

## Understanding The Gambler's Fallacy by David Johnson

The “gambler’s fallacy” is an informal fallacy in which the player notices during a game of chance that is considered “fair” that a certain result has had a “run” and the player then bets against the series continuing. The gambler erroneously assumes that there is a higher probability than mere chance that the streak will come to an end. A simple example will help to clarify: If I toss a fair coin and get heads eight times in a row, then many people would say that the odds of the streak ending are higher than of having it continue, and would therefore bet that the next toss will turn up tails. This conclusion is wrong. It is also wrong to believe that the streak has a higher probability of continuing than mere chance, and this is a related fallacy known as the “hot hand” fallacy. The truth is that each flip of the coin is statistically independent of the other tosses, so the odds for each individual toss is always 50-50 no matter what results we have gotten on previous tosses.

This was quite counterintuitive to me when I first learned of it. In fact, I did not actually believe it. Imagine that you and I play a game of chance using a standard, fair coin, with one side heads, and one side tails. Here are the rules: The game begins when we flip the coin and get a run of either 5 heads or 5 tails in a row. I then enter the game and place a bet that the series will end, and go no further than 5. This means, of course, that I am betting we will get the opposite of the side that produced the run on the next toss of the coin. You are then required to place your bet that the run will stop on exactly 50. If the run stops on any number other than 5, or 50, then that round does not count, and neither one of us win or lose anything. We simply don’t count that game, and then we would start over. But if the series stops at exactly 5 in a row, then you pay me a dollar, and if it stops at exactly 50 in a row, then I will pay you a dollar. Would you want to play this game over and over?

You should say no, of course, because runs of 50 heads/tails in a row while using a fair coin are very rare, whereas runs of 5 in a row are not common, but they still happen far more often than 50 in a row. You would be foolish to take that bet, and that should be obvious to everyone. Okay, so now I will change things a little bit. Now, when I enter the game at 5, betting that the series will end, you are required to place your bet that the series will end exactly at 10 in a row. The same rules apply, meaning that either one of us gains or loses any money unless the series stops at either 5 or 10 exactly. Would you take this bet? Series of 10 in a row happen a lot more often than series of 50 in a row, but still, they happen a lot less often than runs of 5 in a row. So, I would still have a big advantage over you if we had the time to play this game over and over. You would go broke taking this side of the bet over time.

Here is the real question, though. Assuming that you agree that it would be a bad idea to bet that the series will go to 50 rather than end at 5, and it would be less bad, but still bad to bet that it would extend to 10 rather than end at 5, why wouldn’t it also be disadvantageous to you to bet that the series will extend to 6 in a row rather than ending at 5, by the same line of reasoning? The only thing that seems to have changed is that the closer the two numbers are, the less

pronounced the advantage is. Since a series of 6 in a row happens only a bit less frequently than series of 5 in a row, my advantage would narrow, but it would still exist.

Based upon this argument, my claim was simply this: if you know that you are going to play a game of chance where you will toss a fair coin 100 times, and you are forced to make a bet, then it is to your advantage to always bet on it being a shorter series rather than the longer one. I believed you could get that slight advantage by simply picking a series, such as after 2 in a row, and always betting on it to end with the next toss rather than having it continue to 3 in a row. I was basing this upon the fact that it is quite likely that you will have a higher number of series of 2 in a row rather than series of 3 in a row during those 100 tosses of the coin.

I didn't think that I was committing the gambler's fallacy because my projections were not based upon past tosses affecting the current one, it was all simply a projection of the future, and the belief that a short series occurs with greater frequency than a longer one, so how could you not have an advantage by betting accordingly? Here was the main problem for me, though. If my argument was right, then that means that odds for the next toss as a single event must be something other than 50-50 on a fair coin. That obviously cannot be right. It seems like odds of getting 6 heads in a row are less than odds of getting 5 in a row, but at the exact same time odds for getting heads or tails on the next toss as a single event are exactly equal. So, what gives? How does this work?

Well, I am an empiricist, so I finally decided to actually start flipping some coins. I hadn't done it before because I was worried that I couldn't get a large enough sample size, but it was bothering me so much, I finally just did it. I started a game where I would toss a different US coin 10 times every day for 100 days. I figured that this would give me a large enough sample size to see if I noticed any potential advantage, and if there was one, I would do more tests. The rules were that every time I experienced a run of two in a row of either heads or tails I would enter a bet on getting the opposite side for the next toss, thus betting the series would end at 2 in a row rather than continuing. If the bet went against me, I did not continue betting, but instead waited for another series of 2 in a row to enter a new bet. Also, if I got a run of 2 on the 10th toss for that day, I would toss the coin one more time to have 11 total tosses for that day instead of 10.

The experiment actually started with me having a significant advantage, and I started to wonder again if maybe my hypothesis was right. After 3 days, and thus 30 tosses, I had won five times and only lost twice. However, it wasn't long before the results averaged out, and I ended up after 50 tosses of the coin with a winning percentage of exactly 50%, with 6 wins and 6 losses. I had originally planned to toss the coin a total of 1000 times, but I didn't need to go any further than 50 before I finally figured out the flaw in my original hypothesis.

The real heart of the problem is that it isn't as simple as just comparing how frequently you get a series of 2 versus a series of 3. The true comparison is how often you get a series of 2 versus how often you get a series of every length over 2 in a row. This is because I had no way of knowing beforehand whether I was on a short series or a long one, so I had to enter the bet whenever I had 2 in a row. That meant that I lost whenever the series went to 3 and I also lost when it went longer. I should have accounted for that, of course, but I didn't think about it until actually tossing the real coins. I was correct that I had a series of 3 in a row less often than series of 2 in a

row; out of 50 tosses I got a series of 3 four times, and a series of 2 six times. So, if it would have just been between a series of 2 versus a series of 3, I would have enjoyed a significant winning percentage. But I also lost when I got a series that went to 4, and on another one that went all the way to 5 in a row. This gave me a total for my bets of exactly 6 wins and 6 losses.

Even though most people do not articulate it in this way, and probably have not thought about the problem in as much depth as I have, I think that most people who fall for the gambler's fallacy make a similar kind of mistake. They know from experience that longer series occur with less frequency than shorter ones, so they try to bet that it will be like that this particular time.

Unfortunately for them, you can only actually identify these patterns and runs with hindsight.

Once I knew from tossing the real coins that the gambler's fallacy is true, I was also able to discover where I went wrong with the first game scenario I presented. Probability theory says that probability is only a measure of the future, so you cannot jump in and place your bet in the middle of the run. In this case, it would be after a run of 5. Once that run has taken place, it no longer has any effect on future probability. I was skeptical of this because in my original game, where I won if the series stopped at 5, and you won if it went to 50, it didn't seem to make any difference whether I placed the bet on the sixth toss after the run of 5 had already occurred, or whether I placed it before the game ever started. Either way, I would win way more often than you would, regardless of whether I placed my bet before or after the run of 5 had actually happened. This also seemed to be true of the game using 5, and 10. So, it didn't seem to matter whether I entered the series at the very beginning, or came in after the run had happened.

But here is what is really going on in that scenario (which suddenly came to me in a flash of insight at approximately 4 a.m. after a long time of thinking about it):

When I enter the game after a run of 5 heads, my probability of winning is based only on the results for the next toss. That means that my odds of winning the bet would be 1/2. Your odds of winning, on the other hand, are  $\frac{1}{2} \times 45$ . This is because what I am really saying is that I win if the next toss is tails, whereas you win if you have a run of 45 heads in a row. This explains why the odds are still so stacked in my favor even if I have entered the game after the run of 5 has already occurred and why it appears to work to enter the bet after the run has already happened. My odds of winning aren't based on comparing a series of 5 to a series of 50; it is based instead upon me winning if there is a series of 1 while you only win with a series of 45. It is even easier to see this when you make the two series being compared closer. So, let's imagine the same scenario where I win if it ends at 5 (so I win if there is tails on the next one) and you win if there is a series of exactly 7. Well, what is really happening when I enter the bet on that sixth toss is that my odds of winning are again  $\frac{1}{2}$ , and your odds of winning are  $\frac{1}{4}$ , because really, your bet is based upon the probability of you getting 2 heads in a row on the next two tosses (we get 1 out of 4 by multiply the odds for the two events together). So again, it makes sense why I would win more often than you would with that bet, but it does not have anything to do with that prior run of 5 in a row, it is only based upon the probability of future results. So, understanding it in this way, it makes perfect sense why odds are only 50-50 when the two series are only 1 toss away from each other. In that case, both of us could have a potential run of 1, or 0, and neither of us would have any betting advantage over the other one, no matter how we bet.

I am sure that probability theory itself already accounts for all that I have said here, but hopefully this is a common sense, easy to understand explanation of the gambler's fallacy for people who are not great at math and need a more intuitive explanation than what you get simply from learning the formulas. As you can obviously tell, I am not a very good mathematician myself. Still, I didn't know any professional mathematicians I could ask about it, but none of the philosophers that I was able to talk to about the fallacy were able to answer the objections that I had. Those who thought it was correct would simply reiterate what the fallacy was without being able to explain it any further than that. So, hopefully my explanation of why it is so will be useful for some.