$. In other words, there's no benefit of switching.

However, what exactly is wrong with the first approach? It seems reasonable.

In fact, there are circumstances where the first approach above *is* the correct one. It all depends on how the envelopes were prepared. Suppose instead of the procedure above, the preparer placed an amount x in an envelope and handed it to you. They then turned their back to you and flipped a fair coin. If the coin landed Heads they place $2x$ in a second envelope while if it lands Tails they place $x/2$ in that same envelope. In this case, you really *should* switch. The problem is that the envelopes in the original description were not prepared this way.

Follow-Up
OK, so now suppose that you're allowed to look at the contents of the envelope before you make your decision. Is there anything different you can do now? In order to be a bit more precise, let's specify exactly how the two envelopes are prepared:

The preparer generates a value a by sampling from a probability distribution $\psi(a)$. They then place that amount in an envelope as well as $2a$ in another envelope and hand one, chosen at random (50/50), to you.

Now suppose that you look inside the envelope and you see an amount x. You conclude that this envelope could have landed in your hands in one of two ways. Either x was in fact the value drawn from the distribution and the other envelope contains $2x$ or the underlying value drawn was in fact $x/2$ and you just happened to get the larger of the two envelopes.

What we then need to think about is the relative probability of drawing x vs. $x/2$ from the distribution ψ and how that impacts what we ultimately see in our envelope.

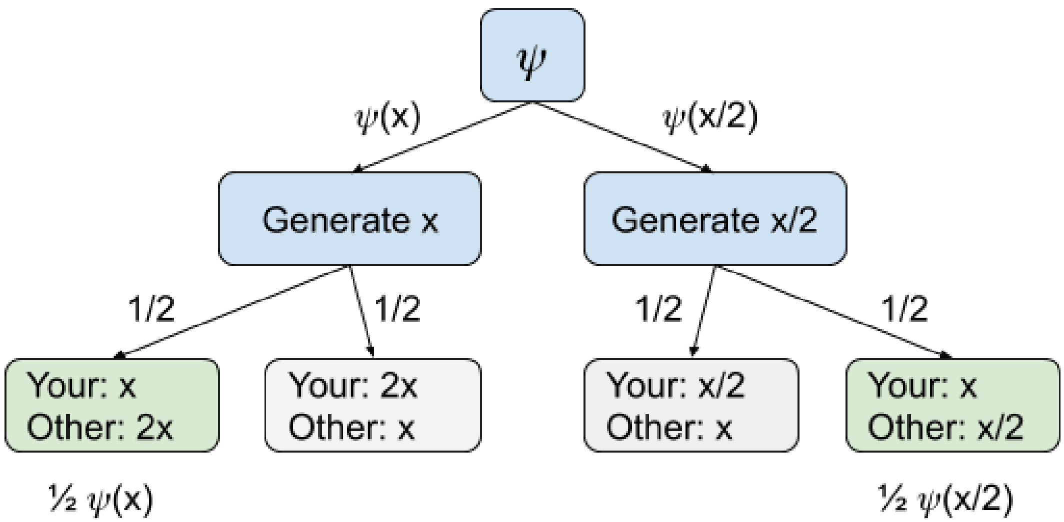

Thus, once we see the value x in our envelope we need to use Bayes' rule to compute the probabilities relevant to make our decision. Let's start with the normalization factor:

$$p(Your = x) = p(Your = x \mid Generate\ x) \cdot p(Generate\ x) + p(Your = x \mid Generate\ x/2) \cdot p(Generate\ x/2)$$
$$p(Your = x) = \tfrac{1}{2}\psi(x) + \tfrac{1}{2}\psi(x/2)$$

Now, we can compute:

$$p(Other = 2x \mid Your = x) = p(Generate\ x \mid Your = x) = p(Your = x \mid Generate\ x)\, p(Generate\ x) / p(Your = x)$$
$$p(Other = 2x \mid Your = x) = \tfrac{1}{2} \cdot \psi(x) / (\tfrac{1}{2}\psi(x) + \tfrac{1}{2}\psi(x/2))$$

A similar calculation can be done for the other direction yielding:

$$p(Other = 2x \mid Your = x) = \psi(x) / (\psi(x) + \psi(x/2))$$

$p(Other = x/2 \mid Your = x) = \psi(x/2) / (\psi(x) + \psi(x/2))$

We now have actual probabilities to ascribe to the outcomes and can compute the expected value of switching:

$$E(outcome \mid switch,\ Your = x) = 2x \cdot \psi(x) / (\psi(x) + \psi(x/2)) + \tfrac{1}{2}x \cdot \psi(x/2) / (\psi(x) + \psi(x/2))$$

We should switch as long as this is greater than x:

$$2x \cdot \psi(x) + \tfrac{1}{2}x\, \psi(x/2) > x \cdot (\psi(x) + \psi(x/2))$$

Slightly simplified, we **should switch if**

$$2\psi(x) > \psi(x/2)$$

Super Weird Example

For simplicity, let's take the possible values as being discrete (so that we don't have to worry about probability distributions etc.)

Now, I'm curious to see if we can find a probability function that satisfies this relationship regardless of what x is. In order for this to be a tractable problem, we want to further restrict the sizes of the envelopes to be powers of two (that way we can ensure that this works for the values 1 and 2, and then use this to propagate it forward to 2, 4 and then to 4, 8, and so on. We thus want:

$2 \cdot p(2) > p(1)$
$2 \cdot p(4) > p(2)$
$2 \cdot p(8) > p(4)$, …

Let's satisfy this by setting $p(2^k) = \frac{1}{2} \cdot p(2^{k-1}) \cdot \alpha$ where $\alpha > 1$. We then have

$$p(2^k) = p(1) \cdot (\alpha/2)^k$$

We can compute $p(1)$ by making sure that this sums to 1. In order for this to work, we also need $\alpha < 2$. Then:

$$p(2^k) = \frac{2-\alpha}{2}(\alpha/2)^k$$

The claim is then that as long as the first envelope is selected from this distribution and we open our randomly selected envelope and see, e.g. the value 32, we would conclude that the expected return by switching would be positive.

For definiteness, let's suppose that we pick $\alpha = 1.6$, then

$$p(2^k) = 0.2 \cdot 0.8^k$$

If you then open your envelope and see the value 2^n, you realize that this could have happened either by the smallest envelope being 2^n and you already having the smaller of the two envelopes or by the smallest envelope being 2^{n-1} and you having the larger of the two envelopes. Since the probability of generating 2^n is exactly 0.8 times as big as generating 2^{n-1}, we conclude that there's a 1/(1+0.8) = 5/9 probability that we have the larger envelope already and a 0.8/(1+0.8) = 4/9 probability that we have the smaller one. The expected outcome if we swap would then be

$$p(outcome \mid swap,\ Your = 2^n) = \tfrac{5}{9} \cdot 2^{n-1} + \tfrac{4}{9} \cdot 2^{n+1} = 2^n \cdot (\tfrac{5}{18} + \tfrac{8}{9}) = 2^n \cdot \tfrac{21}{18} > 2^n$$

In other words, whatever you see in your envelope, the expected value of the other envelope is strictly bigger.[1]

OK, now we're in a position that whatever we see in our first envelope, the decision should be to swap. So, if that's the case why even look at all? Isn't it fair to say that, before even looking in the envelope, we know that we'll swap so we may as well just swap right away?

But if that's the case, aren't we back at the very beginning? Doesn't that mean that the moment we've switched, we ought just to switch back again and so on?

What's the resolution to this?

I *think* the resolution has to do with an infinite expectation value if we don't condition on anything. It is true that once we see what's in the first envelope we should switch. However, taking the leap that this means that we don't need to even look is effectively saying that we can also average over all possible values in the first envelope. However, this involves the difference between two diverging sums:

$$E(outcome \mid swap) - E(outcome \mid don't\ swap) = \Sigma_n (\tfrac{21}{18} \cdot 2^n - 2^n) p(Your = 2^n) \sim \Sigma_n 2^n \cdot 0.8^n \sim \infty$$

In other words, because you end up subtracting two divergent series, you can't draw that conclusion in aggregate, only once you've seen what's in your envelope.

I think this is super weird. Even though whatever we see inside our envelope the decision should be to swap, we can't make the leap to skipping to look in the envelope at all. Somehow, we have to look in order to make the sums finite.

---

[1] Note: This analysis fails if you see the value 1 in your envelope since then there wouldn't be a corresponding smaller value that could have been generated to yield 1 as the bigger of the two. However, in this case it's even stronger: you *know* that the other envelope contains 2 and so the expected outcome if you swap is still larger than 1.