MY TAKE ON THE MONTE HALL PROBLEM

MARK FISHER

This problem has been beaten to death. I won't let that stop me. Consider this overkill.

1. Introduction

The problem. There are three doors. A prize is located behind one and there is nothing behind the other two. The contestant (C) chooses a door after which the game-show host Monte Hall (MH) opens a different door that reveals no prize behind it. C is then allowed to switch the choice to the remaining unopened door. Should C stay or switch?

Just to be clear. There is an equal probability that the prize is located behind each of the three doors before the choice is made. In addition, there are three rules that MH follows: He doesn't open the chosen door; he doesn't open a door that has the prize is behind it; and if the prize is behind the chosen door then he is equally likely to open either of the other two doors.

2. Preliminary thoughts

Repeated play. Imagine two contestants, Stay and Switch, playing at the same time, again and again. Switch always chooses the same door that Stay chooses, but Stay always stays and Switch always switches. Stay wins whenever the prize is behind the chosen door. Switch wins whenever it is not.

Restate the problem. C picks a door and then is immediately allowed to switch to *both* of the other doors, thereby winning the prize if it is behind either of those two doors. Should C stay or switch?

Modification 1. Before C can decide, MH reminds C the prize cannot be behind both of the other two doors. Rather it is either behind one, the other, or neither. With this reminder, should C stay or switch?

Modification 2. Before C can decide and after MH reminds C the prize cannot be behind both of the other two doors, MH opens a door with no prize in order to prove his point. With this proof, should C stay or switch?

The same but different? Suppose instead C switches to the two other doors before MH provides his reminder or his proof. Undeterred, MH proceeds to make his point by opening one of those two doors that has no prize. What effect does that have on the probability that C wins the prize?

Date: February 24, 2026 @ 12:08. *Previous version*: November 2021. *Filename*: `"Monte Hall"`.

Are we done? What else is there to say?

Whether to stay or switch is a *decision* to be made based on what is known and what is learned about an unknown state of the world (i.e., where the prize is located). The rest of this note casts the problem in Bayesian terms, allowing us to solve not only the problem as stated but also related problems.

3. Framework

The analysis is formally the same no matter which door C chooses, so assume C chooses door 1. There are two categorical random variables: the door Behind which the prize is located (the unknown state of the world),

$$\mathsf{B} \in \{1, 2, 3\}, \tag{3.1}$$

and the door that is Opened by MH (the observation),

$$\mathsf{O} \in \{2, 3\}. \tag{3.2}$$

The joint sample space contains six points:

$$(\mathsf{B}, \mathsf{O}) \in \{1, 2, 3\} \times \{2, 3\} = \{(1, 2), (1, 3), (2, 2), (2, 3), (3, 2), (3, 3)\}. \tag{3.3}$$

The joint probability distribution $p(\mathsf{B}, \mathsf{O})$ plays a central role. There are two ways it can be factored into the product of a conditional distribution and a marginal distribution:

$$\begin{aligned} p(\mathsf{B}, \mathsf{O}) &= p(\mathsf{O}|\mathsf{B})\, p(\mathsf{B}) \\ &= p(\mathsf{B}|\mathsf{O})\, p(\mathsf{O}). \end{aligned} \tag{3.4}$$

We will rely on both factorizations. First we will use what we have been told to flesh out $p(\mathsf{B})$ and $p(\mathsf{O}|\mathsf{B})$, which we will then use to construct $p(\mathsf{B}, \mathsf{O})$, in effect filling out a contingency table, row by row. Then we will obtain $p(\mathsf{B}|\mathsf{O})$ from $p(\mathsf{B}, \mathsf{O})$, extracting what we need from that table, column by column.

4. Application

For our purposes, the marginal distribution $p(\mathsf{B})$ plays the role of the *prior distribution*. As such, it summarizes what is known about the location of the prize *before* C makes a choice and (importantly) *before* MH reveals anything:

$$p(\mathsf{B} = 1) = 1/3 \tag{4.1a}$$

$$p(\mathsf{B} = 2) = 1/3 \tag{4.1b}$$

$$p(\mathsf{B} = 3) = 1/3. \tag{4.1c}$$

The conditional distribution $p(\mathsf{O}|\mathsf{B})$ summarizes how MH's observable actions depend on the location of the prize:

$$p(\mathsf{O} = 2|\mathsf{B} = 1) = 1/2 \qquad p(\mathsf{O} = 3|\mathsf{B} = 1) = 1/2 \tag{4.2a}$$

$$p(\mathsf{O} = 2|\mathsf{B} = 2) = 0 \qquad p(\mathsf{O} = 3|\mathsf{B} = 2) = 1 \tag{4.2b}$$

$$p(\mathsf{O} = 2|\mathsf{B} = 3) = 1 \qquad p(\mathsf{O} = 3|\mathsf{B} = 3) = 0. \tag{4.2c}$$

Note $p(\mathsf{O} = 2|\mathsf{B}) + p(\mathsf{O} = 3|\mathsf{B}) = 1$ for all $\mathsf{B} \in \{1, 2, 3\}$.

We can compute the joint distribution by using the first line of (3.4) as a guide to combining (4.1) and (4.2). For example,

$$p(\mathsf{B}=1,\mathsf{O}=2) = p(\mathsf{O}=2|\mathsf{B}=1)\,p(\mathsf{B}=1) = (1/2)\,(1/3) = 1/6. \tag{4.3}$$

This joint distribution can be represented as a contingency table:

	$\mathsf{O}=2$	$\mathsf{O}=3$	$p(\mathsf{B})$
$\mathsf{B}=1$	1/6	1/6	1/3
$\mathsf{B}=2$	0	1/3	1/3
$\mathsf{B}=3$	1/3	0	1/3
$p(\mathsf{O})$	1/2	1/2	1

(4.4)

The six numbers in the central rectangle in the table are the joint probabilities. The prior probabilities appear as the marginal probabilities on the right. We can easily recover $p(\mathsf{O}|\mathsf{B})$ from the table by dividing each B-row of the table by its corresponding marginal probability. In doing so, we would in effect be using a formula obtained by rearranging the first line of (3.4) as follows:

$$p(\mathsf{O}|\mathsf{B}) = \frac{p(\mathsf{B},\mathsf{O})}{p(\mathsf{B})}. \tag{4.5}$$

The second line in (3.4) tells us that we can do the same kind of extraction in the other direction, column by column instead of row by row, producing[1]

$$p(\mathsf{B}|\mathsf{O}) = \frac{p(\mathsf{B},\mathsf{O})}{p(\mathsf{O})}. \tag{4.6}$$

For our purposes, the conditional distribution $p(\mathsf{B}|\mathsf{O})$ plays the role of the *posterior distribution* because it summarizes what we know about the location of the prize *after* we have observed what MH does, combining what we already knew with what we have just observed.

As mentioned earlier, $p(\mathsf{B}|\mathsf{O})$ can be recovered from the columns of the contingency table by dividing each O-column by its corresponding marginal probability:

$$p(\mathsf{B}=1|\mathsf{O}=2) = 1/3 \qquad p(\mathsf{B}=1|\mathsf{O}=3) = 1/3 \tag{4.7a}$$

$$p(\mathsf{B}=2|\mathsf{O}=2) = 0 \qquad p(\mathsf{B}=2|\mathsf{O}=3) = 2/3 \tag{4.7b}$$

$$p(\mathsf{B}=3|\mathsf{O}=2) = 2/3 \qquad p(\mathsf{B}=3|\mathsf{O}=3) = 0. \tag{4.7c}$$

Note $p(\mathsf{B}=1|\mathsf{O}) + p(\mathsf{B}=2|\mathsf{O}) + p(\mathsf{B}=3|\mathsf{O}) = 1$ for all $\mathsf{O} \in \{2,3\}$.

Once we observe which door MH opened, only one of the two distributions is applicable: either $\mathsf{O}=2$ or $\mathsf{O}=3$, but not both.[2] If MH opened door 2 then the probability of the prize being behind door 3 is 2/3. On the other hand, if MH opened door 3 then the probability of the prize being behind door 2 is 2/3.

Solution. Regardless of which door is opened, if C switches then the probability of winning the prize increases from 1/3 to 2/3. Therefore, the *optimal decision* is to switch, no matter which door MH opens.

[1]This formula is often written as $p(\mathsf{B}|\mathsf{O}) = \frac{p(\mathsf{O}|\mathsf{B})\,p(\mathsf{B})}{p(\mathsf{O})}$ and called Bayes rule.

[2]Once we observe which door MH opened, O is fixed in $p(\mathsf{O}|\mathsf{B})$ while B can vary. In this situation $p(\mathsf{O}|\mathsf{B})$ is called the *likelihood* for B.

5. Signaling

No signaling. Under the rules established in Section 1 for MH, the observation (i.e., which door was opened) carried no information about whether or not the prize was behind door 1:[3]

$$p(\mathsf{B}=1|\mathsf{O}=2) = p(\mathsf{B}=1|\mathsf{O}=3) = p(\mathsf{B}=1) = 1/3. \tag{5.1}$$

The observation only carried information about the distribution of the remaining probability (i.e., 2/3) between doors 2 and 3.

We can imbue the observation with additional information by changing the rules to allow MH to signal when the prize is behind door 1.

Yes, signaling. Suppose MH always opens door 2 when the prize is behind door 1 (and C knows this).[4] This signaling doesn't change the prior distribution for location of the prize, $p(\mathsf{B})$, but it does change the way MH behaves based on that location. In other words, it changes $p(\mathsf{O}|\mathsf{B})$.

In this case,

$$p(\mathsf{O}=2|\mathsf{B}=1) = 1 \qquad p(\mathsf{O}=3|\mathsf{B}=1) = 0 \tag{5.2a}$$
$$p(\mathsf{O}=2|\mathsf{B}=2) = 0 \qquad p(\mathsf{O}=3|\mathsf{B}=2) = 1 \tag{5.2b}$$
$$p(\mathsf{O}=2|\mathsf{B}=3) = 1 \qquad p(\mathsf{O}=3|\mathsf{B}=3) = 0 \tag{5.2c}$$

and

	$\mathsf{O}=2$	$\mathsf{O}=3$	$p(\mathsf{B})$
$\mathsf{B}=1$	$1/3$	0	$1/3$
$\mathsf{B}=2$	0	$1/3$	$1/3$
$\mathsf{B}=3$	$1/3$	0	$1/3$
$p(\mathsf{O})$	$2/3$	$1/3$	1

(5.3)

and

$$p(\mathsf{B}=1|\mathsf{O}=2) = 1/2 \qquad p(\mathsf{B}=1|\mathsf{O}=3) = 0 \tag{5.4a}$$
$$p(\mathsf{B}=2|\mathsf{O}=2) = 0 \qquad p(\mathsf{B}=2|\mathsf{O}=3) = 1 \tag{5.4b}$$
$$p(\mathsf{B}=3|\mathsf{O}=2) = 1/2 \qquad p(\mathsf{B}=3|\mathsf{O}=3) = 0. \tag{5.4c}$$

Conclusion. According to (5.4), there is no advantage (or disadvantage) to switching if MH opens door 2. By contrast, if MH opens door 3 then switching guarantees the prize. MH's action clearly conveys information about the probability the prize is behind door 1:

$$\underbrace{p(\mathsf{B}=1|\mathsf{O}=2)}_{1/2} > \underbrace{p(\mathsf{B}=1)}_{1/3} > \underbrace{p(\mathsf{B}=1|\mathsf{O}=3)}_{0}. \tag{5.5}$$

[3]Recall the restatement of the problem given in Section 2. When MH opened one of the two other doors, he was supposedly proving the prize cannot be behind both of them. Since this was already known, his action was uninformative in that regard.

[4]Signaling needs to take into account which door C chooses. For example, if C were to choose door 2 and the prize were behind it, then MH would always open door 3; and if C were to choose door 3 and the prize were behind it, then MH would always open door 1.

Repeated play. In repeated play Switch still wins when the prize is behind door 2 or door 3, which still happens 2/3 of the time. That unconditional probability is determined by $p(\mathsf{B})$, which has not changed. The advantage that signaling creates is conditional: We learn from the observation when the odds are in favor of switching and when they are not.

6. More doors

Suppose there are $n \geq 3$ doors and the contestant chooses door 1. In this case, $\mathsf{B} \in \{1, \ldots, n\}$. The prior probabilities are $p(\mathsf{B}) = 1/n$ for all B and there is no signaling.

Open only one. MH opens *one* door. Let $\mathsf{O} \in \{2, \ldots, n\}$ denote the one door that MH opens. In this case,

$$p(\mathsf{O} = j|\mathsf{B} = i) = \begin{cases} 1/(n-1) & i = 1 \\ 0 & j = i \geq 2 \\ 1/(n-2) & \text{otherwise} \end{cases} . \tag{6.1}$$

As a result, the posterior probabilities are

$$p(\mathsf{B} = i|\mathsf{O} = j) = \begin{cases} 1/n & i = 1 \\ 0 & j = i \geq 2 \\ \frac{n-1}{n\,(n-2)} & \text{otherwise} \end{cases} . \tag{6.2}$$

Switch does not alway win when Stay loses: Sometimes they both lose. Nevertheless, Switch has a better chance of winning than Stay since $\frac{n-1}{n\,(n-2)} > 1/n$. Switch's advantage decreases as n increases, but it never completely goes away.

Open all but one. Now suppose MH opens all doors but one. Let $\mathsf{U} \in \{2, \ldots, n\}$ denote the one door that MH leaves $\mathsf{Unopened}$. In this case,

$$p(\mathsf{U} = j|\mathsf{B} = i) = \begin{cases} 1/(n-1) & i = 1 \\ 1 & j = i \geq 2 \\ 0 & \text{otherwise} \end{cases} . \tag{6.3}$$

As a result, the posterior probabilities are

$$p(\mathsf{B} = i|\mathsf{U} = j) = \begin{cases} 1/n & i = 1 \\ (n-1)/n & j = i \geq 2 \\ 0 & \text{otherwise} \end{cases} . \tag{6.4}$$

We are back to the situation where Switch always wins when Stay loses. As a result, Switch has an advantage that increases as n increases: $(n-1)/n > 1/n$.

Restate the all-but-one problem. C picks a door and then is immediately allowed to switch to *all* of the other doors, thereby winning the prize if it is behind any of those $n-1$ doors. Should C stay or switch to the other $n-1$ doors?

Email address: `mark@markfisher.net`
URL: `http://www.markfisher.net`