

Did a roulette system "break the bank"?

In the book *Thirteen Against the Bank*, Norman Leigh claims to have achieved the impossible and devised a system to consistently return a profit from playing roulette. It sounds too good to be true, and maybe it is. **Graham Kendall** uses computer simulation to sort fact from fiction

• O one can win at roulette unless he steals money from the table while the croupier is not looking." That quote, often attributed to Albert Einstein, was challenged in 1966 by Norman Leigh, who claimed to have trained a team of 12 gamblers to play a roulette system that would guarantee a return.

In a later book, called *Thirteen Against the Bank*, Leigh reported that his team won \$163 000 using this system – and got banned from every casino in France in the process.¹

Since the book was published, there has been debate about whether it is a true story or a work of fiction. The first edition states clearly in its subtitle that it is "The true story of a man who broke the bank at the roulette table with an infallible system". However, later versions of the subtitle moderated this statement somewhat, dropping the "true" to describe it as "The story of a man who broke the bank at the roulette table with an unbeatable system".

In his foreword to the book, Leigh writes: "I am willing to accept, in advance, any challenge whatsoever to the feasibility of my method for winning large sums at roulette, which I call the Reverse Labouchère system." In this article, we set out to provide such a challenge. By simulating the game of roulette, we test the strategy proposed in the book to show whether the profits reported by Leigh could have been due to a biased roulette wheel, fluctuations of short-term play or the use of a profitable system.

The basics

The game of roulette involves the spinning of a numbered wheel. A small ball is then thrown onto the wheel while it is in motion. The ball spirals and bounces around until settling on a number.



Graham Kendall is a professor in the School of Computer Science, University of Nottingham, UK and University of Nottingham Malaysia. He is currently the provost and CEO at the University of Nottingham's Malaysia campus and is also a pro-vice chancellor of the university.

© 2018 The Royal Statistical Society

If the player has bet on that number – or the colour of the number (either red or black), or whether it is odd or even, or one of two value ranges in which the number sits – the player wins.

It is a simple game to play, but a difficult one to beat, as the player is entirely at the mercy of luck. Roulette is a set of independent trials, and one spin has no effect on the next. This means that the probability of winning a given bet does not change with each spin of the wheel, and there are no mathematical tricks that can improve a player's chances of winning over a series of spins.

Also, as in most casino games, the house always has an advantage. A European roulette wheel has 37 numbers: 0 to 36. If a gambler backs one number and wins, they are paid odds of 35–1, when the true odds are 36–1. If the gambler bets on the same number repeatedly, over 37 spins they would lose (on average) one unit, or 2.70% of their money. This is what is known as the "house edge". An American wheel has two zeros, making the house edge 5.26%.

Analysis by F. Downton suggests a roulette system cannot return a profit.² However, many systems have been devised to try to beat the roulette table; Leigh's was not the first. A common one is the Martingale system, in which players are certain to win \$1 at some point. The player starts by betting \$1 on an even-chance bet – either high versus low numbers, odd versus even, or red versus black. If they win, they receive \$2 and have achieved their aim of making a \$1 profit. If they lose, however, they must double their bet on the next spin to \$2. If they win this time, they receive \$4 – which equals the \$3 invested so far plus the \$1 profit. But if they lose their \$2 bet, they must double again. This doubling continues until the player wins, at which point they will have won \$1. This system worked for Eloise Peacock, but it was in an online environment where it was felt that the computer random number generator had an effect.³

The Martingale system suffers from three shortcomings. First, a player needs a large bankroll. They may need to bet \$8192, having already bet \$8191, to stay in the game. Second, when the player has invested \$16 383 (\$8191 + \$8192), they are still not sure they will win, and risk having to bet a further \$16 384. Even if one can afford it, it may be a psychological strain to bet this much just to try to win \$1. And third, the casino will have a house limit. Once a bet exceeds that limit, the player will be unable to place a large enough bet to recover their losses.

In defence of the Martingale system, gamblers will argue that a player is unlikely to have a run of (say) 10 spins without a win, so the player is unlikely to reach the house limit. Indeed, Peacock gives an example where black came up 26 times, and our simulations show that this is not uncommon.³ But the system asks players to assume a lot of risk for so little a reward.

The Reverse Labouchère system, used by Leigh and his team, is based on the Martingale system, but here the player increases their bets when they are winning, not after they lose. They start by writing down the numbers 1, 2, 3, 4. The amount staked on an even-chance bet is given by adding the two end numbers, in this case 4 + 1 = 5. If they win, the player adds the amount staked to the line of numbers, giving 1, 2, 3, 4, 5, and the next bet will be 6. But had they lost, they would cross out

the two end numbers, leaving 2 and 3, and the next bet would also be 5. Once the player crosses out all the numbers, they will have lost 10 units. A successful *progression* ends when the next bet exceeds the house limit. The amount won from a progression equals the sum of the numbers on the notepad, less 10 units representing the initial investment.

This would be fine, except a lone player would regularly lose 10 units while waiting for a progression that would wipe out those losses and deliver a large return. However, Leigh's system has six people playing at the same time, each one betting on one of the even-chance bets. In this set-up, when one of the even-chance bets is losing (say, red), then the opposite bet (black, in this case) is winning, and it will have a progression to offset the losses.

But is the Reverse Labouchère system as unbeatable as Leigh claims?

Simulation of a roulette wheel

We simulated a roulette wheel, and Leigh's system, using the Java programming language. The implementation was straightforward, comprising a number of Java classes, but computer simulations of a roulette wheel are open to suggestions that the random number generator is not fair. To validate our simulated wheel, we bet on a single random number for 50 million spins. We ran 30 trials (a total of 1.5 billion spins) for both the European wheel (which has one zero) and the American wheel (which has two zeros). The losses in Table 1 (page 28) are consistent with what we would expect to see based on the house edge (see box, also on page 28).

We further analysed the number of times each number came up on each wheel. For the European wheel, each number should appear with a frequency of 1/37 = 2.70%, and for the American wheel, each number should appear with a frequency of 1/38 = 2.63%. Our results showed that the simulated wheel was fair and balanced in this respect.

We also analysed each of the trials to look for the longest sequence of possible wins on even-chance bets (red or black; odd or even; low or high). The longest sequence often approached 30 spins (e.g. 30 red numbers in a row). We also looked at the number of times that the same even-chance bet came up on 10 or more consecutive spins. This happened at least 23 000 times for each even-chance bet (6 × 23 000 = 138 000 times in total).

Finally, we tracked how many consecutive occurrences there were of each number. In every one of the trials, each number appeared at least four consecutive times. Some appeared five consecutive times, and there were instances of numbers occurring six consecutive times, although this was rare. Looking at just one of the trials on the European wheel (the others are representative), every one of the 37 numbers appeared twice consecutively on more than 35 000 occasions (min = 35 092, max = 35 856, average = 35 495). Each number also appeared three consecutive times (min = 895, max = 1022, average = 963.41) and four consecutive times (min = 16, max = 41, average = 25.97), but only 17 of the numbers occurred five consecutive times (min = 1, max = 3, average = 1.59).

+

()

۲

While we were satisfied that our simulated wheel was fair and balanced, it is worth noting that roulette wheels in real-life settings may exhibit some amount of bias. A bias is something about the wheel that affects the distribution of the numbers that appear, such that (in the long run) they do not appear with their theoretical frequency. This can occur for several reasons. For example, the wheel may not be completely balanced so that it wobbles a small amount when it spins. Or loose frets between the numbers may affect where the ball lands. These defects, and others, might give a small advantage to individual numbers or a spread of numbers.

TABLE 1 Average roulette losses based on betting on a single random number for 1.5 billion spins, and using the Leigh system for 1.5 billion spins. Results are given to four decimal places to highlight small differences between the simulated values and theoretical values (the house edge).

Betting system	European wheel win/loss (%)	American wheel win/loss(%)
Single random number	-2.6744	-5.2598
Leigh's system	-2.7073	-5.2882
House edge	-2.7027	-5.2632

Roulette and the law of large numbers

It is interesting to note from Table 1 that actual losses from our simulated roulette wheels have not entirely converged to theoretical values even after so many spins. The law of large numbers tells us that the average of the results of a large number of trials should be close to the expected value, and that the average should get closer as more trials are performed. Gamblers will sometimes lean on this idea of an expected value in the mistaken belief that a number is due to come up as it has not appeared for a long time. But as Table 1 demonstrates, even after 1.5 billion spins there is still some deviation from theoretical expectations. If a casino were to carry out those spins by hand, at a rate of 30 an hour, it would take more than 5000 years to achieve – and still the gambler should not expect the wheel to act in a way that enables a predictive method of play.



FIGURE 1 Profit/loss for 50 million spins on different roulette wheel configurations, using the Leigh system. The "no bias" results are the losses from a balanced wheel. "Spread bias (0.05)" shows the effect of P = 0.05 and S = 3. "Spread bias (0.25)" sets P = 0.25 and S = 3. "Custom bias" shows the effect of setting two numbers (12 and 30) to have a 5% probability, and then normalising along with all the other numbers.

Given that Leigh's team were betting on every number (except zero), they may have been experiencing a bias as soon as they sat down to play. This bias may have given one of the even-chance bets a higher probability of appearing, even if the players and the casino were not aware that this was happening. Leigh's team would not care where the bias was as the relevant even-chance bet would, if we believe in the system, exploit it. We simulate bias on a roulette wheel in two ways.

- *Spread bias*. This bias selects a number on the wheel, sets its probability, and then evenly reduces (or increases) the probability for a given spread (a set of neighbouring numbers) so that we have an area of the wheel which shows a bias when it is spun. All the probabilities are then normalised so that they sum to 1. The two parameters to the algorithm are the probability given to the number at the centre of the spread, *P*, and how many numbers either side of the central number are affected, *S*. In this article the spread is always *S* = 3.
- *Custom bias*. This bias enables us to target specific numbers, which means that certain numbers are more likely to appear than others. For example, we could set the probability of two specific numbers to be higher than all others. If we set the probability of these two numbers to P = 0.05 and then normalise all the other numbers, so that the total probability sums to 1, the probability for the two selected numbers becomes 4.78% (European) and 4.77% (American). The remaining numbers will have a probability of 2.58% (European) and 2.51% (American). If we use P = 0.25 for two numbers, after normalisation their probabilities become 17.29% (European) and 17.27% (American). The remaining numbers will have a probability of 1.87% (European) and 1.82% (American).

To test that the two types of bias were operating as expected we again ran 30 trials (50 million spins each) on both wheels. These simulations showed that the expected probability of each number appearing was equal to its assigned probability.

Does the system work?

In Leigh's system there are six players, each betting on a separate even-chance bet. The table limit is set to 4000 units, the same value as in the book. We ran the system on both types of wheel (European and American), using 30 independent runs of 50 million spins (1.5 billion spins in total). The results show that the Leigh system produces losses which are very close to the theoretical losses when using a balanced wheel (see Table 1).

These results support the hypothesis that *Thirteen Against the Bank* is a work of fiction. Over the long run, Leigh's system does not return a profit. But is there a scenario in which the system might have worked as described? Could the system work if the wheel was biased, even if the bias was unknown to the players or the casino?

Figure 1 shows the profit and loss from one of the trials of 50 million spins (which is typical of the others). It includes results from the balanced wheel (labelled "no bias") but also

IN DETAIL

 (\bullet)

۲

shows what happens to the results when different biases are introduced.

When we introduce a small spread bias (P = 0.05) the results are slightly better than on a balanced wheel, but the system still makes a loss (-2.48% and -4.82%). With a much larger spread bias (P = 0.25) the European wheel returns a profit, but the American wheel still makes a loss. Looking at one of the simulations from the European wheel (others are equally representative) we find that seven numbers are affected by the bias, four of which are even. Analysis shows that in this simulation, even-chance bets had many more successful progressions and fewer unsuccessful progressions.

All the simulations avoided zero in the spread. This was done by design as when zero comes up, the even chances lose. Some casinos return half your bet when zero turns up, effectively halving the house edge, but we have ignored this in our simulations. To validate the effect of zero, we ran two further simulations, one on each type of wheel. We centred the spread on zero. Having the spread centred on zero resulted in losses of -4.46% (European) and -6.58% (American). This suggests that it is not enough to find a wheel that is biased, you need to find one where zero is not included in the bias.

A bias may not be spread around a single number, but instead specific numbers may be biased. Using this type of bias, which we refer to as a custom bias, we ran two experiments where we set a higher probability for numbers 12 and 30 to come up. On both wheels, these are red, even numbers, so we are introducing a bias for these two evenchance bets, and a negative bias for black, odd numbers. The bias does not affect the high/low even-chance bet.

We set the probability to P = 0.05 for 12 and 30. These two numbers, after normalisation, have a probability of 4.78% for the European wheel and 4.77% for the American wheel, with all other numbers having a probability of 2.58% (European) and 2.51% (American). For this simulation, we return a loss on both wheels.

As a final experiment, we upped the probability of 12 and 30 to P = 0.25. In this scenario, both wheels produce a positive return; about 26% (European) and 24.5% (American). Given that the combined probability of the two numbers, after normalisation, is 34.58% (European) and 34.55% (American), it is probably not surprising that both wheels returned such a large profit.

So far, our simulations have considered the long-term effect of Leigh's system. But in a real casino the system would only be used for a few hours at a time, and certainly for much less than 50 millions spins – never mind 1.5 billion. Let us assume that the system was used for 12 hours a day, with an average of 30 spins an hour. Would the system perform any better over $12 \times 30 = 360$ spins, when it is less likely that the numbers would be evenly distributed? To investigate this scenario, we simulated 360 spins on a balanced wheel, repeated over 30 independent trials.

Over the 30 trials of the European wheel, three showed a profit, with the largest being 14.31%. The other 27 runs resulted in a loss, with the largest being -9.34% (see Figure 2). The average was -4.73%, which is a greater loss than the theoretical expected loss on a balanced wheel of -2.70%.

Spin, spin, spin...

For the data used in this article, and to read notes on the data, visit bit.ly/roulette13. The American wheel showed an overall loss of -8.02% over 30 trials, again larger than the theoretical loss of -5.26%. Only one trial showed a small profit of 0.28% (see Figure 3).

Fact or fiction?

Based on our simulations, we conclude that the book *Thirteen Against the Bank* by Norman Leigh is a work of fiction – which is a shame as it is a very nice story – and that the system it describes cannot, and does not, consistently return a profit. Our future work will include simulating other betting systems, including those of the game of blackjack, so that the scientific archive, and the general public, have a point of reference for systems which claim to "break the bank".

References

 Leigh, N. (1976) Thirteen Against the Bank: The True Story of a Man who Broke the Bank at the Roulette Table with an Infallible System. London: Weidenfeld & Nicolson.

2. Downton, F. (1980) A note on Labouchère sequences. *Journal of the Royal Statistical Society, Series A*, **143**(3), 363–366.

3. Peacock, E. (2013) When the gambler's fallacy comes true: Beating the online casino. *Significance*, **10**(6), 40–42.





December 2018 | significancemagazine.com | 29