

Books such as Nicholas Taleb's [Fooled by Randomness](#) and the psychological literature on our mental foibles such as gambler's fallacy warn us to beware randomness. Well and good, but randomness actually is one of the most domesticated kinds of uncertainty. In fact, it is one form of uncertainty we can and do exploit.

One obvious way randomness can be exploited is in designing scientific experiments. To experimentally compare, say, two different fertilizers for use in growing broad beans, an ideal would be to somehow ensure that the bean seedlings exposed to one fertilizer were identical in all ways to the bean seedlings exposed to the other fertilizer. That isn't possible in any practical sense. Instead, we can randomly assign each seedling to receive one or the other fertilizer. We won't end up with two identical groups of seedlings, but the differences between those groups will have occurred by chance. If their subsequent growth-rates differ by more than we would reasonably expect by chance alone, then we can infer that one fertilizer is likely to have been more effective than the other.

Another commonplace exploitation of randomness is random sampling, which is used in all sorts of applications from quality-control engineering to marketing surveys. By randomly sampling a specific percentage of manufactured components coming off the production line, a quality-control analyst can decide whether a batch should be scrapped or not. By randomly sampling from a population of consumers, a marketing researcher can estimate the percentage of that population who prefer a particular brand of a consumer item, and also calculate how likely that estimate is to be within 1% of the true percentage at the time.

There is a less well-known use for randomness, one that in some respects is quite counter-intuitive. We can exploit randomness to improve our chances of making the right decision. The story begins with Tom Cover's 1987 chapter which presents what Dov Samet and his co-authors recognized in their [2002 paper](#) as a solution to a switching decision that has been at the root of various puzzles and paradoxes.

Probably the most famous of these is the ["two envelope" problem](#). You're a contestant in a game show, and the host offers you a choice between two envelopes, each containing a cheque of a specific value. The host explains that one of the cheques is for a greater amount than the other, and offers you the opportunity to toss a fair coin to select one envelope to open. After that, she says, you may choose either to retain the envelope you've selected or exchange it for the other. You toss the coin, open the selected envelope, and see the value of the cheque therein. Of course, you don't know the value of the other cheque, so regardless of which way you choose, you have a probability of $\frac{1}{2}$ of ending up with the larger cheque. There's an appealing but fallacious argument that says you should switch, but we're not going to go into that here.

Cover presents a remarkable decisional algorithm whereby you can make that probability exceed $\frac{1}{2}$.

1. Having chosen your envelope via the coin-toss, use a random number generator to provide you with a number anywhere between zero and some value you know to be greater than the largest cheque's value.
2. If this number is larger than the value of the cheque you've seen, exchange envelopes.
3. If not, keep the envelope you've been given.

Here's a "reasonable person's proof" that this works (for more rigorous and general proofs, see [Robert Snapp's 2005](#) treatment or [Samet et al., 2002](#)). I'll take the role of the game-show

contestant and you can be the host. Suppose $\$X_1$ and $\$X_2$ are the amounts in the two envelopes. You have provided the envelopes and so you know that X_1 , say, is larger than X_2 . You've also told me that these amounts are less than \$100 (the specific range doesn't matter). You toss a fair coin, and if it lands Heads you give me the envelope containing X_1 whereas if it lands Tails you give me the one containing X_2 . I open my envelope and see the amount there. Let's call my amount Y . All I know at this point is that the probability that $Y = X_1$ is $\frac{1}{2}$ and so is the probability that $Y = X_2$.

I now use a random number generator to produce a number between 0 and 100. Let's call this number Z . Cover's algorithm says I should switch envelopes if Z is larger than Y and I should retain my envelope if Z is less than or equal to Y . The claim is that my chance of ending up with the envelope containing X_1 is greater than $\frac{1}{2}$.

As the picture below illustrates, the probability that my randomly generated Z has landed at X_2 or below is $X_2/100$, and the probability that Z has landed at X_1 or below is $X_1/100$. Likewise, the probability that Z has exceeded X_2 is $1 - X_2/100$, and the probability that Z has exceeded X_1 is $1 - X_1/100$.

0----- X_2 ----- X_1 -----100

The proof now needs four steps to complete it:

1. If $Y = X_1$ then I'll make the right decision if I decide to keep my envelope, i.e., if Y is less than or equal to X_1 , and my probability of doing so is $X_1/100$.
2. If $Y = X_2$ then I'll make the right decision if I decide to exchange my envelope, i.e., if Y is greater than X_2 , and my probability of doing so is $1 - X_2/100$.
3. The probability that $Y = X_1$ is $\frac{1}{2}$ and the probability that $Y = X_2$ also is $\frac{1}{2}$. So my total probability of ending up with the envelope containing X_1 is $\frac{1}{2}$ of $X_1/100$, which is $X_1/200$, plus $\frac{1}{2}$ of $1 - X_2/100$, which is $\frac{1}{2} - X_2/200$. That works out to $\frac{1}{2} + X_1/200 - X_2/200$.
4. But X_1 is larger than X_2 , so $X_1/200 - X_2/200$ must be larger than 0. Therefore, $\frac{1}{2} + X_1/200 - X_2/200$ is larger than $\frac{1}{2}$.

Fine, you might say, but could this party trick ever help us in a real-world decision? Yes, it could. Suppose you're the director of a medical clinic with a tight budget in a desperate race against time to mount a campaign against a disease outbreak in your region. You have two treatments available to you but the research literature doesn't tell you which one is better than the other. You have time and resources to test only *one* of those treatments before deciding which one to adopt for your campaign.

Toss a fair coin, letting it decide which treatment you test. The resulting cure-rate from the chosen treatment will be some number, Y , between 0% and 100%. The structure of your decisional situation now is identical to the two-envelope setup described above. Use a random number generator to generate a number, Z , between 0 and 100. If Z is less than or equal to Y use your chosen treatment for your campaign. If Z is greater than Y use the other treatment instead. Your chance of having chosen the treatment that would have yielded the higher cure-rate under your test conditions will be larger than $\frac{1}{2}$ and you'll be able to defend your decision if you're held accountable to any constituency or stakeholders.

In fact, there are ways whereby you may be able to do even better than this in a real-world situation. One is by shortening the range, if you know that the cure-rate is not going to exceed

some limit, say L , below 100%. The reason this would help is because $X_1/2L - X_2/2L$ will be greater than $X_1/200 - X_2/200$. The highest it can be is $1 - X_2/2X_1$. Another way, as Snapp (2005) points out, is by knowing the probability distribution generating X_1 and X_2 . Knowing that distribution boosts your probability of being correct to $\frac{3}{4}$.

However, before we rush off to use Cover's algorithm for all kinds of decisions, let's consider its limitations. Returning to the disease outbreak scenario, suppose you have good reasons to suspect that one treatment (T_a , say) is better than the other (T_b). You could just go with T_a and defend your decision by pointing out that, according to your evidence the probability that T_a actually is better than T_b is greater than $\frac{1}{2}$. Let's denote this probability by P .

A reasonable question is whether you could do better than P by using Cover's algorithm. Here's my claim:

- If you test T_a or T_b and use the Cover algorithm to decide whether to use it for your campaign or switch to the other treatment, your probability of having chosen the treatment that would have given you the best test-result cure rate will converge to the Cover algorithm's probability of a correct choice. This may or may not be greater than P (remember, P is greater than $\frac{1}{2}$).

This time, let X_1 denote the higher cure rate and X_2 denote the lower cure-rate you would have got, depending on whether the treatment you tested was the better or the worse.

1. If the cure rate for T_a is X_1 then you'll make the right decision if you decide to use T_a , i.e., if Y is less than or equal to X_1 , and your probability of doing so is $X_1/100$.
2. If the cure rate for T_a is X_2 then you'll make the right decision if you decide to use T_b , i.e., if Y is greater than X_2 , and your probability of doing so is $1 - X_2/100$.
3. We began by supposing the probability that the cure rate for T_a is X_1 is P , which is greater than $\frac{1}{2}$. The probability that the cure rate for T_a is X_2 is $1 - P$, which is less than $\frac{1}{2}$. So your total probability of ending up with the treatment whose cure rate is X_1 is

$$P \cdot X_1/100 + (1 - P) \cdot (1 - X_2/100).$$

The question we want to address is when this probability is greater than P , i.e.,

$$P \cdot X_1/100 + (1 - P) \cdot (1 - X_2/100) > P.$$

It turns out that a rearrangement of this inequality gives us a clue.

4. First, we subtract $P \cdot X_1/100$ from both sides to get $(1 - P) \cdot (1 - X_2/100) > P - P \cdot X_1/100$.
5. Now, we divide both sides of this inequality by $1 - P$ to get $(1 - X_2/100)/P > P \cdot (1 - X_1/100)/(1 - P)$, and then divide both sides by $(1 - X_1/100)$ to get $(1 - X_2/100)/(1 - X_1/100) > P/(1 - P)$.

We can now see that the values of X_2 and X_1 have to make the odds of the Cover algorithm larger than the odds resulting from P . If $P = .6$, say, then $P/(1 - P) = .6/.4 = 1.5$. Thus, for example, if $X_2 = 40\%$ and $X_1 = 70\%$ then $(1 - X_2/100)/(1 - X_1/100) = .6/.3 = 2.0$ and the Cover algorithm will improve your chances of making the right choice. However, if $X_2 = 40\%$ and $X_1 = 60\%$ then the algorithm offers no improvement on P and if we increase X_2 above 40% the algorithm will return a lower probability than P . So, if you already have strong evidence that one alternative is better than the other then don't bother using the Cover algorithm.

Nevertheless, by exploiting randomness we've ended up with a decisional guide that can apply to real-world situations. Faced with being able to test only one of two alternatives, if you are undecided about which one is superior but can only test one alternative, test one of them and use Cover's algorithm to decide which to adopt. You'll end up with a higher probability of making the right decision than tossing a coin.