

Section Two

The Wheels

It doesn't require an extensive mathematical background to look at the 38 identically-sized spaces on an American roulette wheel (note the 35-1 payoff on a single number) and conclude that the game is unbeatable. With a $1/38$ chance of having a number come up on the next spin and the 35-1 payoff, it is easy to calculate the often-quoted expectancy of the player of -5.26 . The odds for other wheels, especially the Wheel of Fortune, appear even more against the player.

The unbeatability of the roulette wheel is based on the mechanical perfection of the wheel—such a conclusion is based on the assumption that the ball has an equal chance of landing in each pocket. This may or may not be true, although Allan Wilson, in *The Casino Gambler's Guide*, and others give fairly convincing evidence for the existence of biased wheels—wheels sufficiently biased to overcome the house advantage.

The very mechanical perfection of the wheel, however, would suggest the applicability of the laws of physics to prediction of the next number, whether the game is roulette or the Wheel of Fortune. Just as the future position of a planet can be predicted quite accurately, so can an understanding of the physical laws at work minimize the uncertainty surrounding the

resting place of the ball or the final position of the wheel.

It is not possible, of course, to obtain an exact prediction. But this is not absolutely necessary to assure a profit. As Marvin Karllins has pointed out in his book *Psyching Out Vegas*, "Simply being able to predict which half of the wheel the ball will plunk into would give the player such a whopping edge that he could go for the chandeliers . . . and make it."

The following two chapters investigate the promise of this approach to beating the wheel as well as discussing some of the difficulties that might arise implementing such a strategy in the casino environment.

Roulette

It was the spring of 1955. I was finishing my second year of graduate physics at U.C.L.A. In the course of the next year I would make three decisions that would shape my life for the next 28 years. I married (my present wife, Vivian), I changed my field of study from physics to mathematics, and I began to toy with the fantasy that I could shatter the chains of poverty through a scientifically-based winning gambling system.

I was living in Robison Hall, the student-owned cooperative. For \$50 a month and four hours work a week, we got our room and board. I had lived in the co-ops for nearly six years of undergraduate and graduate work, on a budget of about \$100 a month. Part of this came from scholarships and, in the early years, I got some help from home. But I was basically self-supporting like most of the other 200 or so co-op residents.

I attended classes and studied from 50 to 60 hours a week, generally including Saturdays and Sundays. I had read about the psychology of learning in order to be able to work longer and harder. I found that "spaced learning" worked well: study for an

hour, then take a break of at least ten minutes (shower, meal, tea, errands, etc.). One Sunday afternoon about 3 p.m., I came to the co-op dining room for a tea break. The sun was streaming through the big glass windows. (Robison, designed by Richard Neutra in the '30s, was very radical for that time. It had so many big sheet glass windows that it was often called "the glass house.") My head was bubbling with physics equations, and several of my good friends were sitting around chatting.

In our mutual poverty the conversation readily turned to fantasies of easy money. We began to speculate on whether there was a way to beat the roulette wheel. In addition to me, the group included math majors Mel Rosenfeld and Andy Bruckner (now professors of mathematics at U.C. Santa Barbara), Tom Scott, and engineering major Rick Rushall. After all these years it's hard to be sure of exactly who said what, but we began the discussion by acknowledging that mathematical systems were impossible. I'll demonstrate this in a future chapter.

Then we kicked around the idea of whether croupiers could control where the ball will land well enough to significantly affect the odds. I will show later that this is impossible under the usual conditions of the game. (The incredible thing is that logical reasoning could even be used to settle such a question.) It was a short brainstorming step to wondering whether wheels were imperfect enough to change the odds to favor the player. Those in the group who "knew" assured me that the wheels are veritable jeweled watches of perfection, carefully machined, balanced and maintained. This is false. Wheels are sometimes imperfect enough so they can be beaten. I had no experience with gambling, or with casinos, or with roulette wheels, so I accepted the mechanical perfection of roulette wheels.

But mechanical perfection, for a physicist, means predictability. You can't have it both ways, I argued. If these wheels are very imperfect the odds will change enough so we can beat them. If they are perfect enough we can predict (in principle) approximately where the ball will land. Suddenly the orbiting roulette ball seemed like the planets in their stately and precise, predictable paths. In

my mind there was that intuitive "click" of discovery that I would experience again and again. Unknowingly, I had just taken the first step on a long journey in which I would discover winning systems such as those for blackjack and for the options market, and I would accumulate a wealth I never imagined.

One side argued that it is a long way from prediction in principle to practical prediction. My group said that, over and over, the story of science has been a rapid leap from a theoretical vision ($E=MC^2$) to an unexpected practical result (nuclear power plants). By now our initial group of people agreed that the idea had merit and might well work. The novel debate attracted listeners, some of them cynical. They challenged us to prove the idea worked. The ten minute "study break" had run into a couple of hours. We adjourned with the half definite idea of "doing something."

In the following weeks the idea kept coming back to me: measure the position and velocity of the roulette ball at a fixed time and (maybe) you can then predict its future path, including when and where the ball will spiral into the rotor. (The rotor is the spinning circular central disc where the ball finally comes to rest in numbered pockets.) Also measure the rotor's position and velocity at a (possibly different) fixed time and you can predict the rotor's rotation for any future time. But then you will know what section of the rotor will be there when the ball arrives. So you know (approximately) what number will come up!

You can see that the system requires that bets be placed *after* the ball and rotor are set in motion and somehow timed. That means that the casinos have a simple, perfect countermeasure: forbid bets after the ball is launched. However, I have checked games throughout the world, including Reno, Las Vegas, London, Venice, Monte Carlo, and Nice. Only in a few cases were bets forbidden after the ball was launched. A common practice instead was to call "no more bets" a revolution or two before the ball dropped into the center.

The simple casino countermeasure meant that there were two problems: (1) find out whether exact enough predictions could be

made to get a winning edge, first in theory and then in the casino itself, and (2) camouflage the system so the casinos would be unaware of its use. If we could solve the prediction problem, the camouflage was easy. Have an observer standing by the wheel recording the numbers that came up, as part of a "system." Many do this so it doesn't seem out of place. But the observer also wears a concealed computer device with timing switches. His real job is to time the ball and rotor. (Much later we settled on toe-operated switches, leaving both hands free and in the open.) The computer would make the prediction and transmit it by radio to the bettor. The bettor, at the far end of the layout, would appear to have no connection to the observer-timer. The bettor would have a poor view of ball and rotor and would not pay much attention to them. To further break any link between timer and bettor, I would have several of each, with identical devices. They would each come and go "at random."

The important bets have to be placed after the ball is launched. A bettor who only bet then, and who consistently won, would soon become suspect. To avoid that, I planned to have the bettor also make bets before the ball was launched. These would be limited so their negative expectation didn't cancel all the positive expectation of the other bets. I became a radio amateur (W6VVM) when I was 13 (back in 1945 when there weren't easy novice-class tests), so I thought I could build the radio link and other electronic gadgetry.

This left me with the prediction problem to solve. More than a year passed without much time for roulette: I got my Master's degree in Physics (June 1955) and wrote the first part of my Ph.D. thesis on nuclear shell structure (Mayer-Jensen theory). The mathematical problems that I ran into led me in the fall of 1955 to take graduate math courses. I needed so many that I got my Ph.D. in math instead! And early in 1956 I got married. I had been working as a tutor and one of my "students" was T.T. Thornton. He was an independently wealthy, knowledge-loving bachelor of about 45, who had degrees in English and chemistry. Now he was getting a degree in mathematics, just for the pleasure

of it. He was an excellent student who didn't need a tutor but had hired me simply to learn faster and more efficiently.

We shared bits and pieces of our hopes, dreams, and enthusiasms.

After I had mentioned the roulette project, I was surprised and touched by his gift of a half-sized wheel. It was black plastic (bakelite?) made in France. I learned later that it cost the enormous sum of \$25. Though I had thought about the roulette system off and on, the gift of this wheel (sometime in 1958, I recall) got me to work more seriously on it. My first idea was to use a home movie camera to film the orbiting ball. I then plotted the amount the ball had traveled versus the number of the frame of the film. I expected that the pictures were taken at a uniform rate of 24 (?) frames per second so I could plot (angular) distance traveled versus time as in Figure 4-1. Instead of a smooth graph like the solid line in Figure 4-1, my first film showed a peculiar wavy structure, like the dashed line.

After thinking about this, I guessed that this was because the camera did not run at uniform speed. By taking a movie of a stopwatch that timed in hundredths of a second, I found that the camera did vary in speed. Photo stores confirmed this. The distortion of the curve in Figure 4-1 is analogous to the way a musical tone is distorted by a phono turntable whose speed varies slightly.

My next move was to take a movie of the rotating ball and the stopwatch. This gave me an accurate time for each frame. (I still have a roll of these pictures, postmarked January 16, 1959.) But there was still some "ripple" to the curves. (I later learned that even a slight tilt would cause this.) Worse, I found that the curves were not consistent from spin to spin. The situation was something like Figure 4-2. This meant the ball behaved differently from spin to spin. This meant that the distance it traveled varied even with the same initial velocity. This doomed predictability on my wheel.

I found with further experiments that my half-sized wheel was really very irregular. The track was curved like a tube and the ball

Figure 4-1





Figure 4-2



“rattled around” erratically, up and down, as it orbited. The slick bakelite surface was moulded, not machined. The ball also skidded and bounced. And there was a horizontal junction which added irregularities to the track.

But full-sized wheels were not like that. In December 1958, I made my first visit to the casinos. I observed several regulation wheels and found that the ball moved smoothly in its track. Also the track was a pair of flat-beveled, carefully-machined surfaces, not a tube. When I saw how good the casino wheels were, I was more convinced than ever that prediction was possible. But I needed a full-sized wheel and some good laboratory equipment to continue. How could I pay for it? I got my Ph.D. in June of 1958 and was teaching at U.C.L.A. Though my wife was finally able to stop working, we had no savings and I barely supported us. I couldn't ask her to go back to work to buy me a roulette wheel and to finance my pipe dream.

But I persisted. I simulated the study of the problem of whether the roulette ball would, for the same starting velocity, travel about the same distance along the track. I set up a little vee-shaped inclined trough. I would start a marble from a fixed height (a mark on the trough) and measure how far across the floor it rolled. I was encouraged but not surprised to find that the distance the marble went could be predicted closely from the starting height.

One memorable evening when my in-laws were due for dinner, I ran overtime on a marble experiment. They came into the kitchen wondering why I hadn't come to greet them at the door. They found me rolling marbles down a little wooden trough and across the floor. All over the floor were little distance markers and pieces of tape.

In early 1959 Vivian and I spent time with Mel and Judy Rosenfeld, working on a radio link for the casino test of my yet to be completed roulette system. We took model airplane radio control equipment and altered it somewhat. We succeeded in getting a workable but somewhat inconvenient radio link.

Then around March or April of 1959, I pushed the roulette pro-

ject aside. Twelve man years of blackjack calculations arrived, courtesy of Baldwin, Cantey, Maisel and McDermott. I had convinced myself (as described in *Beat the Dealer*) that I could devise a winning blackjack card counting system and now I set to work on this intensely. The impractical marble roller now said he could beat the casinos at blackjack. What next?

I wrote my blackjack computer programs in the summer and fall of 1959. Testing, then debugging followed, and then from late 1959 through early 1960 my computer production runs produced the basic results that gave me the five-count system in early 1960. Then during 1960 I worked out most of the ten-count system and the ideas for the ultimate strategy. I also made the computer runs and worked out the methodology so that all of today's so-called “one parameter” blackjack systems could be readily devised by anyone versed in the use of computers. In December 1960, The Notices of the American Mathematical Society carried the abstract of my upcoming talk, “Fortune's Formula: The Game of Blackjack.” Life would never be the same again. The intense professional and public interest aroused by the abstract, even before the talk, led me to seek quick publication in a scientific journal. I chose to try the Proceedings of the National Academy of Sciences. I needed a member of the Academy to communicate (i.e. approve and forward for recommended publication), so I sought out the one mathematics member of the Academy at M.I.T., Claude Shannon.

Claude Shannon: Genius

Shannon, then in his early forties, was and is one of the most famous applied mathematicians in the world. As one genius among many, he was relatively unnoticed as a graduate student—until he handed in his master's thesis. It developed the mathematical theory of switching electrical networks (e.g. telephone exchanges) and became the landmark paper in the subject. After receiving his doctorate, Shannon worked at Bell labs for several years and then became world-famous for papers

establishing the mathematical foundations of information theory.

I was able to arrange a short appointment early one chilly December afternoon. But the secretary warned me that Shannon was only going to be in for a few minutes, not to expect more, and that he didn't spend time on subjects (or people) that didn't interest him (enlightened self-interest, I thought to myself).

Feeling both awed and lucky, I arrived at Shannon's office for my appointment. He was a thinnish alert man of middle height and build, somewhat sharp featured. His eyes had a genial crinkle and the brows suggested his puckish incisive humor. I told the blackjack story briefly and showed him my paper. We changed the title from "A Winning Strategy for Blackjack" to "A Favorable Strategy for Twenty-One" (more sedate and respectable). I reluctantly accepted some suggestions for condensation, and we agreed that I'd send him the retyped revision right away for forwarding to the Academy.

Shannon was impressed with both my blackjack results and my method and cross-examined me in detail, both to understand and to find possible flaws. After my few minutes were up, he pointed out in closing that I appeared to have made the big theoretical breakthrough on the subject and that what remained to be discovered would be more in the way of details and elaboration. And then he asked, "Are you working on anything else in the gambling area?"

I decided to spill my other big secret and told him about roulette. Several exciting hours later, as the wintery sky turned dusky, we finally broke off with plans to meet again on the roulette project. Shannon lived in a huge old three story wooden house on one of the Mystic Lakes, several miles from Cambridge. His basement was a gadgeteer's paradise. It had perhaps a hundred thousand dollars worth of electronic, electrical and mechanical items. There were hundreds of categories, like motors, transistors, switches, pulleys, tools, condensers, transformers, and on and on.

Our work continued there. We ordered a regulation roulette

wheel from Reno and assembled other equipment including (most important) a strobe light and a large clock with a second hand that made one revolution in one second. The dial was divided into hundredths of a second and still finer time divisions could be estimated closely. We set up shop in "the billiard room," where a massive old dusty slate billiard table made a perfect solid stable mounting for the roulette wheel.

Analyzing the Motion

My original plan was to divide the various motions of ball and rotor into parts and analyze each one separately. They were:

- The ball is launched by the croupier. It orbits on a horizontal track on the stator until it slows down enough to fall off this (sloped) track towards the center (rotor). Assume at first that (a) the wheel is perfectly level, and (b), the velocity of the ball depends on how many revolutions it has left before falling off. Referring to Figure 4-2, (b) means that every spin would produce the same curve, not different ones like my half-sized wheel. Put another way, this means that if you timed one revolution of the ball on the stator, you could tell how many more revolutions and how much more time until the ball left the track. If these assumptions turned out to be poor, we would attempt to modify the analysis.

- Next analyze the portion of the ball orbit from the time the ball leaves the track until it crosses from the stator to the rotor. If the wheel is perfectly level and there are no obstacles, then it seems plausible that this would always take the same amount of time. (We later learned that wheels are often significantly tilted. This tilt, when it occurs, can affect the analysis substantially. We eventually learned how to use it to our advantage.) There are, however, vanes, obstacles, or deflectors on this portion of the wheel. The size, number, and arrangement vary from wheel to wheel.

On average, perhaps half the time these have a significant effect on the ball. Sometimes they knock it abruptly down into the

rotor, tending to cause it to come to rest sooner. This is typical of “vertical” deflectors (ones approximately perpendicular to the ball’s path). Other times they “stretch out” the ball’s path, causing it to enter the rotor at a more grazing angle and to come to rest later, on average. This is typical of “horizontal” deflectors (ones approximately parallel to the ball’s path).

- Assume the rotor is stationary (not real), and beat that situation first. Reasoning: if you can’t beat a stationary rotor, you can’t beat the more complex moving rotor. Here the uncertainty is due to the ball being “spattered” by the frets (the dividers between the numbered pockets). Sometimes a ball will hit a fret and bounce several pockets on, other times it will be knocked backwards. Or it may be stopped dead. Occasionally the ball will bounce out to the edge of the rotor and move most of a revolution there before falling back into the inner ring of pockets. Thus, even if we knew where the ball would enter the rotor, the “spattering” from the frets causes considerable uncertainty regarding where it finally stops. This tells you that there is no possible reliable “physical” method for predicting ahead of time which pocket the ball is going to land in, unless the wheel is grossly defective or crooked. That makes the roulette method “used” in the movie “The Honeymoon Machine,” where the players forecasted the exact pocket, an impossibility. It also tells you that successful physical prediction can at most forecast with an advantage which sector of the wheel the ball will end in.

- Assume now that the rotor is moving. Generally the ball and rotor move in opposite directions; increasing the velocity of the ball relative to the rotor. We’ll assume this is always the case. I’ve never seen or heard of a casino spinning ball and rotor in the same direction. If this were done, the relative motion of ball and rotor would be even less than with a stationary rotor and prediction would be easier yet. With a moving rotor, the amount of ball “spattering” is increased and predictability is further reduced. Note that this change depends on the rotor velocity. Since that varies from time to time and from croupier to croupier, this adds further complexity. It turns out that the velocity of the rotor changes very slowly, so it is possible to predict with high accuracy

which part of the rotor will be “there” at the predicted time and place that the ball leaves the stator.

I will now take you through a simplified version of what we first tried to do. Later, with that overview to guide us, I’ll explain some of the modifications we had to make and describe our casino experiences.

First, let’s consider part 1, the motion of the ball on the track. The actual function $x(t)$, which describes the number of remaining revolutions x versus the remaining time t , is theoretically very complex.*

Our first problem, and the key one, was to predict when and where on the stator the ball would leave the track. This problem was key because once we knew this, everything else except rotor velocity was a “constant.” And rotor velocity is easy to measure in advance and incorporate into the prediction, as we shall see. Our method was to measure the time of one ball revolution. If the time were short, the ball was “fast” and had a long way to go. If the time were “long,” the ball was “slow” and would soon fall from the track.


We hit a microswitch as the ball passed a reference mark on the stator. This started the electronic clock. This was at time t_1 (to go) with x_1 revolutions to go. (There are many such “marks” available on all actual casino wheels.) When the ball passed the reference mark the second time we hit the switch again, stopping the electronic clock. That was at a time t_0 (left to go) before the ball left the track) with x_0 revolutions left. The clock measured $t_1 - t_0$, the time T for one revolution (so $x_1 - x_0 = 1$.)*

Movie Experiments

The function $x(t)$ which we are using in this illustration is not the actual one. The actual $x(t)$ can be determined by a “movie experiment” like the ones I described earlier which I did in 1959 on my half-size wheel. To do this experiment today, get a full-size roulette wheel, a large clock which reads accurately in hundredths of a second or better, and a video camera or movie camera. Then take a movie of the orbiting ball. The successive frames give

values for t and $x(t)$, which can be plotted to get an $x(t)$ curve like that of Figure 4-3. Several movies should be made to see how much the $x(t)$ curve varies from one spin to another. This uncertainty is a source of errors in determining T , that I'll discuss later on. These $x(t)$ errors can be incorporated into the theory in the same way as the timing errors. They each cause some uncertainty in the predicted $X_0(T)$ value. The data from the movie experiment can be improved if the camera frames are synchronized to a strobe so that the motion of both ball and clock is "stopped" rather than blurry. I didn't do this in my original movies, so I got a short blurry arc, instead of a ball, in each frame.

If an appropriate clock is not available, you can use a high quality phonograph turntable instead. These rotate at very uniform speeds which can be verified for your turntable with a strobe. Now get a stiff paper disc and mark the edges in equal small units. Number these units (much as you would a "circular" ruler) for ease in reading. Now place a thin fixed pointer just



above the disc. When the disc rotates, you have a very accurate clock whose hand is fixed and whose face moves. If you use a paper disc of polar coordinate graph paper (glued, perhaps, to an old record), there will be 360 equally spaced degree marks.

At $33\frac{1}{3}$ r.p.m., each mark is $1/200$ sec. At 45 r.p.m., each mark is $1/270$ sec., and at 78 r.p.m., each mark is $1/468$ sec. On a 12-inch disc, the 360 marks will be spaced about a tenth of an inch apart so additional marks can be used or the pictures can simply be read to a fraction of an interval. Record test discs with equally spaced "spokes," for use with a strobe for testing turntables, are also available and can be used.

Timing Errors

Shannon and I used the switch which measured T to flash a strobe as well as start and stop the clock. We discovered the lights and the strobe flash "stopped" the ball at each of the two instants the switch was hit. This allowed us to see how much the ball was off the reference mark. Since we knew approximately how fast the ball was moving, we could tell about how much in time we were early or late in hitting the switch. This enabled us to correct the times recorded on the clock, thereby making the data much more accurate. We also learned from the visual feedback how to become much more accurate at timing.

Here's an illustration. Suppose the track of the wheel was 25 inches in diameter. (I don't have any of this equipment now so I'm remembering back over 20 years and recalling about what the sizes, velocities, etc. seemed to be. They'll be close enough to be representative and good enough to show you how to do it all again, better for you if you want to.) Suppose the ball is $\frac{3}{4}$ inch in diameter and T , the time for one revolution, is 0.8 seconds. Then the track is 78.54 inches in length, or 98.17 ball diameters. If the ball center is one diameter away from the reference mark when the strobe flashes, then the timing error is about $1/98.17$ of T or about $8/1000$ of a second. There will be one of these errors when the switch is first hit and another when it is hit the second time. With practice we were able to reduce each error to a typical (root

mean square) size of one ball diameter or about 8/1000 seconds. According to the theory of errors, the two errors together give a typical (root mean square) size of $\sqrt{2}/1000$ or about 11.2/1000 seconds.

These errors would be unobservable in casino play, so we couldn't correct for them there. The critical question is how do they affect the prediction? *

A Simple Casino Countermeasure

It should be clear that for this method to work, we have to time the ball (and rotor) before placing our potentially winning bets. (Earlier bets are losing, on average, so are only camouflage.) Thus, the casino must allow us to continue to bet for a time after the ball is launched. I have observed roulette wheels all over the world: Monte Carlo (our final goal), Nevada, Puerto Rico, Nice, Venice, and London. The practice has been, generally but not always, to allow bets until the ball was almost ready to fall off the track. This was much longer than we needed. Be warned again, though; all the casino needs to do to prevent our method is to forbid bets once the ball is launched. That simple perfect countermeasure is the Achilles heel of the system and a major reason why I never made a total effort to implement it. (People who use the system in casino play say the casinos don't catch on and don't use the countermeasure. But if the player is not really careful, I would expect the casino to catch on.)

The ball timing errors cause errors in predicting both the time and place the ball leaves the track. Even if the spiral path of the ball down the stator into the rotor is always the same in time and distance, this still yields errors in predicting when and where on the rotor the ball enters. *

Error Analysis

We have a long list of sources for errors in the prediction of the ball's final position. They are:

E1 Rotor timing—use 1.4 pockets to illustrate.

E2 Ball timing—use 5.5 pockets to illustrate.

E3 Variations in ball “paths” on rotor (see Figure 4-1). Error size is unknown, call it *X*.

E4 Ball path down stator: error due primarily to “deflectors” and varies with the type and placement. Use seven pockets to illustrate.

E5 Variation in distance ball travels on rotor: error due primarily to frets between pockets “spattering” ball, plus occasional very long paths along the rim of the rotor “outside” the pockets. Use six pockets to illustrate.

E6 Tilted wheel. (We didn't know about this yet.)

For illustrative purposes, assume the errors approximately obey the normal probability distribution. Then the standard deviation (typical size) of the sum of several errors is the square root of the sum of all the squared errors. For instance, using “pockets” as our unit, combined errors *E4* + *E5* have typical size $\sqrt{6^2 + 7^2} = \sqrt{85} = 9.2$ pockets. Now add on the timing errors: *E1* + *E2* + *E4* + *E5* have typical size $\sqrt{1.4^2 + 5.5^2 + 6^2 + 7^2} = \sqrt{117.21} = 10.8$ pockets. Thus the timing errors in this example cause very little additional error: just $10.8 - 9.2$, or 1.6 pockets.*

Of course, we haven't added in *E3* yet and, if *X* is big enough, it could ruin everything. Possible variations in the ball orbit behavior on the stator were difficult for us to measure because we found it hard to tell at exactly what point the ball lost contact with the outer wall of the wheel. We also learned from both our own lab experiences and from watching in the casinos why the orbit varied somewhat. Once a drunken, cigar-smoking bettor knocked his ash onto the track. This was hard to clear out. It got on the ball and spread out on the track. That immediately changed the ball's

The Mathematics of Gambling

behavior. Skin oil from our fingers or the croupier's would slowly "poison" ball and track and seem to affect the orbit behavior.

If we or the croupier gave the ball lots of axial "spin" (in the sense of tennis or ping pong), it could take several revolutions around the track before this abnormal spin energy was converted to orbit energy. (We named this effect after the famous quantum mechanics concept of "spin-orbit coupling.") On the other hand, the ball might be launched with no spin or backspin, so it would skid for a while before spin and orbit got "into synch."

Advantage Versus Error

Obviously, the greater the error, the less the advantage. If we assume the total prediction error E is (approximately) normally distributed, then we can construct a table showing the player's expected gain or loss as a function of E .

Table 4-4 gives the results for a bet on the best pocket and also for a bet on the best "octant." The best octant is a set of five pockets, two on each side of the best pocket.

The Table shows that, when the prediction error is normally distributed, the typical forecast error (standard deviation) must be 16 pockets or less, in order for the bettor to have an advantage. This is 16/38, or about 0.42 revolutions. This is true both for bets on the best pocket and the best octant. Since the best octant includes four pockets that aren't quite as good as the best, the advantage is somewhat less for a given typical error E . However, as we will see later in discussing the Kelly-Breiman system for money management, it is generally better for a small to medium-sized bankroll to bet the best octant.

Kimmel and the Dealer's Signature

Stephen Kimmel asserted that a dealer who works eight hours a day, 50 weeks a year, tends to spin the ball and rotor in a habitual, regular way. This would make possible accurate predictions—a bet on ten pockets, Kimmel contended, would have a 50% chance of success. His views were contained in an article "Roulette and Randomness" in the December, 1979 issue of *Gambling Times*.

Table 4-4

Typical Error E (No. of Pockets)	Percent Advantage Betting on Best		Typical Error E (No. of Pockets)	Percent Advantage Betting on Best	
	Pocket	Octant		Pocket	Octant
0	3500.00	620.00	16	0.46	0.30
1	1278.53	611.06	17	— 1.62	— 1.72
2	610.69	467.86	18	— 3.01	— 3.07
3	376.52	328.65	19	— 3.90	— 3.94
4	258.12	236.98	20	— 4.46	— 4.49
5	186.76	175.71	21	— 4.81	— 4.82
6	139.09	132.62	22	— 5.01	— 5.02
7	105.00	100.89	23	— 5.13	— 5.13
8	79.41	76.65	24	— 5.19	— 5.19
9	59.54	57.60	25	— 5.23	— 5.23
10	43.77	42.38	26	— 5.24	— 5.25
11	31.19	30.18	27	— 5.25	— 5.25
12	21.24	20.52	28	— 5.26	— 5.26
13	13.54	13.03	29	— 5.26	— 5.26
14	7.73	7.37	30	— 5.26	— 5.26
15	3.47	3.24	00	— 5.26	— 5.26

I don't believe Kimmel's approach works. Here's why: there are three important conditions that must remain roughly constant throughout play for the player to take advantage of the regularity of the dealer's signature. These conditions are (1) the rotor velocity should be approximately the same each time the ball is spun, (2) the spinning ball should make approximately the same number of revolutions each time, and (3) the initial position of the rotor when the dealer launches the ball should be approximately the same each time. This third condition, which is not mentioned in Kimmel's article, is crucial.

By way of illustration, suppose that the rotor velocity was exactly the same each time and that the dealer spun the ball exactly the same number of revolutions in each instance. Suppose further that the ball spun exactly eight revolutions and the rotor four revolutions during this time. Given those assumptions, the ball would land about 12 revolutions beyond the point where it was launched. In other words, if the number 13 was passing the ball as the dealer released it, the ball would arrive 12 revolutions later, *relative to the spinning rotor*, at approximately the number 13. You can see, however, that if the number 2 on the rotor was closest to the ball at the instant it was released, the ball would then end up near that number 12 revolutions later.

If the dealer releases the ball without regard to which number on the spinning rotor is closest to the launch point, the ball would randomly fall on the rotor 12 revolutions later. In this case, there would be no predictability whatsoever, even though the rotor velocity is absolutely fixed and the number of ball revolutions constant. Any variance in rotor velocity or number of ball revolutions would further guarantee a random outcome. Because Kimmel did not discuss variations in the point of release, I do not believe in his method.

There is a better approach to this statistical analysis of roulette. Watch a dealer and count the number of revolutions the ball makes on the stator from the time of release until it crosses onto the rotor. Note how constant that number of revolutions is. The results of your observations can be statistically stated as some average number of revolutions plus an error term.

Next, count the number of revolutions the rotor makes during the time the ball is on the stator. This will give you another average for the number of rotor revolutions, plus a second error term. Finally, count how far the ball travels on the rotor after it has crossed the divider between the rotor and stator. You can summarize these results as some average number of revolutions or pockets plus an error term.

In order for this approach to work, it is necessary that the square root of the sums of the squares of the error terms be less

than 17 pockets. The proof of this appears in Table 4-4 which shows what the rate of return is, given various root mean square errors. That table demonstrates that a positive return is possible only when that root mean square error is less than 17 pockets.

Now for the improved method. In the unlikely event that the root mean square error is less than 17 pockets, then—and only then—you have a chance to win. The key lies in using the position of the rotor when the ball is launched as your starting point for predicting where the ball will fall out on the wheel.

For example, suppose you find that for a certain dealer the ball travels eight revolutions with a root mean square error of five pockets. Suppose also that during this time, the rotor travels four revolutions, with a root mean square error of six pockets. And suppose still further that once the ball is on the rotor, it travels 13 pockets with a root mean square error of eight pockets. Given these suppositions, you can predict that the ball will travel eight revolutions plus four revolutions plus 13 pockets from the launch position, or 13 pockets beyond that point. The root mean square error is the square root of five squared plus six squared plus eight squared. This turns out to be 11.2 pockets, well within the required error of less than 17 pockets. In this case, the prediction system would work.

However, I think you will find that when you collect this data, the errors at each stage are several times as large as I have used in this example. My own observation is that the dealer error in the number of revolutions for the ball spin is about 20 pockets for the more consistent dealers; it is much larger with a less consistent one. I also noticed that the rotor velocity is not nearly as constant as Kimmel would like. That is because the dealer gives it an extra kick every few spins to rebuild its velocity.

It is also true that the deflecting vanes on the sides of the rotor add considerable randomness to the outcome, as do the frets or spacers between the pockets. The upshot is that I don't believe that any dealer is predictable enough to cause a root mean square error of less than 17 pockets. I'm willing to examine proof to the contrary, but I would be very surprised if anyone could ever pro-

duce it.

If a dealer dutifully practiced spinning the ball a fixed number of revolutions, and if a motor drive spun the rotor at a constant velocity, and if we have a very good way of deciding exactly which number is opposite the ball just as it is released, it might be barely possible to gain a small prediction advantage. I consider even that very unlikely.

In closing, I'll give you the perfect casino countermeasure to the strategy of the dealer's signature, pretending for the moment that the strategy worked. First, the casino halts the betting before the dealer spins the ball. Second, the dealer closes his eyes or looks away from the wheel when he releases the ball so that he has no knowledge of which number on the rotor is closest to the ball when it is launched. Then, for the reasons explained above, the result will be perfectly random.

Chapter 5

The Wheel of Fortune

In the last chapter, I described a system for winning at roulette based on physical prediction. That system was developed largely in 1961 and 1962 in collaboration with Claude Shannon at MIT. One by-product was an even simpler system for physical prediction of the Wheel of Fortune. A story about me and blackjack card-counting in *Life* magazine, March 27, 1964, reported on this in a section entitled "Beating the Wheel of Fortune with the Big Toe."

While I was at the Fifth Annual Conference on Gambling and Risk Taking at Caesars Tahoe in October of 1981, I collected data on a Wheel of Fortune at Caesars. I wanted to see whether their wheels could still be predicted in the same way.

My Casio C-80 watch has a digital stop watch feature which times to 1/100 of a second. I used it to time one revolution of the wheel and then recorded how many revolutions it went. I collected the data in Table 5-1 at the Wheel of Fortune nearest to Caesars' cashier cage.

To see how predictable the Wheel was, I looked for a

Table 5-1

Wheel of Fortune Data Caesars Tahoe											
Observation Number	Time T	Further Raw Data	Revs R Decimal	Prediction P	Error P-R	Error In "Pegs"	Expected Time E	Time Difference T-E (Sec)			
1	5.11	3.5 + 22p	3.907	3.856	-.051	-2.8	5.08	+.03			
2	5.33	3.5 - 2p	3.463	3.527	+.064	3.5	5.38	-.05			
3	5.09	4 - 6p	3.889	3.889	.000	0.0	5.09	.00			
4	4.83	4 + 15p	4.278	4.345	.067	3.6	4.87	-.04			
5	5.87	3 - 9p	2.833	2.876	.043	2.3	5.91	-.04			
6	5.67	3 + 10p	3.185	3.095	-.090	-4.9	5.59	+.08			
7	7.14	2 - 6p	1.889	1.901	.012	0.6	7.16	-.02			
8	5.78	3 + 3p	3.056	2.972	-.084	-4.5	5.70	.08			
9	4.66	4.5 + 8p	4.648	4.687	.039	2.1	4.68	-.02			
10	5.16	4 - 9p	3.833	3.778	-.055	-3.0	5.12	.04			
11	6.37	2.5 - 9p	2.333	2.419	.086	4.6	6.48	-.11			
12	6.93	2 + 3p	2.056	2.024	-.032	-1.7	6.88	.05			

mathematical curve which would best fit these data points. A curve which worked well was $R = A \text{ times } T \text{ to the } B \text{ power}$ where $A = 121.545$ and $B = -2.11153$. In the equation, T is the time for the wheel to make one revolution and R is the number of additional revolutions which it then travels. Intuitively, if T is short, the wheel did one revolution quickly so it will go far and R will be large. But if T is long, the wheel was slow and will stop soon so R will be small.

The letter p in the third column of the Table ("raw data") stands for "pegs." The wheel has pegs separating the payoff numbers. As the wheel rotates, the pegs push past a flexible "flapper." This gradually slows the wheel. When the wheel stops, the winning number is the one with the flapper between its pegs.

The raw data column gives $3.5 + 22p$ for observation number 1. This means that the wheel traveled 3.5 revolutions plus 22 pegs or further numbers. Since there are 54 numbers in all, it went $3.5 + 22/54$ or 3.907 revolutions in all. That is shown under "decimal" in column 4.

The prediction P is made from the equation. The "error" P-R is the amount the prediction is off from what actually happened. Strictly speaking, what I am calling a prediction is only a fit to the data. The fit approaches a "true" fit more closely as more data is included. However, there is generally a difference between the "true" fit and the actual fitted equation.

New data tends to cluster around this slightly different unknown true fit, so it will tend to deviate from the actual fit to the data by this extra amount. Thus, we expect future data to be predicted by the equation not quite as well as the data in Table 5-1.

The error P-R has a standard deviation ("typical size") of .0587 revolutions, or 3.2 numbers. The true curve location (standard deviation of the curve) is probably within .0169 revolutions or 0.9 numbers, on average. Considering this and the greatest positive and negative values in the column, error in "pegs" suggests that the prediction will almost always be within five "pegs" or positions of the actual outcome.

Table 5-2 shows the actual arrangement of numbers on the

wheel. They are listed in order, clockwise, as seen by the player. Each number gives the profit per unit bet. Thus, a player who bets on 2 wins \$2 for each \$1 bet. The number marked 40A, and called Caesars, pays 40 to 1 and the number 40B, called Cleo, also pays 40 to 1. A bet on one of them does not win if the other one comes up.

There are 24 "ones" in Table 5-2. Thus, if each of the 54 numbers comes up once, "one" wins 24 times and loses 30 times for a loss of 6 units in 54 unit bets, or an expected loss rate of $-6/54 = -1/9 = 11.1\%$. Similar calculations lead to Table 5-3. For the player who doesn't predict, the house edge is enormous. This is a game to avoid.

Table 5-2

2	1	40A	2	1	2
1	2	1	10	1	5
1	2	1	20	1	2
1	5	2	1	10	1
2	5	1	2	1	40B
1	2	1	5	2	1
10	1	5	1	2	1
20	1	2	1	5	2
1	10	1	2	5	1

Table 5-3

Number	House Edge
1	6/54 11.1%
2	9/54 16.7%
5	12/54 22.2%
10	10/54 18.5%
40A	13/54 24.1%
40B	13/54 24.1%

Now let's see what the player advantage might be from predictions. Suppose for the sake of discussion that the final wheel position is always within five numbers of the predicted wheel position. For any prediction in the eleven number strip centered around 40A, we should bet on 40A. In 54 spins where each final position occurred once, we will place 11 bets on 40A and win one of them for a gain of $40 - 10 = 30$ units.

The discussion is the same for 40B. For any prediction in either of the eleven number strips surrounding each 20, twenty-two numbers in all, we bet on 20. In twenty-two bets we expect to win 20 units twice and lose one unit twenty times for a net gain of

twenty units. This leaves 54-44 or ten predicted positions where we need instructions.

There are four 10s in this left-over set of ten positions. Suppose we bet the 10 each time one of these positions is predicted. It seems plausible to suppose that we would win ten units four times and lose 1 unit six times for a net gain of $40 - 6 = 34$ units. (Actually, since the 10s in this case are either the predicted number or within one position of the predicted number, we expect to do better still.

Finally, in 54 unit bets we net 30 units from 40A, 30 units from 40B, 20 units from the two 20s, and 34 units from the four 10s, for a total of 114 units/54 units or a 211% rate of return.

It may be possible to improve both the timing procedure and the method of exploiting predictability. This would improve the results.

We see now that the Caesars wheel can be predicted well enough so that we can beat it if the casino will let us put down bets after the wheel has been set in motion.