

Anyone who considers arithmetical  
methods of producing random  
digits is, of course, in a state of sin.

---

John von Neumann

# Monte Carlo Simulation

You are the new kid on the block and it is time to choose teams at the rec center. You think you are pretty good so you say, “I’m a 90% free throw shooter.” This is quite impressive. Someone hands you a basketball and says, “Prove it.”

You shoot 100 free throws, but it does not go as well as you hoped. You make only 75. Someone says, “You are not a 90% free throw shooter.” You insist, however, that you really are. “It was just bad luck. Honestly,” you say, “I really am a 90% free throw shooter.”

The question is: Should we believe you? Anyone who has ever shot free throws knows there is luck involved. We would not expect you to make 9 out of every 10 shots, like clockwork. So, it could be that by chance, you missed a few more than expected. Another way to ask the question is: Can randomness explain this poor outcome? Or, in yet other words, how uncommon is missing this many free throws for a 90% shooter?

We are not having this conversation if you had made 89 or even 88 out of a 100. Then it would be easy to believe you are actually a 90% shooter and it was just bad luck. But how do we handle the fact that you missed 15 more than expected? That seems like a lot, but how rare is that?

There is a way to answer this question analytically, that is, with mathematics. We will not go that route. Instead, we will use the method of simulation.

Monte Carlo simulation simply means the repeated running of a chance process and then direct examination of the results. It can be used in frontier research work, but we will use it just like we used numerical methods to solve optimization problems—simulation enables us to understand complicated concepts without advanced mathematics.

Monte Carlo simulation is based on brute force—repeat the chance process and examine the results. It requires no imagination or mathematics at all. It will be our go-to method for understanding randomness and answering questions like “Do we believe you are a 90% free throw shooter if you make only 75 out of a 100?”.

## Gauss and Two Approaches

Carl Friedrich Gauss (1777-1855) was perhaps the greatest mathematician of all time. Before the euro, Germany’s ten Deutsche Mark note featured him along with a graph of the normal curve (which he made famous and is called the Gaussian distribution). Look carefully in Figure 3.1 and you can see that it even displays the equation of the normal distribution.



Figure 3.1: German currency featuring Gauss.

There is a story, probably apocryphal, of how he amazed his kindergarten teacher. Apparently, the children were especially unruly one day, so the teacher assigned a dreary problem as punishment. He told them to add all the numbers from 1 to 1,000. This starts easily, but gets tedious and painful pretty quickly.  $1 + 2 = 3$ ,  $3 + 3 = 6$ ,  $4 + 6 = 10$ , and  $5 + 10 = 15$ . It will take a long time to get to 1,000.

Gauss waited a minute, then stood up and announced the answer, 500,500. The stunned teacher asked him where he got that number, which is correct, and Gauss said he noticed a pattern. Remember, he is five years old.

If you make a list of the numbers, then create a second list, but flipped, the pairs always add up to 1,001. 1 goes with 1,000, 2 with 999, 3 with 998, and so on until the end when 998 is with 3, 999 with 2, and 1 with 1,000.

The rest is easy (well, maybe not for the usual five-year old, but this is Gauss). Multiply 1,001 by 1,000 (since there are 1,000 pairs) and divide by 2 to get 500,500. As they say, QED.

This is clever, remarkable, and beautiful. It is like Michelangelo and the Sistine Chapel. Is Monte Carlo simulation like this? No.

Monte Carlo simulation is a different approach to problems which uses little creativity or subtlety. It is a direct attack on a question.

Monte Carlo simulation is like solving the teacher's tedious problem by using a spreadsheet to add the numbers.

***STEP*** Make a list from 1 to 1000 in cells A1:A1000 (using fill down, of course) and then, in cell B1, enter the formula `=SUM(A1:A1000)`.

Excel displays 500,500. This is nowhere as magnificent as what Gauss did, but it does give the answer.

Monte Carlo simulation was developed during World War II by physicists working on the Manhattan Project. Nicholas Metropolis coined the term because his colleague, Stanislaw Ulam, was an avid poker player. They were simulating how radiation propagates and incorporating randomness. The connection to chance and gambling is why Metropolis named the method after the famous Monte Carlo casino in Monaco.

So Monte Carlo simulation, or simulation for short, is an alternative to the analytical approach. Instead of using equations and algebraic manipulations, simulation uses computers to repeat the chance process many times and then directly observes the outcomes.

You can think of simulation as the much younger sibling of analytical methods. Let's apply it to free throw shooting to show how it works.

## Are You Really a 90% Shooter?

Our solution strategy will be simulation, but make no mistake, Gauss would not have needed simulation. He would have immediately rejected your claim. He knows a formula that can be used,  $\sqrt{n} \sigma$ , which answers the question quickly. The formula is the product of sophisticated mathematics and can be called beautiful, but most people find it extremely difficult to understand and cannot use it to answer the question.

All Monte Carlo simulations use a random number generator(RNG). Excel's RNG function is RAND(). This draws uniformly distributed random numbers in the interval from zero to one.

**STEP** Insert a sheet in your Excel workbook and, in cell A1, enter the formula =RAND().

You see a number with several decimal places displayed that is between zero and one. The number is actually much longer.

**STEP** Widen the column and add decimal places to see this. Keep adding decimal places (widening column A as needed) until you start seeing zeroes.

As you learned when we explained Solver's false precision, most modern spreadsheets use 64-bit double-precision floating point format. If you count carefully, you will see that RAND() has a zero, then a decimal point, and then 15 decimal places with values from zero to nine. After that, they are all zero so we have reached the maximum precision. It is important to understand that our spreadsheet's random number is finite, but also that it has many more decimal digits than what was originally displayed.

**STEP** Repeatedly press F9 (you may have to hold down the *fn* key on your keyboard). F9 is the keyboard shortcut to recalculate the sheet.

The number in cell A1 changes each time you recalculate the sheet. This is the beating heart of the simulation.

The bouncing numbers show that, although RAND() is finite, it has a massive set of numbers to choose from. If it had only one decimal place, RAND() would have 10 possible numbers (from 0.0, 0.1, 0.2, and so on until 0.9). Six

decimal places would give it 1 million different numbers. Twelve gives a trillion numbers. Fifteen is a quadrillion!

So, you can think of `RAND()` as a plucking a number from a humongous box with a quadrillion numbers in it.

Full disclosure: This is not exactly right because `RAND()` is using an algorithm to produce the next number. This is why computer-generated random numbers are called *pseudorandom*, where the prefix *pseudo* means false.

To model a 90% free throw shooter, we use an IF statement.

**STEP** In cell B1, enter the formula `=IF(A1<0.9, 1, 0)`.

The IF function has three arguments (or inputs), separated by commas. The first argument is the test, the second is what happens if the test is true (or yes), and the third is what happens if it is false (or no). If the random number in cell A1 is less than 90%, then cell B1 shows a one, which means the free throw was made; otherwise, it shows a zero, which means it was missed.

Some students (usually really smart, careful ones) obsess about whether A3 should be less than ( $<$ ) or less than or equal to ( $\leq$ ) 90%. This does not matter because `RAND()` has so many random numbers available to it. The chances of drawing exactly 0.9000000000000000 are ridiculously small.

When the IF function evaluates to 1, the free throw is made and when it is 0, it is missed. This is a binomial random variable since it can only take on two values, 0 or 1.

We do not need to actually see the random number generated so we can embed `RAND()` directly in the IF statement.

**STEP** In cell B2, enter the formula `=IF(RAND() $<$ 0.9,1,0)`. Fill down this formula to cell B100. Press F9 a few times to see the 0 and 1 values bouncing around.

We have implemented the *data generation process* (DGP) in Excel. The DGP tells us how our data are produced.

**STEP** Rename the sheet (double-click on the sheet tab) *DGP* and save the workbook as *FreeThrowSim*.

As you scroll back up to the top row, you will see many ones and a few zeroes. With a 90% success rate, roughly one in ten cells will have random number greater than 0.9 and, therefore, show a 0.

The fact that each cell in column B stands alone and does not depend on or influence other cells means we are assuming *independence*. In our model, a miss or make does not affect the chances of hitting the next shot.

If you believe in the hot hand (Cohen, 2020), this implementation of the chance process is wrong. If making the previous shot increases the chances of making the current shot, there is *autocorrelation* and we cannot use 0.9 as the threshold value for every shot. We assume independence from one shot to the next.

How many shots out of 100 will a 90% shooter make?

**STEP** Enter the formula `=SUM(B1:B100)` in cell C1 and the label, “Number made” in cell D1.

You will see a number around 90 in cell C1. Each press of F9 gives you the result of a new outcome from 100 attempted free throws.

The number of made free throws from the virtual shooter you have constructed in Excel is not always exactly 90 because you incorporated `RAND()` in each shot.

**STEP** Use your keyboard shortcut, F9, to recalculate the sheet a few times to get sense of the variability in the number of made shots from 100 free throws.

There is no doubt that the number of made shots is a random variable since it is bouncing around when you recalculate the sheet. It makes common sense that adding 100 bouncing numbers will produce a random outcome.

A *statistic* is a recipe for the data. Cell C1 is a sample statistic because the recipe is to add up the results from a sample of 100 shots. We are interested in the distribution of the sum of 100 free throws from a 90% free throw shooter, including its central tendency, dispersion, and shape of a histogram of outcomes. With this, we can decide if 75 made shots is merely unlikely or so rare that we reject the claim that you are a 90% shooter.

Simulation is simply repeating the experiment many times so we can approximate the center, dispersion, and distribution of the outcomes. Since we are working with a sample statistic, the distribution of the sum of 100 free throws is called a *sampling distribution*.

To process the many outcomes, we need software. A free Excel add-in that does Monte Carlo simulation is available here: [tiny.cc/addin](http://tiny.cc/addin).

**STEP** Download the *MCSim.xla* file from the link above and use the Add-Ins Manager (File → Options → Add-ins → Go) to install it. Click the *Add-ins* tab and click *MCSim*.

**EXCEL TIP** The keyboard shortcut to call the Add-ins Manager is alt, t, i (press these keys in order without holding any of them down).

Figure 3.2 shows the MCSim add-in dialog box. On the left are three required choices. You must select a cell to track (C1 in our example), the number of repetitions (the default is 1,000), and the random number generator to use. The MCSim add-in comes with its own RNG, RANDOM. Selecting it will replace all RAND in the sheet with RANDOM. The default is no changes.

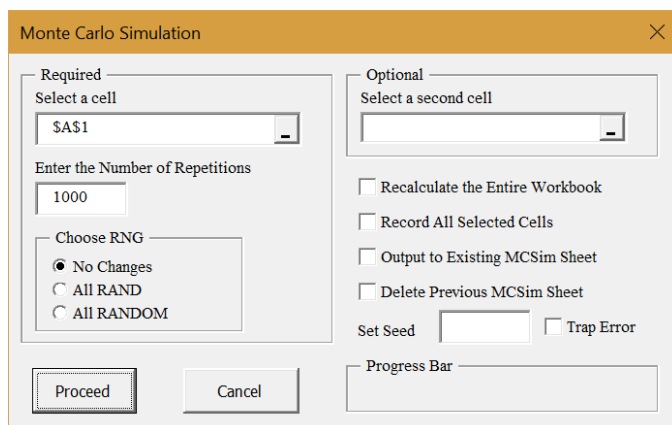


Figure 3.2: The Monte Carlo Simulation Excel Add-in.

On the right are some advanced options. Some of these will be used in future work. The *Set Seed* option forces the RNG to begin from the same initial position, which allows for replication of results.

**STEP** Click in the *Select a cell* box, clear it, and click cell A1 (displaying RAND()). Click in the *Set Seed* box and enter 123. Click *Proceed*.

A new sheet is inserted in the workbook. It shows the first 100 outcomes in column B, summary statistics, and a histogram. It is roughly, but not exactly a rectangle. If you ran more repetitions, it would be less jagged.

**STEP** Return to the *DGP* sheet and repeat the simulation of cell A1.

The results are exactly the same as before because the *Set Seed* option started the RNG from the same initial value.

**STEP** Return to the *DGP* sheet and click the *MCSim* button in the *Add-ins* tab. Select cell C1 (the sum of made free throws) and clear the *Set Seed* box. Click *Proceed*.

Figure 3.3 shows the results. Yours will be different because we cleared the *Set Seed* box. However, your results will be quite close in the sense that your average is near 90 and the SD is around 3.




Figure 3.3: Simulation Results for Sum of 100 Attempts from a 90% Shooter.

The average and SD are approximations to their true exact analogues, the expected value (EV) is exactly 90 and the standard error (SE) is exactly 3. The EV is the center of the sampling distribution, while the SE is the typical deviation, or bounce, in the statistic.



We would say that we expect a 90% free throw shooter to make 90 out of 100 attempts, plus or minus 3 free throws. The plus or minus is critical because it tells us the variability in the number of made free throws.

The sim also shows the maximum and minimum shots made from 100 free throws in 1,000 repetitions. In Figure 3.3, the max is 98 and min is 78.

This gives an answer to our question. In 1,000 repetitions, the worst a 90% shooter did was 78. It certainly looks like you are not a 90% free throw shooter.

What if we did more repetitions?

**STEP** Run a simulation of 10,000 sets of 100 free throws.

Once again, you are unlikely to see 75 or fewer. It seems the bad luck defense is not going to work. While it is possible that you really are a 90% free throw shooter and had an incredibly unlikely run of bad luck, such an outcome is incredibly rare. So rare, in fact, that we do not believe your claim to be a 90% free throw shooter.

The average and SD values changed a little with the second simulation. This shows that simulation always gives an approximate answer with some variability. Simulation can never give us an exact answer because we cannot run an infinity of repetitions. As the number of repetitions increases, the approximation gets better, but it is never exact.

By the way, as mentioned earlier, Gauss and statisticians using his work would have answered this question differently. A simple formula would lead immediately to rejection of your claim.

The procedure begins by computing the SE with the formula  $\sqrt{n}\sigma = \sqrt{100} \times 0.3 = 3$ . Next, express the observed from the expected difference in standard units:  $\frac{75 - 90}{3} = -5$ . This is so far in the tail of the normal (Gaussian) curve that the claim is rejected.

In other words, 75 out of 100 when 90 was claimed is 5 standard units away from what we expected to see and this is ridiculously unlikely so, sorry, we do not believe that you are a 90% free throw shooter who had some bad luck.

In fact, neither analytical methods nor simulation can ever give a definitive, guaranteed answer. Both agree that, given the evidence, 75 out of 100 means we do not believe the claim that you are a 90% shooter. Since chance is involved, it is possible that you are a 90% shooter and missed every shot. We are not interested in what is possible. We want to know how to use the evidence to decide whether or not we believe a claim.

If you have taken a statistics course, you might recognize that we are doing hypothesis testing without explicitly saying so. The null is that you are a 90% shooter and the alternative hypothesis is that you are not. 75 out of 100 produces a test statistic far from the expected 90 so the  $p$ -value is really small. Thus, we reject the null.

## Max Streak

A second example of simulation involves streaks, also known as runs. A streak in this case is a consecutive set of made shots.

***STEP*** Return to the *DGP* sheet. Starting from cell B1, find the first 1 (it could be cell B1), and then count how many 1s in a row you see before you encounter a miss. Write that number down and count the next streak. Continue until you reach the 100<sup>th</sup> shot attempt. The longest streak is the *max streak*.

The question is: what is the length of the typical max streak in a set of 100 free throws from a 90% free throw shooter?

This is an exceedingly difficult question. It is not asking to count the streaks (also a hard question), but to find the biggest streak in 100 shots. You do not simply add up the number made, you have to find the length of all of the streaks and then identify the longest one.

Unlike how many free throws in 100 attempts a 90% shooter will make, you have no easy way to guess the typical max streak. It could be 20, 40, or maybe 50. Who knows? How can we answer this question?

The analytical approach is a bit of a dead end. There are formulas that approximate a solution (Feller, 1968, p. 325), but the math is somewhat complicated. No exact analytical solution has been found.

Simulation can be used if we can figure out a way to ask the question in Excel so that a cell displays the answer. This means simulation requires some ingenuity. We need a cell that computes the max streak so we can use the MCSim add-in. We do this in two steps: first we figure out how to report the current streak, then we use the MAX function to find the longest streak.

**STEP** In cell E1, enter the formula `=B1`. In cell E2, enter the formula `=IF(B2=1,E1+1,0)`. Fill it down a few cells.

Now you can see what the formula is doing. If the shot is made, we add it to the previous running sum, but if it is missed, it resets the running sum to 0. The `B2=1` part of the formula tests if the current shot is made and `E1 + 1` increases the current streak length by one. The zero means you missed and the streak is now zero.

**STEP** Fill the formula down to E100 and look at the values as you scroll back up.

You should see several streaks in a set of 100 free throws. We want the longest streak. That is the second step in our implementation of the question in Excel and it is easy.

**STEP** In cell F1, enter the formula `=MAX(E1:E100)` and enter the label “max streak” in cell G1. Press F9 a few times.

Cell F1 is displaying the max streak from each set of 100 free throws. Max streak is a statistic, just like the sum, because it is a recipe—albeit much more complicated than the sum.

It has an expected value, standard error, and sampling distribution. We can approximate all of these with simulation.

**STEP** Run a simulation, with 10,000 repetitions, of cell F1.

Figure 3.4 shows results. Yours will be a little different. The average is an approximate answer to our question: the max streak is about 27 or so. The exact answer is the expected value, but we have no way of computing it.

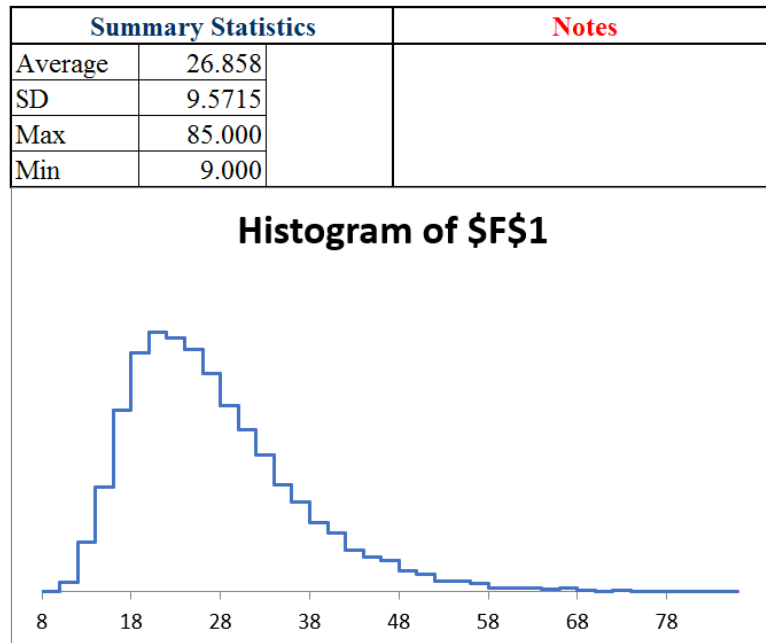


Figure 3.4: Sim Results for Max Streak in 100 Attempts from a 90% Shooter.

The histogram is an approximation to the exact sampling distribution (which no one has figured out how to exactly derive). The graph tells us which values are unlikely: roughly 10 or fewer and 50 or more.

Notice that the sampling distribution of the max streak statistic, unlike the sum, does not appear to follow the normal curve. The max streak has a long right tail and is not symmetric.

## Takeaways

Sometimes, a function or problem is deterministic, but other times we are faced with *stochastic data*—the numbers depend on chance, luck, and randomness. The values we observe are produced by a data generation process, DGP, and they are volatile.

Monte Carlo simulation is a brute force approach to answering questions involving stochastic data. A much older alternative, the analytical approach, relies on brainpower to derive formulas.

To run a simulation, the problem must be implemented in Excel (or some software that can generate random numbers). Of course, one can manually flip a coin many times in the real world, but this is tedious. Simulation did not become a powerful tool until modern computers were invented and enabled a great many repetitions in a short period of time.

Simulation is always only an approximation. By running more repetitions, the approximation improves, but it can never give an exact answer because it would have to run forever.

Often, we are searching for the sampling distribution of a statistic. This tells us the chances of each outcome, the typical result (called the expected value), and the dispersion in values (called the standard error).

The MCSim add-in always produces summary statistics and a histogram. If the cell that is tracked is a statistic, then the average is the approximate expected value and the SD is the approximate SE.

In case you think streaks are a waste of time, look at this headline to an article in *The Wall Street Journal* on July 27, 2023, on page B1:

---

## Dow Sets Longest Winning Streak Since '87

The Dow's streak of 13 consecutive sessions with gains ended the next day. The longest streak ever (as of this writing) is 14, back in 1897.

## References

The epigraph is a famous quote from 1951, when computer science was taking off: "Various Techniques Used in Connection With Random Digits". There are many algorithms for random number generation, some are better and others worse. Excel's RAND() is not great.

Cohen, B. (2020) *The Hot Hand: The Mystery and Science of Streaks*, Custom House. Russ Roberts interviews Cohen in an August 10, 2020 episode of EconTalk available at [www.econtalk.org/ben-cohen-on-the-hot-hand](http://www.econtalk.org/ben-cohen-on-the-hot-hand).

Feller, W. (1968, 3<sup>rd</sup> ed.) *An Introduction to Probability Theory and Its Applications*, John Wiley & Sons. [archive.org/details/introductiontopr0001feller](http://archive.org/details/introductiontopr0001feller)