

THE THREE DEGREES OF IGNORANCE

BENEATH LONDON'S RITZ HOTEL LIES A HIGH-STAKES CASINO. It's called the Ritz Club, and it prides itself on luxury. Croupiers dressed in black oversee the ornate tables. Renaissance paintings line the walls. Scattered lamps illuminate the gold-trimmed decor. Unfortunately for the casual gambler, the Ritz Club also prides itself on exclusivity. To bet inside, you need to have a membership or a hotel key. And, of course, a healthy bankroll.

One evening in March 2004, a blonde woman walked into the Ritz Club, chaperoned by two men in elegant suits. They were there to play roulette. The group weren't like the other high rollers; they turned down many of the free perks usually doled out to big-money players. Still, their focus paid off, and over the course of the night, they won £100,000. It wasn't exactly a small sum, but it was by no means unusual by Ritz standards. The following night the group returned to the casino and again perched beside a roulette table. This time their winnings were much larger. When they eventually cashed in their chips, they took away £1.2 million.

Casino staff became suspicious. After the gamblers left, security looked at the closed-circuit television footage. What they saw was enough to make them contact the police, and the trio were soon arrested at a hotel not far from the Ritz. The woman, who turned out to be from Hungary, and her accomplices, a pair of Serbians, were accused of obtaining money by deception. According to early media reports, they had used a laser scanner to analyse the roulette table. The measurements were fed into a tiny hidden computer, which converted them into predictions about where the ball would finally land. With a cocktail of gadgetry and glamour, it certainly made for a good story. But a crucial detail was missing from all the accounts. Nobody had explained precisely how it was possible to record the motion of a roulette ball and convert it into a successful prediction. After all, isn't roulette supposed to be random?

THERE ARE TWO WAYS to deal with randomness in roulette, and Henri Poincaré was interested in both of them. It was one of his many interests: in the early twentieth century, pretty much anything that involved mathematics had at some point benefited from Poincaré's attention. He was the last true 'Universalist'; no mathematician since has been able to skip through every part of the field, spotting crucial connections along the way, as he did.

As Poincaré saw it, events like roulette appear random because we are ignorant of what causes them. He suggested we could classify problems according to our level of ignorance. If we know an object's exact initial state—such as its position and speed—and what physical laws it follows, we have a textbook physics problem to solve. Poincaré called this the first

degree of ignorance: we have all the necessary information; we just need to do a few simple calculations. The second degree of ignorance is when we know the physical laws but don't know the exact initial state of the object, or cannot measure it accurately. In this case we must either improve our measurements or limit our predictions to what will happen to the object in the very near future. Finally, we have the third, and most extensive, degree of ignorance. This is when we don't know the initial state of the object or the physical laws. We can also fall into the third level of ignorance if the laws are too intricate to fully unravel. For example, suppose we drop a can of paint into a swimming pool. It might be easy to predict the reaction of the swimmers, but predicting the behaviour of the individual paint and water molecules will be far more difficult.

We could take another approach, however. We could try to understand the effect of the molecules bouncing into each other without studying the minutiae of the interactions between them. If we look at all the particles together, we will be able to see them mix together until—after a certain period of time—the paint spreads evenly throughout the pool. Without knowing anything about the cause, which is too complex to grasp, we can still comment on the eventual effect.

The same can be said for roulette. The trajectory of the ball depends on a number of factors, which we might not be able to grasp simply by glancing at a spinning roulette wheel. Much as for the individual water molecules, we cannot make predictions about a single spin if we do not understand the complex causes behind the ball's trajectory. But, as Poincaré suggested, we don't necessarily have to know what causes the ball to land where it does. Instead, we can simply watch a large number of spins and see what happens.

That is exactly what Albert Hibbs and Roy Walford did in 1947. Hibbs was studying for a math degree at the time, and his friend Walford was a medical student. Taking time off from their studies at the University of Chicago, the pair went to Reno to see whether roulette tables were really as random as casinos thought.

Most roulette tables have kept with the original French design of thirty-eight pockets, with numbers 1 to 36, alternately coloured black and red, plus 0 and 00, coloured green. The zeros tip the game in the casinos' favour. If we placed a series of one-dollar bets on our favourite number, we could expect to win on average once in every thirty-eight attempts, in which case the casino would pay thirty-six dollars. Over the course of thirty-eight spins, we would therefore put down thirty-eight dollars but would only make thirty-six dollars on average. That translates into a loss of two dollars, or about five cents per spin, over the thirty-eight spins.

The house edge relies on there being an equal chance of the roulette wheel producing each number. But, like any machine, a roulette table can have imperfections or can gradually wear down with use. Hibbs and Walford were on the hunt for such tables, which might not have produced an even distribution of numbers. If one number came up more often than the others, it could work to their advantage. They watched spin after spin, hoping to spot something odd. Which raises the question: What do we actually mean by 'odd'?

WHILE POINCARÉ WAS IN France thinking about the origins of randomness, on the other side of the English Channel Karl Pearson was spending his summer holiday flipping coins. By the time

the vacation was over, the mathematician had flipped a shilling twenty-five thousand times, diligently recording the results of each throw. Most of the work was done outside, which Pearson said ‘gave me, I have little doubt, a bad reputation in the neighbourhood where I was staying’. As well as experimenting with shillings, Pearson got a colleague to flip a penny more than eight thousand times and repeatedly pull raffle tickets from a bag.

To understand randomness, Pearson believed it was important to collect as much data as possible. As he put it, we have ‘no absolute knowledge of natural phenomena,’ just ‘knowledge of our sensations’. And Pearson didn’t stop at coin tosses and raffle draws. In search of more data, he turned his attention to the roulette tables of Monte Carlo.

Like Poincaré, Pearson was something of a polymath. In addition to his interest in chance, he wrote plays and poetry and studied physics and philosophy. English by birth, Pearson had travelled widely. He was particularly keen on German culture: when University of Heidelberg admin staff accidentally recorded his name as Karl instead of Carl, he kept the new spelling.

Unfortunately, his planned trip to Monte Carlo did not look promising. He knew it would be near impossible to obtain funding for a ‘research visit’ to the casinos of the French Riviera. But perhaps he didn’t need to watch the tables. It turned out that the newspaper *Le Monaco* published a record of roulette outcomes every week. Pearson decided to focus on results from a four-week period during the summer of 1892. First he looked at the proportions of red and black outcomes. If a roulette wheel were spun an infinite number of times—and the zeros were ignored—he would have expected the overall ratio of red to black to approach 50/50.

Out of the sixteen thousand or so spins published by *Le Monaco*, 50.15 per cent came up red. To work out whether the difference was down to chance, Pearson calculated the amount the observed spins deviated from 50 per cent. Then he compared this with the variation that would be expected if the wheels were random. He found that a 0.15 per cent difference wasn’t particularly unusual, and it certainly didn’t give him a reason to doubt the randomness of the wheels.

Red and black might have come up a similar number of times, but Pearson wanted to test other things, too. Next, he looked at how often the same colour came up several times in a row. Gamblers can become obsessed with such runs of luck. Take the night of August 18, 1913, when a roulette ball in one of Monte Carlo’s casinos landed on black over a dozen times in a row. Gamblers crowded around the table to see what would happen next. Surely another black couldn’t appear? As the table spun, people piled their money onto red. The ball landed on black again. More money went on red. Another black appeared. And another. And another. In total, the ball bounced into a black pocket twenty-six times in a row. If the wheel had been random, each spin would have been completely unrelated to the others. A sequence of blacks wouldn’t have made a red more likely. Yet the gamblers that evening believed that it would. This psychological bias has since been known as the ‘Monte Carlo fallacy’.

When Pearson compared the length of runs of different colours with the frequencies that he’d expect if the wheels were random, something looked wrong. Runs of two or three of the same colour were scarcer than they should have been. And runs of a single colour—say, a black sandwiched between two reds—were far too common. Pearson calculated the probability of observing an outcome at least as extreme as this one, assuming that the roulette

wheel was truly random. This probability, which he dubbed the p value, was tiny. So small, in fact, that Pearson said that even if he'd been watching the Monte Carlo tables since the start of Earth's history, he would not have expected to see a result that extreme. He believed it was conclusive evidence that roulette was not a game of chance.

The discovery infuriated him. He'd hoped that roulette wheels would be a good source of random data and was angry that his giant casino-shaped laboratory was generating unreliable results. 'The man of science may proudly predict the results of tossing halfpence,' he said, 'but the Monte Carlo roulette confounds his theories and mocks at his laws.' With the roulette wheels clearly of little use to his research, Pearson suggested that the casinos be closed down and their assets donated to science. However, it later emerged that Pearson's odd results weren't really due to faulty wheels. Although *Le Monaco* paid reporters to watch the roulette tables and record the outcomes, the reporters had decided it was easier just to make up the numbers.

Unlike the idle journalists, Hibbs and Walford actually watched the roulette wheels when they visited Reno. They discovered that one in four wheels had a bias of some sort. One wheel was especially skewed, so betting on it caused the pair's initial one-hundred-dollar stake to grow rapidly. Reports of their final profits differ, but whatever they made, it was enough to buy a yacht and sail it around the Caribbean for a year.

There are plenty of stories about gamblers who've succeeded using a similar approach. Many have told the tale of the Victorian engineer Joseph Jagger, who made a fortune exploiting a biased wheel in Monte Carlo, and of the Argentine syndicate that cleaned up in government-owned casinos in the early 1950s. We might think that, thanks to Pearson's test, spotting a vulnerable wheel is fairly straightforward. But finding a biased roulette wheel isn't the same as finding a profitable one.

In 1948, a statistician named Allan Wilson recorded the spins of a roulette wheel for twenty-four hours a day over four weeks. When he used Pearson's test to find out whether each number had the same chance of appearing, it was clear the wheel was biased. Yet it wasn't clear how he should bet. When Wilson published his data, he issued a challenge to his gambling-inclined readers. 'On what statistical basis,' he asked, 'should you decide to play a given roulette number?'

It took thirty-five years for a solution to emerge. Mathematician Stewart Ethier eventually realized that the trick wasn't to test for a non-random wheel but to test for one that would be favourable when betting. Even if we were to look at a huge number of spins and find substantial evidence that one of the thirty-eight numbers came up more often than others, it might not be enough to make a profit. The number would have to appear on average at least once every thirty-six spins; otherwise, we would still expect to lose out to the casino.

The most common number in Wilson's roulette data was nineteen, but Ethier's test found no evidence that betting on it would be profitable over time. Although it was clear the wheel wasn't random, there didn't seem to be any favourable numbers. Ethier was aware that his method had probably arrived too late for most gamblers: in the years since Hibbs and Walford had won big in Reno, biased wheels had gradually faded into extinction. But roulette did not remain unbeatable for long.

WHEN WE ARE AT our deepest level of ignorance, with causes that are too complex to understand, the only thing we can do is look at a large number of events together and see whether any patterns emerge. As we've seen, this statistical approach can be successful if a roulette wheel is biased. Without knowing anything about the physics of a roulette spin, we can make predictions about what might come up.

But what if there's no bias or insufficient time to collect lots of data? The trio that won at the Ritz didn't watch loads of spins, hoping to identify a biased table. They looked at the trajectory of the roulette ball as it travelled around the wheel. This meant escaping not just Poincaré's third level of ignorance but his second one as well.

This is no small feat. Even if we pick apart the physical processes that cause a roulette ball to follow the path it does, we cannot necessarily predict where it will land. Unlike paint molecules crashing into water, the causes are not too complex to grasp. Instead, the cause can be too small to spot: a tiny difference in the initial speed of the ball makes a big difference to where it finally settles. Poincaré argued that a difference in the starting state of a roulette ball—one so tiny it escapes our attention—can lead to an effect so large we cannot miss it, and then we say that the effect is down to chance.

The problem, which is known as 'sensitive dependence on initial conditions', means that even if we collect detailed measurements about a process—whether a roulette spin or a tropical storm—a small oversight could have dramatic consequences. Seventy years before mathematician Edward Lorenz gave a talk asking 'Does the flap of a butterfly's wings in Brazil set off a tornado in Texas?' Poincaré had outlined the 'butterfly effect'.

Lorenz's work, which grew into chaos theory, focused chiefly on prediction. He was motivated by a desire to make better forecasts about the weather and to find a way to see further into the future. Poincaré was interested in the opposite problem: How long does it take for a process to become random? In fact, does the path of a roulette ball ever become truly random?

Poincaré was inspired by roulette, but he made his breakthrough by studying a much grander set of trajectories. During the nineteenth century, astronomers had sketched out the asteroids that lay scattered along the Zodiac. They'd found that these asteroids were pretty much uniformly distributed across the night sky. And Poincaré wanted to work out why this was the case.

He knew that the asteroids must follow Kepler's laws of motion and that it was impossible to know their initial speed. As Poincaré put it, 'The Zodiac may be regarded as an immense roulette board on which the Creator has thrown a very great number of small balls.' To understand the pattern of the asteroids, Poincaré therefore decided to compare the total distance a hypothetical object travels with the number of times it rotates around a point.

Imagine you unroll an incredibly long, and incredibly smooth, sheet of wallpaper. Laying the sheet flat, you take a marble and set it rolling along the paper. Then you set another going, followed by several more. Some marbles you set rolling quickly, others slowly. Because the wallpaper is smooth, the quick ones soon roll far into the distance, while the slow ones make their way along the sheet much more gradually.

The marbles roll on and on, and after a while you take a snapshot of their current positions.

To mark their locations, you make a little cut in the edge of the paper next to each one. Then you remove the marbles and roll the sheet back up. If you look at the edge of the roll, each cut will be equally likely to appear at any position around the circumference. This happens because the length of the sheet—and hence the distance the marbles can travel—is much longer than the diameter of the roll. A small change in the marbles' overall distance has a big effect on where the cuts appear on the circumference. If you wait long enough, this sensitivity to initial conditions will mean that the locations of the cuts will appear random. Poincaré showed the same thing happens with asteroid orbits. Over time, they will end up evenly spread along the Zodiac.

To Poincaré, the Zodiac and the roulette table were merely two illustrations of the same idea. He suggested that after a large number of turns, a roulette ball's finishing position would also be completely random. He pointed out that certain betting options would tumble into the realm of randomness sooner than others. Because roulette slots are alternately coloured red and black, predicting which of the two appears meant calculating exactly where the ball will land. This would become extremely difficult after even a spin or two. Other options, such as predicting which half of the table the ball lands in, were less sensitive to initial conditions. It would therefore take a lot of spins before the result becomes as good as random.

Fortunately for gamblers, a roulette ball does not spin for an extremely long period of time (although there is an oft-repeated myth that mathematician Blaise Pascal invented roulette while trying to build a perpetual motion machine). As a result, gamblers can—in theory—avoid falling into Poincaré's second degree of ignorance by measuring the initial path of the roulette ball. They just need to work out what measurements to take.

THE RITZ WASN'T THE first time a story of roulette-tracking technology emerged. Eight years after Hibbs and Walford had exploited that biased wheel in Reno, Edward Thorp sat in a common room at the University of California, Los Angeles, discussing get-rich-quick schemes with his fellow students. It was a glorious Sunday afternoon, and the group was debating how to beat roulette. When one of the others said that casino wheels were generally flawless, something clicked in Thorp's mind. Thorp had just started a PhD in physics, and it occurred to him that beating a robust, well-maintained wheel wasn't really a question of statistics. It was a physics problem. As Thorp put it, 'The orbiting roulette ball suddenly seemed like a planet in its stately, precise and predictable path.'

In 1955, Thorp got hold of a half-size roulette table and set to work analysing the spins with a camera and stopwatch. He soon noticed that his particular wheel had so many flaws that it made prediction hopeless. But he persevered and studied the physics of the problem in any way he could. On one occasion, Thorp failed to come to the door when his in-laws arrived for dinner. They eventually found him inside rolling marbles along the kitchen floor in the midst of an experiment to find out how far each would travel.

After completing his PhD, Thorp headed east to work at the Massachusetts Institute of Technology. There he met Claude Shannon, one of the university's academic giants. Over the previous decade, Shannon had pioneered the field of 'information theory', which revolutionised how data are stored and communicated; the work would later help pave the way

for space missions, mobile phones and the Internet.

Thorp told Shannon about the roulette predictions, and the professor suggested they continue the work at his house a few miles outside the city. When Thorp entered Shannon's basement, it became clear quite how much Shannon liked gadgets. The room was an inventor's playground. Shannon must have had \$100,000 worth of motors, pulleys, switches and gears down there. He even had a pair of huge polystyrene 'shoes' that allowed him to take strolls on the water of a nearby lake, much to his neighbours' alarm. Before long, Thorp and Shannon had added a \$1,500 industry-standard roulette table to the gadget collection.

MOST ROULETTE WHEELS ARE operated in a way that allows gamblers to collect information on the ball's trajectory before they bet. After setting the centre of the roulette wheel spinning counterclockwise, the croupier launches the ball in a clockwise direction, sending it circling around the wheel's upper edge. Once the ball has looped around a few times, the croupier calls 'no more bets' or—if casinos like their patter to have a hint of Gallic charm—'*rien ne va plus.*' Eventually, the ball hits one of the deflectors scattered around the edge of the wheel and drops into a pocket. Unfortunately for gamblers, the ball's trajectory is what mathematicians call 'nonlinear': the input (its speed) is not directly proportional to the output (where it lands). In other words, Thorp and Shannon had ended up back in Poincaré's third level of ignorance.

Rather than trying to dig themselves out by deriving equations for the ball's motion, they instead decided to rely on past observations. They ran experiments to see how long a ball travelling at a certain speed would remain on the track and used this information to make predictions. During a spin, they would time how long it took for a ball to travel once around the table and then compared the time to their previous results to estimate when it would hit a deflector.

The calculations needed to be done at the roulette table, so at the end of 1960, Thorp and Shannon built the world's first wearable computer and took it to Vegas. They tested it only once, as the wires were unreliable, needing frequent repairs. Even so, it seemed like the computer could be a successful tool. Because the system handed gamblers an advantage, Shannon thought casinos might abandon roulette once word of the research got out. Secrecy was therefore of the utmost importance. As Thorp recalled, 'He mentioned that social network theorists studying the spread of rumors claimed that two people chosen at random in, say, the United States are usually linked by three or fewer acquaintances, or "three degrees of separation."' The idea of 'six degrees of separation' would eventually creep into popular culture, thanks to a highly publicised 1967 experiment by sociologist Stanley Milgram. In the study, participants were asked to help a letter get to a target recipient by sending it to whichever of their acquaintances they thought were most likely to know the target. On average, the letter passed through the hands of six people before eventually reaching its destination, and the six degrees phenomenon was born. Yet subsequent research has shown that Shannon's suggestion of three degrees of separation was probably closer to the mark. In 2012, researchers analysing Facebook connections—which are a fairly good proxy for real-life acquaintances—found that there are an average of 3.74 degrees of separation between any two people. Evidently, Shannon's fears were well founded.

TOWARD THE END OF 1977, the New York Academy of Sciences hosted the first major conference on chaos theory. They invited a diverse mix of researchers, including James Yorke, the mathematician who first coined the term ‘chaotic’ to describe ordered yet unpredictable phenomena like roulette and weather, and Robert May, an ecologist studying population dynamics at Princeton University.

Another attendee was a young physicist from the University of California, Santa Cruz. For his PhD, Robert Shaw was studying the motion of running water. But that wasn’t the only project he was working on. Along with some fellow students, he’d also been developing a way to take on the casinos of Nevada. They called themselves the ‘Eudaemons’—a nod to the ancient Greek philosophical notion of happiness—and the group’s attempts to beat the house at roulette have since become part of gambling legend.

The project started in late 1975 when Doyne Farmer and Norman Packard, two graduate students at UC Santa Cruz, bought a refurbished roulette wheel. The pair had spent the previous summer toying with betting systems for a variety of games before eventually settling on roulette. Despite Shannon’s warnings, Thorp had made a cryptic reference to roulette being beatable in one of his books; this throwaway comment, tucked away toward the end of the text, was enough to persuade Farmer and Packard that roulette was worth further study. Working at night in the university physics lab, they gradually unravelled the physics of a roulette spin. By taking measurements as the ball circled the wheel, they discovered they would be able to glean enough information to make profitable bets.

One of the Eudaemons, Thomas Bass, later documented the group’s exploits in his book *The Eudaemonic Pie*. He described how, after honing their calculations, the group hid a computer inside a shoe and used it to predict the ball’s path in a number of casinos. But there was one piece of information Bass didn’t include: the equations underpinning the Eudaemons’ prediction method.

MOST MATHEMATICIANS WITH AN interest in gambling will have heard the story of the Eudaemons. Some will also have wondered whether such prediction is feasible. When a new paper on roulette appeared in the journal *Chaos* in 2012, however, it revealed that someone had finally put the method to the test.

Michael Small had first come across *The Eudaemonic Pie* while working for a South African investment bank. He wasn’t a gambler and didn’t like casinos. Still, he was curious about the shoe computer. For his PhD, he’d analysed systems with nonlinear dynamics, a category that roulette fell very nicely into. Ten years passed, and Small moved to Asia to take a job at Hong Kong Polytechnic University. Along with Chi Kong Tse, a fellow researcher in the engineering department, Small decided that building a roulette computer could be a good project for undergraduates.

It might seem strange that it took so long for researchers to publicly test such a well-known roulette strategy. However, it isn’t easy to get access to a roulette wheel. Casino games aren’t generally on university procurement lists, so there are limited opportunities to study roulette. Pearson relied on dodgy newspaper reports because he couldn’t persuade anyone to fund a trip to Monte Carlo, and without Shannon’s patronage, Thorp would have struggled to carry out his

roulette experiments.

The mathematical nuts and bolts of roulette have also hindered research into the problem. Not because the maths behind roulette is too complex but because it's too simple. Journal editors can be picky about the types of scientific papers they publish, and trying to beat roulette with basic physics isn't a topic they usually go for. There has been the occasional article about roulette, such as the paper Thorp published that described his method. But though Thorp gave enough away to persuade readers—including the Eudaemons—that computer-based prediction could be successful, he omitted the details. The crucial calculations were notably absent.

Once Small and Tse had convinced the university to buy a wheel, they got to work trying to reproduce the Eudaemons' prediction method. They started by dividing the trajectory of the ball into three separate phases. When a croupier sets a roulette wheel in motion, the ball initially rotates around the upper rim while the centre of the wheel spins in the opposite direction. During this time, two competing forces act on the ball: centripetal force keeping it on the rim, and gravity pulling it down toward the centre of the wheel.

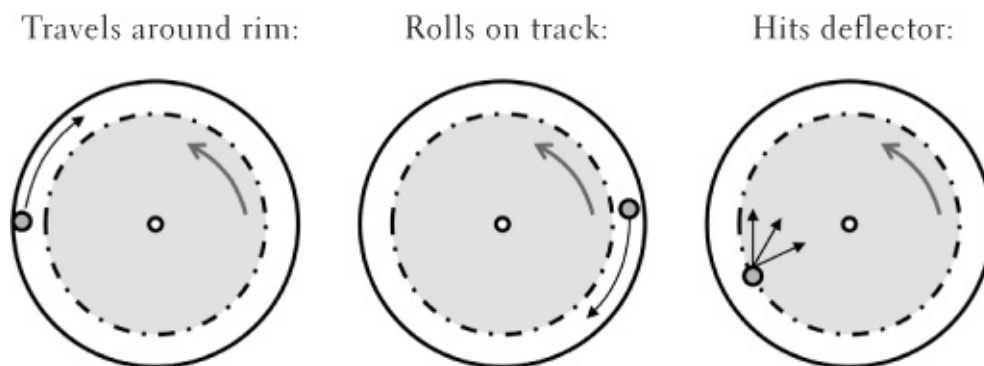


FIGURE 1.1. The three stages of a roulette spin.

The pair assumed that as the ball rolls, friction slows it down. Eventually, the ball's angular momentum decreases so much that gravity becomes the dominant force. At this point, the ball moves into its second phase. It leaves the rim and rolls freely on the track between the rim and the deflectors. It moves closer to the centre of the wheel until it hits one of the deflectors scattered around the circumference.

Until this point, the ball's trajectory can be calculated using textbook physics. But once it hits a deflector, it scatters, potentially landing in one of several pockets. From a betting point of view, the ball leaves a cosy predictable world and moves into a phase that is truly chaotic.

Small and Tse could have used a statistical approach to deal with this uncertainty. However, for the sake of simplicity, they decided to define their prediction as the number the ball was next to when it hit a deflector. To predict the point at which the ball would clip one of the deflectors, Small and Tse needed six pieces of information: the position, velocity and acceleration of the ball, and the same for the wheel. Fortunately, these six measurements could be reduced to three if they considered the trajectories from a different standpoint. To an onlooker watching a roulette table, the ball appears to move in one direction and the wheel in the other. But it is also possible to do the calculations from a 'ball's-eye view', in which case it's only necessary to measure how the ball moves relative to the wheel. Small and Tse did this

by using a stopwatch to clock the times at which the ball passed a specific point.

One afternoon, Small ran an initial series of experiments to test the method. Having written a computer program on his laptop to do the calculations, he set the ball spinning, taking the necessary measurements by hand, as the Eudaemons would have done. As the ball travelled around the rim a dozen or so times, he gathered enough information to make predictions about where it would land. He only had time to run the experiment twenty-two times before he had to leave the office. Out of these attempts, he predicted the correct number three times. Had he just been making random guesses, the probability he would have got at least this many right (the p value) was less than 2 per cent. This persuaded him that the Eudaemons' strategy worked. It seemed that roulette really could be beaten with physics.

Having tested the method by hand, Small and Tse set up a high-speed camera to collect more precise measurements about the ball's position. The camera took photos of the wheel at a rate of about ninety frames per second. This made it possible to explore what happened after the ball hit a deflector. With the help of two engineering students, Small and Tse spun the wheel seven hundred times, recording the difference between their prediction and the final outcome. Collecting this information together, they calculated the probability of the ball landing a specified distance away from the predicted pocket. For most of the pockets, this probability wasn't particularly large or small; it was pretty much what they'd have expected if picking pockets at random. Some patterns did emerge, however. The ball landed in the predicted pocket far more often than it would have if the process were down to chance. Moreover, it rarely landed on the numbers that lay on the wheel directly before the predicted pocket. This made sense because the ball would have to bounce backward to get to these pockets.

The camera showed what happened in the ideal situation—when there was very good information about the trajectory of the ball—but most gamblers would struggle to sneak a high-speed camera into a casino. Instead, they would have had to rely on measurements taken by hand. Small and Tse found this wasn't such a disadvantage: they suggested that predictions made with a stopwatch could still provide gamblers with an expected profit of 18 per cent.

After announcing his results, Small received messages from gamblers who were using the method in real casinos. 'One guy sent me detailed descriptions of his work,' he said, 'including fabulous photos of a "clicker" device made from a modified computer mouse strapped to his toe.' The work also came to the attention of Doyne Farmer. He was sailing in Florida when heard about Small and Tse's paper. Farmer had kept his method under wraps for over thirty years because—much like Small—he disliked casinos. The trips he made to Nevada during his time with the Eudaemons were enough to convince him that gambling addicts were being exploited by the industry. If people wanted to use computers to beat roulette, he didn't want to say anything that would hand the advantage back to the casinos. However, when Small and Tse's paper was published, Farmer decided it was time to finally break his silence. Especially because there was an important difference between the Eudaemons' approach and the one the Hong Kong researchers had suggested.

Small and Tse had assumed that friction was the main force slowing the ball down, but Farmer disagreed. He'd found that air resistance—not friction—was the main reason for the ball slowing down. Indeed, Farmer pointed out that if we placed a roulette table in a room

with no air (and hence no air resistance), the ball would spin around the table thousands of times before settling on a number.

Like Small and Tse's approach, Farmer's method required that certain values be estimated while at the roulette table. During their casino trips, the Eudaemons had three things to pin down: the amount of air resistance, the velocity of the ball when it dropped off the rim of the wheel, and the rate at which the wheel was decelerating. One of the biggest challenges was estimating air resistance and drop velocity. Both influenced the prediction in a similar way: assuming a smaller resistance was much like having an increased velocity.

It was also important to know what was happening around the roulette ball. External factors can have a big effect on a physical process. Take a game of billiards. If you have a perfectly smooth table, a shot will cause the balls to ricochet in a cobweb of collisions. To predict where the cue ball will go after a few seconds, you'd need to know precisely how it was struck. But if you want to make longer-term predictions, Farmer and his colleagues have pointed out it's not enough to merely know about the shot. You also need to take into account forces such as gravity—and not just that of the earth. To predict exactly where the cue ball will travel after one minute, you have to include the gravitational pull of particles at the edge of the galaxy in your calculations.

When making roulette predictions, obtaining correct information about the state of the table is crucial. Even a change in the weather can affect results. The Eudaemons found that if they calibrated their calculations when the weather was sunny in Santa Cruz, the arrival of fog would cause the ball to leave the track half a rotation earlier than they had expected. Other disruptions were closer to home. During one casino visit, Farmer had to abandon betting because an overweight man was resting against the table, tilting the wheel and messing up the predictions.

The biggest hindrance for the group, though, was their technical equipment. They implemented the betting strategy by having one person record the spins and another place the bets, so as not to raise the suspicions of casino security. The idea was that a wireless signal would transmit messages telling the player with the chips which number to bet on. But the system often failed: the signal would disappear, taking the betting instructions with it. Although the group had a 20 per cent edge over the casino in theory, these technical problems meant it was never converted into a grand fortune.

As computers have improved, a handful of people have managed to come up with better roulette devices. Most rarely make it into the news, with the exception of the trio who won at the Ritz in 2004. On that occasion, newspapers were particularly quick to latch on to the story of a laser scanner. Yet when journalist Ben Beasley-Murray talked to industry insiders a few months after the incident, they dismissed suggestions that lasers were involved. Instead, it was likely the Ritz gamblers used mobile phones to time the spinning wheel. The basic method would have been similar to the one the Eudaemons used, but advances in technology meant it could be implemented much more effectively. According to ex-Eudaemon Norman Packard, the whole thing would have been pretty easy to set up.

It was also perfectly legal. Although the Ritz group were accused of obtaining money by deception—a form of theft—they hadn't actually tampered with the game. Nobody had

interfered with the ball or switched chips. Nine months after the group's initial arrest, police therefore closed the case and returned the £1.3 million haul. In many ways, the trio had the UK's wonderfully archaic gambling laws to thank for their prize. The Gaming Act, which was signed in 1845, had not been updated to cope with the new methods available to gamblers.

Unfortunately, the law does not hand an advantage only to gamblers. The unwritten agreement you have with a casino—pick the correct number and be rewarded with money—is not legally binding in the UK. You can't take a casino to court if you win and it doesn't pay up. And although casinos love gamblers with a losing system, they are less keen on those with winning strategies. Regardless of which strategy you use, you'll have to escape house countermeasures. When Hibbs and Walford passed \$5,000 in winnings by hunting for biased tables in Reno, the casino shuffled the roulette tables around to foil them. Even though the Eudaemons didn't need to watch the table for long periods of time, they still had to beat a hasty retreat from casinos on occasion.

AS WELL AS DRAWING the attention of casino security, successful roulette strategies have something else in common: all rely on the fact that casinos believe the wheels are unpredictable. When they aren't, people who have watched the table for long enough can exploit the bias. When the wheel is perfect, and churns out numbers that are uniformly distributed, it can be vulnerable if gamblers collect enough information about the ball's trajectory.

The evolution of successful roulette strategies reflects how the science of chance has developed during the past century. Early efforts to beat roulette involved escaping Poincaré's third level of ignorance, where nothing about the physical process is known. Pearson's work on roulette was purely statistical, aiming to find patterns in data. Later attempts to profit from the game, including the exploits at the Ritz, took a different approach. These strategies tried to overcome Poincaré's second level of ignorance and solve the problem of roulette's outcome being incredibly sensitive to the initial state of the wheel and ball.

For Poincaré, roulette was a way to illustrate his idea that simple physical processes could descend into what seems like randomness. This idea formed a crucial part of chaos theory, which emerged as a new academic field in the 1970s. During this period, roulette was always lurking in the background. In fact, many of the Eudaemons would go on to publish papers on chaotic systems. One of Robert Shaw's projects demonstrated that the steady rhythm of droplets from a dripping tap turns into an unpredictable beat as the tap is unscrewed further. This was one of the first real-life examples of a 'chaotic transition' whereby a process switches from a regular pattern to one that is as good as random. Interest in chaos theory and roulette does not appear to have dampened over the years. The topics can still capture the public imagination, as shown by the extensive media attention given to Small and Tse's paper in 2012.

Roulette might be a seductive intellectual challenge, but it isn't the easiest—or most reliable—way to make money. To start with, there is the problem of casino table limits. The Eudaemons played for small stakes, which helped them keep a low profile but also put a cap on potential winnings. Playing at high-stakes tables might bring in more money, but it will also

bring additional scrutiny from casino security. Then there are the legal issues. Roulette computers are banned in many countries, and even if they aren't, casinos are understandably hostile toward anyone who uses one. This makes it tricky to earn good profits.

For these reasons, roulette is really only a small part of the scientific betting story. Since the shoe-computer exploits of the Eudaemons, gamblers have been busy tackling other games. Like roulette, many of these games have a long-standing reputation for being unbeatable. And like roulette, people are using scientific approaches to show just how wrong that reputation can be.