Aaron Brown

Gambler's Ruin

Ruination leads to enlightenment; tough, but true

The term "gambler's ruin" is used for a number of statistical ideas whose common denominator is predicting the eventual outcome of a series of repeated bets. There are some beautiful and counterintuitive results in which bets that seem unappealing individually lead inevitably to good outcomes, and bets that seem too attractive to refuse, lead just as inexorably to disaster.

The main gambler's ruin results, from most to least obvious, are:

- If a bet has an absorbing state and you make the bet often enough, you will reach the absorbing state. Technically we need a condition that the probability of the absorbing state is non-zero and does not decrease too quickly, but I will ignore such technicalities for this column.
- If a bet has a negative expected return and you bet enough, you will be ruined.

The next two results apply even to positive expected return bets:

- If you raise your bet proportionate to your winnings, but do not cut it when you lose, you will eventually be ruined.
- If you bet more than the Kelly amount, you will eventually be ruined.

Each of these results can be applied in either forward or inverse form. Sometimes we have good information about the statistical properties of individual bets, and can use it to predict the eventual outcomes for people who make the bets repeatedly. In other cases, we have good data on eventual outcomes, and can use it to deduce properties of the individual bets. It's a good statistical principle to study what you want to know about – individual bets if you want to know individual bet properties, or eventual outcomes if you want to predict eventual outcomes. But sometimes we have better information on the thing we are less interested in, and that is when gambler's ruin can be helpful.

What's ruin got to do with it?

The genesis of the gambler's ruin problem is a letter from Blaise Pascal to Pierre Fermat in 1656 (two years after the more famous correspondence on the problem of points). Pascal's version was summarized in a 1656 letter from Pierre de Carcavi to Huygens:

> *Let two men play with three dice, the first player scoring a point whenever 11 is thrown, and the second whenever 14 is thrown. But instead of the points accumulating in the ordinary way, let a point be added to a player's score only if his opponent's score is nil, but otherwise let it be subtracted from his opponent's score. It is as if opposing points form pairs, and annihilate each other, so that the trailing player always has zero points. The winner is the first to reach twelve points; what are the relative chances of each player winning?*

Pascal thought this problem was much harder than the earlier problems he discussed with Fermat (the ones that led to the development of modern probability theory). In fact, he was confident Fermat could not solve it.

Although Pascal's reason is lost to history, it seems to me that this problem was designed to evade the trick Fermat used to solve the earlier problem of the points. In that problem two gamblers are competing to see who can get to seven points first, and one gambler has six points while the other has five. Fermat noted that the game has to be over after two throws, and you can enumerate the four possible outcomes, even though two of them (the two in which the gambler with six points wins the first throw) would never be played to completion. Pascal's variant of the problem creates an infinite number of possible paths to the end state, so it is impossible to enumerate them. This is the philosophic core to gambler's ruin, similar to what Ludwig Boltzmann invented the word "ergodic" to describe.

Huygens reformulated Pascal's problem and published it in *De ratiociniis in ludo aleae* ("On Reasoning in Games of Chance", 1657):

> *Problem (2-1) Each player starts with 12 points, and a successful roll of the three dice for a player (getting an 11 for the first player or a 14 for the second) adds one to that player's score and subtracts one from the other player's score; the loser of the game is the first to reach zero points. What is the probability of victory for each player?*

This is the classic gambler's ruin formulation: two players begin with fixed stakes, transferring points until one or the other is "ruined" by getting to zero points. However, the term "gambler's ruin" was not applied until many years later. The theorem Huygens proved to solve this problem was an important result in the early development of probability.

Why did Huygens restate Pascal's problem in equivalent terms, terms that introduced the idea of ruin? It's sheer speculation, but my guess is he started by imagining gamblers transferring coins back and forth. This, of course, is how people gambled and it makes the problem both more natural and clearer to state. Each player starts with 12 florins and either wins or loses a florin on rolls of 11 or

14, and the game continues until one player loses his entire stake. But when it came time to write the book, Huygens decided to remove the reference to gambling for money in order to avoid giving unnecessary offence, leaving us with a formulation even clumsier than Pascal's. Whatever the reason, this is how we got the concept of gambler's ruin.

Absorbing states, oysters, and cigarettes

If you eat an oyster, there is about one chance in a million you will die from it. Therefore, if you eat enough oysters, it will kill you.

It's harder to estimate the chance that smoking a cigarette will kill you, but it is likely between one in ten million and one in two million. Anyway, it's lower than the chance of dying from eating an oyster. So, why is smoking considered irresponsibly dangerous, while eating oysters escapes attention from health zealots?

The first thing that comes to mind is risks other than death. But by this criteria, oysters are still much more dangerous than cigarettes. Even if you consider that every smoker gets sick from cigarettes, that's still only 4 to 20 sick people per cigarette death. Oysters make 200 people sick for every one they kill. On other grounds oysters beat cigarettes: they kill much faster, and they kill younger and healthier people. Measuring second-hand smoke deaths is even harder than measuring cigarette deaths, but it is certainly less than 20 percent of the number of smokers killed by cigarettes, and adding in other indirect deaths doesn't increase the percentage much. It's much smaller than the five people killed obtaining and processing oysters for every fatal poisoning of an oyster eater.

The reason we consider cigarettes dangerous and oysters safe is that people smoke a lot more cigarettes. Eating three dozen oysters a week for 50 years is 93,921 oysters, while smoking two packs of cigarettes a day for the same 50 years is 730,500 cigarettes. If each cigarette is half as dangerous smoking is 3.4 times as likely to kill you as eating oysters.

However, there is a problem of categorization here. There are many forms of tobacco use other than cigarettes: cigars, pipes, chewing tobacco, snuff, and others. People combine these when talking about the dangers of smoking, although few people do enough of any of them for gambler's ruin considerations to kick in. Considered individually, they are more like eating oysters than smoking cigarettes in that respect. On the other hand, smoking marijuana or other substances is usually separated out, and, as far as I know, no one has ever suggested combining cigarettes and oysters in a single category.

Another issue is that the chances of dying cited above are averages. Eating an oyster in a good seaside restaurant in December is safer than getting one from a careless discount stand in July, far from oceans or refrigerators. Smoking your first cigarette is likely much less dangerous than smoking your one-hundred-thousandth, or your sixty-first of the day, because some damage from tobacco is cumulative. The idea of gambler's ruin is that if you do something enough, individual odds don't matter. Unless you can reduce the chance of the absorbing state to zero, enough repetitions will send you there, whether you are careless or reckless.

Ain't no mountain high enough

A good example of that is mountain climbing. A serious attempt at a major peak carries about a 5 percent chance of dying. Mountaineers are quick to point out that deaths occur disproportionately among the least experienced climbers, and are often the result of avoidable risks such as attempting climbs without oxygen, or proceeding in the face of a bad weather forecast. The problem with that argument, however, is the more experienced and careful climbers make more attempts. About 40 percent of serious mountain climbers will die from their avocation. Some of the remainder either sustain serious injuries while climbing that prevent them from making more attempts, or die young for other reasons. The result is that serious mountain climbers who continue climbing are more likely to die from it than from all other causes put together.

The 5 percent probability of dying from attempting a climb of a major peak is about the same as the probability of dying from attempting suicide. However, from a gambler's ruin perspective, mountain climbing is about four times as deadly as suicide, since mountain climbers make more climb attempts than suicide attempters try to kill themselves. About 10 percent of people who attempt suicide succeed eventually, compared to the 40 percent of mountain climbers who die in a climbing attempt.

This raises two related questions. First is whether we care about individual bet odds or eventual outcome. For mountain climbing, the answer is probably eventual outcome. You don't climb just one major peak. You choose a lifestyle that requires hard physical training, acquisition of difficult skills, purchase of expensive equipment and time commitment that eats into social and professional development. More important than any of that is you will accept a set of values distinct from those of mainstream society, and join a like-minded group. A significant chance of dying on a mountain is part of the admission price, regardless of the specific dangers of your current attempt, or whether you are personally more or less careful than average.

Eating oysters, on the other hand, is something most people can take or leave. It requires no fixed investment, no skill (other than the ability to order "a half dozen Little Skookums and some Naked Roy's" without laughing), no adjustment of values. It makes sense to evaluate each oyster one at a time, is the pleasure of eating it worth the risk of vibrosis, norovirus infection and other shellfish poisonings? Smoking is somewhere in between, taking an occasional cigar is more like eating oysters, regular cigarette smoking more like climbing mountains.

The related question is how to estimate the risk of the activity. Do we ask, "What fraction of people who attempt to climb Annapurna or Siula Grande each year die in the attempt?" or, "What fraction of people who have ever climbed an 8,000 meter peak die in their beds from old age?" While this will always depend partly on the data we have available, and in general we would integrate both kinds of information to come to an informed judgment, the general statistical principle is to study what you care about. Mountain climbers should look around to see if there are any old people like themselves, oyster eaters should check the calendar for an "R" in the month.

This discussion of the simplest gambler's ruin principle, that everything with non-zero probability will happen if you wait long enough, has identified four issues: how do we categorize events, how do we average probabilities, do we care about individual bets or eventual outcomes and do we estimate probabilities by individual bets or eventual outcomes? We're now going to move on to less obvious gambler's ruin principles. The same issues will arise in more complicated forms.

Expect the unexpected

The next most obvious form of gambler's ruin is that if you repeat a negative expected value bet often enough, you will eventually be ruined. This is the form of the principle that makes roulette systems hopeless. You can switch from red to black, raise and lower your bet, make whatever stopping rule you want, but your long-term average result is stubbornly fixed at a loss equal to the house edge times the total amount you bet.

If someone wins a multimillion dollar prize in the lottery, it is a news event. If someone buys a ticket and does not win a multimillion dollar prize, it is not. So judging from news accounts could lead to the conclusion that lottery tickets are good investments. But we know the organizers pay out only a small fraction of ticket prices in prizes, so the expected value of a ticket is negative.[1] If we average over a large enough number of tickets, we know we'll find more losses than gains.

If you bet more than Kelly for long enough, you don't go broke in the sense of getting zero bankroll, because you never get to zero

Similarly, we know that the average investor in S&P500 stocks earns the return on an S&P500 index fund, minus whatever extra fees, expenses and taxes the account runs up. Active trading in financial markets carries a large cost. That doesn't mean everyone who trades actively loses money from it, but it does mean the losses outweigh the gains.

This matters outside a casino and financial trading when we care about long-term outcomes, but don't have long-term evidence.

Speaking of casinos, it is a natural tendency of bettors to increase their wagers as their bankrolls expand due to winning. After all, when luck is running their way, it's foolish not to take advantage. And if you've won \$10,000 in an evening, a \$100 bet doesn't carry much thrill any more. You're playing with the house's money at that point, so why not play? Casinos encourage this strategy by "chipping up," replacing your lower denomination chips with higher ones.

Losing is not symmetrical. It's true that runs of bad luck cause some people to reduce wagers, but most people instead "chase," raising bets in order to get back to peak bankroll, or later in the evening, to get back to even.

Suppose, for example, you always bet 5 percent of your peak bankroll. In that case, 20 losses in a row will wipe you out, and this will happen sooner or later, even if each bet has a positive expectation. In fact, any run of bets of any length with 20 more losses than wins will wipe you out.

This ruinous strategy is not confined to casinos. Successful risk taking in investments, sports, romance and almost everything else pushes most people to increase their wagers. Losses are not as effective at encouraging retrenchment.

Historically, this was the first result to be called "gambler's ruin," more than 200 years after Huygens publication.

Kelly's heroes

The least obvious version of gambler's ruin applies to positive expected return bets and symmetrical bet sizing decisions. Suppose, for example, in a casino promotion the house offers to reverse the odds in roulette so the bettor has the edge. If you bet on red or black, you win if either of the green zero or double zero come up, or if you get your color. Thus your odds of winning are 20/38 for an even payout game.

Clearly you have a positive edge of 2/38 = 5.26 percent in this game. We know repeated play leads to large predictable profits for the casino when it has the edge. So how much should you bet if you can play this game for a long period?

Suppose you decide to wager 20 percent of your bankroll on each spin. Your expected profit is 20 percent x 5.26 percent = 1.05 percent of your bankroll each spin. After 220 spins, you expect to have ten times what you started with. The trouble is that 93 percent of the time you end up with less than ten times your starting wealth, and 78 percent of the time you actually lose money. Nearly half of your expected gain comes from the 0.1 percent of the time when you get more than 1,000 times your starting bankroll.

It's volatility drag that gets you. Every time you win and then lose (or lose and then win) a bet, you lose 20 percent squared = 4 percent of your bankroll. With average luck, you'll win 116 spins and lose 104, which you can think of as 104 matched wins and losses plus 12 unmatched wins. The 104 matched wins and losses reduce your bankroll to $0.96^{104} = 0.014$ of its initial value. The 12 unmatched wins multiply your bankroll by $1.2^{12} = 8.9$. But 0.014 x 8.9 = 0.13, meaning you lose 87 percent of your bankroll if you get the expected outcome.

If you bet more than Kelly for long enough, you don't go broke in the sense of getting zero bankroll, because you never get to zero. Your probability of losing money, however, goes to 1. Your expected value climbs to the stars, but it becomes composed of ever more microscopic chances of winning ever more astronomical amounts.

Kelly bets 2/38 of bankroll each spin, which leads to an expected return of 5.26 percent x 5.26 percent = 0.28 percent each spin. After 220 spins, you expect to have $1.0028^{220} = 1.84$ times your initial stake instead of the 10 times you get from betting 20 percent each spin. But if you get the expected 116 wins you'll have 1.39 times your initial bankroll, instead of the 0.13 you'd have from betting 20 percent each time. You'll have a profit after 220 spins 67 percent of the time, and beat the guy who bets 20 percent, 85 percent of the time.

Whether you are a gambler or a mathematician or, like most readers of *Wilmott*, both, you need to think a lot about ruin. It's an important way to link local uncertainty to eventual predictability, which is the key to a lot interesting stuff.

ENDNOTE

1. This does not rule out the possibility that some lottery tickets have positive expected value, for example on days with abnormally large jackpots or in scratch off games with known patterns in the cards.