The Two Envelopes Problem

Rich Turner and Tom Quilter

The "**Two Envelopes Problem**", like its better known cousin, the Monty Hall problem, is seemingly paradoxical if you are not careful with your analysis. In this note we present analyses of both the careless and careful kind, providing pointers to common pitfalls for authors on this topic.

As with the Monty Hall problem, the key is to **condition** on all the facts you have when you crank the handle of Bayesian inference. Sometimes facts unexpectedly **have** to provide you with information: and you have to take them into account.

This point is well known, and most explanations of the paradox point this out. However authors then tend to make a number of mistakes. Firstly they claim the conventional analysis corresponds to using a uniform-improper distribution. This is wrong: it corresponds to a log-uniform improper distribution. Secondly they claim the root of the paradox is in the use of an un-normalisable distribution. This is merely a conjecture and the toy examples provided here suggest that improper distributions may be found which do not result in the paradox.

The "Two Envelopes Problem"'s ultimate lesson is that **inference always involves making assumptions** and any attempt to blindly use "uniformative" priors as a method for "avoiding" making assumptions can run into trouble. After all the innocuous uniform (or log-uniform) prior corresponds to a **very strong assumption about your data**.

0.1 The problem

You are playing a game for money. There are two envelopes on a table. You **know** that one contains $\$X$ and the other $\$2X$, [but you do not know which envelope is which or what the number X is]. Initially you are allowed to pick one of the envelopes, to open it, and see that it contains $\$Y$. You then have a choice: walk away with the $\$Y$ or return the envelope to the table and walk away with whatever is in the other envelope. What should you do?

0.2 The Paradox

Clearly, to decide whether we should **switch** or **stick** we want to compute the **expected return** from switching ($\langle R \rangle$) and compare it to the return from sticking (trivially, \$$Y$). Initially, the chances of choosing the envelope with the highest contents $P(\mathrm{C{=}H})$ are equal to the chances of choosing the envelope containing the smaller sum $P(\mathrm{C{=}L})$, and they are both $\frac{1}{2}$. A **careless** analysis analysis might then reason that the expected return from switching is therefore:

$$\begin{aligned}
\langle R \rangle &= \frac{1}{2}Y \times P(\mathrm{C{=}H}) + 2Y \times P(\mathrm{C{=}L}) && (1)\\
&= \frac{1}{2}Y \times \frac{1}{2} + 2Y \times \frac{1}{2} && (2)\\
&= \frac{5}{4}Y && (3)
\end{aligned}$$

This would mean we expect to gain $\frac{1}{4}Y$ from switching on average. This is paradoxical as it says "It doesn't matter which envelope you choose initially - you should always switch".

0.3 A thought experiment

So goes the usual explanation of the paradox. Where's the flaw in the analysis?

In particular, a key question we have to resolve is does **observing the contents of the first envelope provides us with useful information**. If it does not, then it will not matter whether we stick or switch [just as it didn't matter which envelope we chose to start with].

First I'm going to pursuade you that observing the contents **does** provide you with useful information. Then I'm going to show you how to use Bayes' theorem to correctly use this information to compute the optimal decision.

Imagine you make your first choice, and open up the envelope only to discover that it contains a bill - the envelope is charging you money - Y is negative. Not only is this annoying, it is also very unexpected. Who said we could lose money playing this game? If we use the above analysis for the new situation, we should now stick and not switch. So the paradox - that we should switch regardless of Y - (partly) melts away if the assumption that Y is always positive is broken. Put another way: if Y can be negative or positive then we'd have to observe Y before we decided whether switch. The moral of the story is that we should pay closer attention to our assumptions and, in particular, to what Y is telling us.

0.4 Condition on all the infomation

So we made a mistake: When computing the expected reward above we should have taken into account [technically, **condition on**] all the information we have available to us at the time. Equation 1 should therefore have read:

$$\langle R \rangle = \frac{1}{2} Y \times P(\text{C=H}|Y) + 2Y \times P(\text{C=L}|Y) \tag{4}$$

We didn't do this because we naively we thought the amount of money that the envelope contained didn't provide any useful information (ie. $P(\text{C=H}|Y) = P(\text{C=H})$, conditional independence). However the thought experiment above shows us that this is not necessarily true, and moreover whether it is or not depends on our assumptions. The cure is that we should work through the analysis stating the assumptions explicitly. The way to incorporate these assumptions explcitly is to use Bayes' theorem:

$$P(\text{C=i}|Y) = \frac{P(Y|\text{C=i})P(\text{C=i})}{P(Y)} \tag{5}$$

$$P(\text{C=i}|Y) = \frac{P(Y|\text{C=i})P(\text{C=i})}{P(Y|\text{C=H})P(\text{C=H}) + P(Y|\text{C=L})P(\text{C=L})} \tag{6}$$

We have already noted that the probability that $P(\text{C=H}) = P(\text{C=L}) = \frac{1}{2}$ and so the factors of $\frac{1}{2}$ in the denominator and numerator all cancel and we can plug the result into the formula for the expected reward:

$$\langle R \rangle = \frac{\frac{1}{2} Y P(Y|\text{C=H}) + 2Y P(Y|\text{C=L})}{P(Y|\text{C=H}) + P(Y|\text{C=L})} \tag{7}$$

$$= \frac{1}{2} Y \frac{1 + 4\gamma(X)}{1 + \gamma(X)} \tag{8}$$

The expected reward therefore depends on $\gamma = \frac{P(Y|\text{C=L})}{P(Y|\text{C=H})}$ and we are now forced to state our beliefs about the two distributions:$P(Y|\text{C=L})$ and $P(Y|\text{C=H})$ to calculate the expected reward.

0.5 Uniform beliefs

At first sight a sensible distribution - in the sense that it is non-commital - is a uniform prior between two limits, Y_{min} and Y_{max} say:

$$
\begin{aligned}
p(Y|\text{C=L}) &= \frac{1}{Y_{max} - Y_{min}} &(9)\\
&= \frac{1}{Z} &(10)\\
p(Y|\text{C=H}) &= \frac{1}{2Y_{max} - 2Y_{min}} &(11)\\
&= \frac{1}{2Z} &(12)
\end{aligned}
$$

For the sake of clarity let's think about the situation where Y is always greater than zero ($Y_{min} = 0$). What are the possible values for γ?

Well, if $0 < Y < Y_{min}$ then $\gamma = \frac{p(Y|\text{C=L})dY}{p(Y|\text{C=H})dY} = 1/2$, so: $\langle R \rangle = 3Y/2$ and we should always switch. This makes sense - it is twice as likely that the envelope contains the lower amount in this case, so switching should yield $\frac{2}{3} \times 2 + \frac{1}{3} \times \frac{1}{2}$, which it does.

However, if $Y_{min} < Y < Y_{max}$ then $\gamma = 0$ and $\langle R \rangle = \frac{1}{2}$ - we **must** have chosen the envelope containing the larger amount, the expected reward from switching is therefore \$ $\frac{1}{2}$ and so we should stick.

There is no paradox.

0.6 Improper uniform beliefs

We might now say, "well, I'm very unsure about what value Y would take and therefore I'll make Y_{max} extremely large, and take the limit $Y_{max} \to \infty$"[1]. This turns out to be a **terrible assumption** for any practical situation. As we take this limit $p(Y|\text{C=L})$ and $p(Y|\text{C=H})$ become very ill matched to **any possible real data** in such a way that this causes problems. All data we observe will lie beneath Y_{max} (as it's tending to infinity). However, in this regime we "believe" it is twice as likely that the envelope contains the lower amount, and so we switch whatever we see. Put another way, in order to stick we would have to observe a Y which is larger than Y_{max}, but we'll never encounter an envelope containing more than infinity dollars [2] so we'll always switch.

Although we tried to make the most non-commital assumptions possible - we ended up making very definite, absurd assumptions which

[1]In this limit the uniform distribution is un-normalisable and such distributions are called "improper"

[2]although there are numbers greater than infinity, it's tough to write legal cheques for these amounts

makes sensible inference impossible. [3]

Before we move on, it's worth mentioning that it might have been tempting to set γ equal to one in the above case. After all, in the limit both $p(Y|\text{C=L})$ and $p(Y|\text{C=H})$ are (improper) uniform distributions between zero and infinity. We would then have $\gamma = 1$ and recover the old result $\langle R\rangle = 5Y/4$. However, setting the two densities equal is not consistent with the assumption that one envelope contains an amount twice that in the other. This assumption demands one density be half of the other (and to have twice the range). So the original analysis does not correspond to assuming a uniform improper distribution (as some authors claim).

To conclude, it's impossible to use this particular improper distributions without being careful (or else we introduce a contradiction with our other assumptions) and even if we are careful they correspond to absurd assumptions (where we switch unless Y is greater than infinity).

0.7 The log-uniform beliefs

If the $\gamma = 1$ case does not correspond to pair of uniform distributions $p(Y|\text{C=H})$ and $p(Y|\text{C=L})$, does it correspond to another choice? The answer, somewhat suprisingly, is yes. To see this we need to do some more work to simplify the form of γ. We can do this by realising that specifying $p(Y|\text{C=L})$ immediately specifies $p(Y|\text{C=H})$ (due to the doubling constraint). We can derive a relation which relates the two via manipulation which is easy to skrew up (as, for example, many authors on this paradox do - even in publications). Here's one way to make sure you get the right result: We want to sum up all the probability at Y' in the old distribution which gets mapped to $Y = 2Y'$ in the new distribution. So we write down:

$$
\begin{aligned}
p(Y|\text{C=H}) &= \int_0^\infty p(Y'|\text{C=L})\delta(2Y'-Y)dY' && (13)\\
&= \int_{-\infty}^\infty p(U/2|\text{C=L})\delta(U-Y)dU/2 && (14)\\
&= p(Y/2|\text{C=L})/2 && (15)
\end{aligned}
$$

This makes intuitive sense: A chunk of probability in the old distribution $p(Y'|\text{C=L})dY'$ is mapped to a region centred at $Y = 2Y'$ but it is smeared out to extend over a length $2dY'$ so the density at $p(Y|\text{C=L})$ must be half that of $p(Y/2|\text{C=L})$.

[3] Note, however, that as we take the limit the paradox (that we should always switch regardless of Y) does not rear it's ugly head.

If γ equals one then this means: $p(Y|\text{C=L}) = p(Y|\text{C=H}) = p(Y/2|C = L)/2$ for all Y, which holds when $p(Y|\text{C=L}) = 1/Y$. This is a uniform distribution on $\log Y$ and says, "I have know idea about the **scale** of Y. I think Y is as likely to be between 1 and 10 as it is to be between 10 and 100". This seems more sensible than the uniform prior before - for one thing, it decreases with Y.

However, this distribution is not normalisable (without introducing a lower and upper cut-off and then γ is not 1 for all Y). Limiting arguments fail for the same reason they did in the uniform case: in the region where both the distributions are non-zero, you should always switch. So taking Y_{max} to infinity pushes the regime where you should stick to a place where no data will land. Again, the limit corresponding to the improper distribution corresponds to a terrible set of assumptions.

Does using an improper distribution **always** lead to the paradox? Perhaps, but we haven't proved it. More over the intuition from the above examples is that it wasn't the improper nature of the distribution which lead to the paradox (as some authors glibly state) - it was the fact that they corresponded to stupid assumptions. I could imagine there are distributions which corresponded to reasonable assumptions in the limit where they become un-normalisable. [4]

0.8 A sensible approach

Is there a sensible approach? In a word yes. However, it will depend on what assumptions **you** want to make. In a way the previous discussion is a red-herring: you might be quite happy to use a uniform distribution, or a log-uniform distribution - and everything will be ok so long as you think carefully about the choice for the ranges (am I playing with a friend for his poket money, or am I through to the final stages of a game-show called "Who wants to be a billionaire"?). [5] One final alternative (which does not have hard cut-offs) is to place an exponential distribution over Y:

$$P(Y|\text{C=L}) \quad = \quad \lambda \exp(-\lambda Y) \tag{16}$$

$1/\lambda$ is the characteristic length scale over which the density decays. From eqn. 8 we should switch when:

[4] Another point noted by authors is that the improper distribution above has an infinite mean value - this may correspond to another undesireable assumption too.

[5] One alternative is to place a prior over the ranges and integrate them out: $P(Y|\text{C=L}) = \int P(Y|\text{C=L}, Y_{min}, Y_{max}) P(Y_{min}, Y_{max}) dY_{min} dY_{max}$

$$\begin{aligned} \frac{1}{2} &\leq \gamma(Y) && (17) \\ &= 2\exp(-\lambda Y/2) && (18) \end{aligned}$$

So we switch when $Y \leq \frac{2}{\lambda}\ln 4 \simeq \frac{3}{\lambda}$, which is proportional to the characteristic length scale, as expected.