Introduction to Gaming Mathematics

Eliot Jacobson, Ph.D.

A brief introduction to the mathematics of casino games and other games of chance, with applications to game development, casino marketing and table games management.

Jacobson Gaming, LLC

PO Box 41532
Santa Barbara, CA, 93140

Introduction to

Gaming Mathematics

by

Eliot Jacobson, Ph.D.

Introduction

The opportunity to play a casino game is a product that the consumer purchases with his gaming dollar. The view I take is that the player wants to have a certain experience when he sits at a game and is willing to pay for that experience by means of the inevitable bite the house edge takes. To fully understand the product the player is purchasing requires a journey into the mathematics of casino games. Marketing must understand this material when developing effective strategies.

Every part of a player's interaction with a casino game has its own mathematics. Wagers have "odds" that determine how much to pay the player if he wins. The percentage of time the player wins is called the "hit frequency." The "house edge" is the fraction of the player's total initial wagers that the house expects to keep over the long run if the player uses perfect strategy. The "theo" or "theoretical" is the actual dollar amount the casino expects to win from the player. The "return to player" measures the fraction of the player's total wagers the player is paid back in live casino play.

X-for-1, X-to-1, 1-in-X

Most table game wagers are paid in an "X-to-1" fashion. What this means is that if the player wins the round, his wager is returned and he is paid an additional amount for his win. If the player loses, then his wager is removed from the table.

For example, if the player bets on red-7 straight-up in Roulette then the payout is 35-to-1. If his wager is \$10, then he either wins \$350 (35 × \$10) and his \$10 is returned, or he loses \$10.

This type of pay is different from slot machines. For slots, every time the player pushes the button or pulls the handle, an amount is deducted from his credit balance. That's his wager. Then the reels spin, and when they stop he is paid an amount based on the pattern displayed by the reels by reference to some pay table. No matter what the outcome of the reels, the player's original wager is forfeited. We say that slots are "X-for-1."

Most wagers on keno, lotteries, slots and video poker are paid "X-for-1." This confusion has been used effectively, for example, in blackjack slot machines. Blackjack slots typically show a payout of 2 credits when the player is dealt a blackjack. The player knows that blackjack usually pays 3-to-2, so this seems like a sweet deal. However, this payout is 2-for-1. Thus the player forfeits his original 1 credit and then wins back 2 credits if he is dealt a blackjack. The player nets 1 credit in this way. Thus 2-for-1 is the same as 1-to-1. Few would play the table game of blackjack if the payout for a blackjack was reduced to 1-to-1, yet people often play these slots with a payout of 2-for-1.

The progressive wager in Caribbean Stud is an example of an "X-for-1" wager appearing in a well known table game. The player's wager is automatically collected as it is made, and a light illuminates on the table indicating the player has made this wager. Another common "X-for-1" wager is the bad-beat jackpot in the poker room. The pot contributes \$1 each hand to the progressive; that money is forfeit.

It is customary that wagers pay out "X-to-1" on a table game; examples of successful "X-for-1 wagers on table games are rare (one example are the hop bets in craps). Return the player's original wager if the player wins or pushes; otherwise collect the player's wager. Table games and slots are different animals and the most fundamental difference is that slots pay "X-for-1."

Some wagers in table games are very similar to slots. For example the Pair Plus wager in Three Card Poker has a straight pay table based on the three cards dealt to the player. It usually pays "1-to-1" for a pair, "3-to-1" for a flush, "6-to-1" for a straight, "30-to-1" for trips and "40-to-1" for a straight flush. This side bet could just as well be a slot, as it requires no strategy, but it would then have the rather strange looking pay table: "2, 4, 7, 31, 41."

Finally, the terminology "1-in-X" is commonly used, but not in the context of payouts. For example, in craps, the shooter will get a 7 (over the long run) 1-in-6 rolls. In double-0 Roulette, the number red-7 will come up (over the long run) 1-in-38 spins. The player is dealt a blackjack (over the long run) approximately 1-in-21 hands. The terminology "1-in-X" has nothing to do with how much a game pays for a winning wager. It is used to describe the rate of occurrence of an event at a table game.

Hit Frequency

The concept of "hit frequency" is best applied to games that pay "X-for-1" like keno, lotteries, slots and video poker. It is a little fuzzier when applied to table games, especially those with a strategy or multiple wagering opportunities.

If the player is playing keno, the lottery, slots or video poker then the player makes a bet (say $1). That $1 is forfeit. It's gone. Now the only question is if he will get something back after the round has resolved. Sometimes yes, sometimes no. In the real world, nothing is exact. But to the mathematician analyzing these games, they can determine very exact numbers. We'll use Keno as an example.

Suppose you are playing Keno and you mark 10 spots on a card with 80 spots. Suppose 20 balls are dropped in the Keno game. This is all very standard, exactly how Keno usually works in many casinos.

Now suppose the pay table for Keno awards some money back if you hit 3 or more of the 10 spots you marked. You always lose your original bet. It's gone. The question is if you get to collect anything at the end of the ball drop. It's time for decimals and percents. Don't ask me how I figured it out, but for this exact situation the hit frequency (probability of 3 or more hits) is 47.94%, or 0.4794. If we modify the Keno pay table so that it pays back only when you hit 4 or more of the 10 spots you marked, then the hit frequency is 21.20%, or 0.2120. If we have a payout any time 2 or more of the 10 spots are hit, then the hit frequency is 77.46%, or 0.7746.

For the Pair Plus wager in Three Card Poker, the hit frequency is 25.61%, or 0.2561. That is, about 25.61% of the time the player will be dealt three cards that give a pair or higher.

The "hit frequency" is the rate of occurrence of a winning result during the play of a casino game. It is often given as a percent.

For example, to say that the Pair Plus wager has a hit frequency of 25.61% means that over the long run, the player can expect to have a winning result approximately 25.61% of the time. If the player plays 10,000 hands of the Pair Plus wager, he can expect to get a pair or higher (to win) on approximately 2,561 of those wagers. The more rounds the game gets played, the more accurately the following equation describes the hit frequency:

$$\text{Hit frequency} \approx (\text{Number of wins}) / (\text{Number of rounds}).$$

For the Pair Plus wager, the number 25.61% is an exact number. The hit frequency is a theoretical number; it is computed by a mathematician. In practice, it gives a good sense for how the game will perform on the floor, but the game's actual performance has nothing to do with determining the hit frequency.

The hit frequency gives a sense for how often a player is going to get something back on his wager. Obviously a player wants to win as often as possible, so from the player's perspective a high hit frequency is a good thing. Naturally, the player also wants high payouts. These two desires act against each other: giving the player everything he wants is impossible.

A game that has a low hit frequency with high payouts may frustrate the player because he seems to lose and lose while waiting for a combination to occur that seems ever more unlikely. A game that has a high hit frequency with low payouts may frustrate the player because he feels like he is going nowhere fast; it becomes obvious how the house is making its money. The answer is always somewhere in between.

We are going to create two lottery games to illustrate these psychological points. A low hit frequency, high payout lottery and a high hit frequency, low payout lottery. We've designed these as extreme examples just to drive home the point about how this feels from the player's perspective.

Suppose we've been given the task of creating a lottery with one million tickets; the cost of each ticket is one dollar. We know that if we sell every ticket, then we will take in $1,000,000 gross income. We are directed to return $900,000 in cash prizes, leaving a net profit of $100,000. We consider two prize schemes.

Example: Lottery 1 (low hit frequency, high payout): One million tickets are produced. Each costs $1 to play. Out of these 1 million tickets, suppose exactly 9 of them pay 100,000-for-1 and the others lose. That is, 9 of them say "You win $100,000!" and the other 999,991 of them say "Please play again!" The player will consider the $1 purchase of the ticket measured against the possibility of winning $100,000.

In this case the hit frequency is 9/1,000,000 = 0.0009%.

Example: Lottery 2 (high hit frequency, low payout): One million tickets are produced. Each costs $1 to play. Out of these 1 million tickets, suppose exactly 900,000 pay 1-for-1 and the others lose. That is, 900,000 of them say "You win $1!" and the other 100,000 say "Please play again!" The player will consider the $1 purchase of a ticket measured against the possibility of winning $1.

In this case the hit frequency is 900,000/1,000,000 = 90.00%.

It is easy to see that each lottery makes the same amount of profit if every ticket is sold. Lottery players like the opportunity to win big, so lottery 1 is more likely to succeed in the marketplace. But if lottery 1 was a casino game, player after player would lose his entire buy-in a matter of minutes (even playing just $1 per round). A few players would win big eventually, but that would be extremely rare. The game would die soon after it was put on the floor. If a casino game was designed in lottery 2 style, players might play for the entertainment and free drinks, but that game too would not last long. Players of lottery 2 have no chance to ever get ahead of the game.

Hit frequencies, as described above, are precisely defined for keno, lotteries, slots and video poker. This term also has this meaning for certain wagers that pay for a specific hand or event at a table game. As we saw above, it makes sense to discuss the hit frequency for the Pair Plus wager in Three Card Poker. What does "hit frequency" mean for a more complicated wager involving strategy at a table game?

Suppose you are playing blackjack. You place a $10 bet and are dealt 8-8 against the dealer's 6. You split, and re-split to four hands. You double on three of the four splits. You put out a total of $70 on the hand. A lot of other stuff happens and you look down at the end at what you've been paid. Either there is more than $70 in front of you, less than $70, or you managed to get a big push and have $70 returned to you. What does "hit frequency" mean in a situation like this? Honestly, I have no idea. No matter how ill-defined, hit frequency has become part of the vernacular of table games. Some features of games have hit frequencies that can be computed. In many cases, hit frequency just doesn't make sense according to the formal definition.

House Edge

Figure 1 indicates the house edge (or "house advantage," or "H/A") for some well known games. The house edge is commonly understood as the percentage of total initial player wagers the casino expects to earn over the long run if all players use perfect strategy against the game. For those games that have exactly one wagering opportunity, the house edge refers to the edge for that wager. This is the case with the percents given for Casino War, and roulette, as well as keno, slots and video poker games. If a table game has a main game and a side bet, then it is usually the house edge of the main game that is quoted as the house edge for the overall game. For example, Three Card Poker has both the Ante/Play and the Pair Plus wagers; the stated house edge is for the main game, the Ante/Play bet. Blackjack and craps also have more than one wagering opportunity.

House Edge for Common Table Games	
Game	**Edge**
Blackjack	0.46%
Casino War	2.90%
Craps (Pass line)	1.41%
Let it Ride	3.51%
Pai Gow Poker	2.66%
Roulette	5.26%
Three Card Poker	3.37%

Figure 1.

The house edge for blackjack is based on the specific rules for the game played; it is not a number that is valid for all blackjack games. For the percent presented in Figure 1, blackjack is played with 8 decks, the dealer stands on soft 17, the player can split twice (up to three hands) and double after a split. Split aces receive one card each, and the player cannot surrender.

Blackjack actually has a second wagering opportunity that very few players use and it is not considered in the house edge. Players, dealers and casino management don't think about insurance properly; insurance is a separate wagering opportunity in blackjack that is independent of the main bet. It is a side bet. It is a wager that is available to the player any time the dealer shows an Ace. The player is paid 2-to-1 on the wager if the dealer has a face card under the Ace, otherwise the player loses. There is no theoretical link between the main bet and the insurance bet. The house edge for the insurance wager on an 8 deck game is 7.47%.

For craps it is not possible to give an overall house edge. Craps is a hodge-podge of wagers, some good and some bad. If we focus on one wager then we can indicate the house edge. The pass line wager has a house edge of 1.41% and that's what is included in Figure 1. The "Any Craps" wager has a house edge of 11.11%.

For Three Card Poker, there are two wagers: the Ante/Play bet and the Pair Plus bet. The Ante/Play bet has a house edge of 3.37% (using the 1-4-5 bonus pay table) and the Pair Plus bet has a house edge of 7.28% (using the 1-3-6-30-40 pay table). For some games, the player has to play perfect basic strategy to get the house edge. For example, blackjack, Let it Ride, Pai Gow Poker, and Three Card Poker all are games with a player strategy. Casino War, craps and roulette have no player strategy, so the house edge is inescapably the long-run experience the players will have. Any variations in this experience are simply expressions of the volatility.

Note that even if a strategy is trivial, there will still be player errors. In Three Card Poker, correct strategy is to raise with Q-6-4 off-suit or higher and to fold all other hands. The strategies for blackjack, Let it Ride and Pai Gow Poker are complex. The more complicated the strategy, the more errors the players will make. Some errors are very costly on simple games. Error rates translate to returns to the casino higher than the "theoretical win" predicted by the house edge.

For slots and video poker, the house edge is also called the "hold" of the game. The use of the word "hold" as a synonym for house edge causes great confusion. The word "hold" has an entirely different meaning for table games, to be discussed later. The house edge for each slot is set by the casino by referencing several optional pay tables supplied with the slot.

Pay tables and other features of a slot are created by a mathematician who analyzed the game and built a spread sheet for the slot machine, called a PARS sheet (PARS = "Pay Table and Reel Strip). A PARS shows the game manufacture how to create the machine by indicating exactly how each reel strip is set up and how much each combination pays when it is hit. A common misconception is that somehow slots are set up to be hot or cold at certain times. In fact, each location on each reel is equally likely to occur on each spin, and this doesn't change. The changes are set in advance by engineering the design of the reels and the payouts for combinations.

Typically a casino can choose from a menu of possibilities and get a house edge for a slot anywhere from 2% to 20%. Slots that have a house edge under 4% are called "loose." Slots that have a house edge of 8% or higher are often referred to as "tight."

For video poker games, the game is dealt from a single deck of cards. By law, the deal of these cards must behave statistically in every way just like the game as dealt live. There can be no weighting of the cards or other deviations. Each card must be equally likely to be dealt any time it remains in the deck.

Video poker comes in a wide variety of versions. The most common and most popular variants are Jacks or Better, Bonus, Double Bonus, Deuces Wild, and Joker Poker. In each case the house edge is determined by the payouts for various poker hands.

For example, playing Jacks or Better, there is a payout for any final poker hand that consists of a pair of Jacks or higher. Figure 2 gives the pay table that knowledgeable video poker players seek out. The player must play 5 credits to get the listed payouts. Note that these payouts are "for-1." The player's credits are deducted from his total before the hand begins. For those who wager less than “max coin” the player receives a lower payout for a Royal Flush, typically 250 instead of 800.

Common Pay Table for Jacks or Better VP	
Hand	**Pays (for 1)**
Royal Flush	800
Straight Flush	50
Four of a Kind	25
Full House	**9**
Flush	**6**
Straight	4
Three of a Kind	3
Two Pairs	2
Jacks of Better	1
Nothing	0

Figure 2.

The house edge for this pay table is 0.46%, based on the player playing perfect strategy. Perfect strategy is quite complicated and all but a few experts play a considerably worse game. In practice, a house edge of about 1% is earned by the casino. This pay table is referred to as Jacks or Better “9/6,” with the 9/6 referring to the payouts for a Full House (9) and Flush (6). The most common pay table for Jacks or Better is 8/5, which has a house edge of 2.70% playing perfect strategy. At some casinos you will see a 7/5 (house edge = 3.85%) or even a 6/5 (house edge = 5.00%) pay table for Jacks or Better. I have never seen a pay table tighter than 6/5.

For most video poker games, pay tables exist that have close to a 0% break even house edge, or in some cases, have a slight player edge. For example, 10/6 Jacks or Better (player edge = 0.70%) can occasionally be found. Such video poker games are offered at a low denomination and may be placed close to the front of the casino, as a loss leader. In some cases, professional players will target video poker games that offer these great odds. It should be clear that if a casino is using a “cash back for points” marketing system, together with offering a video poker game with a break even house edge, the expert player will turn the house edge well in his favor.

The house edge is an exact number that can be worked out precisely for most games. For many games, the edge can be determined by using a spread sheet. For some games, the house edge is determined by doing an exhaustive analysis of every possible combination that can occur in the game. This may require extensive spread sheet analysis. Typically the mathematics for slot machines is done this way. The analysis may require a computer to go through each combination of events, work out the proper strategy for that combination, and keep track of the overall results. This is commonly referred to as “running a cycle.” Video poker is analyzed by running a cycle.

Whenever possible, the house edge is obtained by a complete spread sheet analysis or by running a cycle (or both, just to be sure). A computer simulation is often used to confirm the exact theoretical results. For a few games, the toughest games, the rules and strategy are so complex that the house edge can only be approximately determined by computer simulation.

Once casino management knows the house edge, they will have a sense for how much money they can make from the game based on the total wagers. But, the house edge tells other stories as well. If the house edge is too close to 0%, then the casino won't make enough money from the game. If the house edge is too high, then players will lose their money too fast and abandon the game, and once again the casino won't make enough money from the game.

The house edge is another balancing act, just like the "hit frequency" and "payout" dilemma described in the lottery examples. The house edge has to be low enough to make players want to play the game and high enough so that the casino earns a reasonable amount of profit from the game.

The standard table games consisting of baccarat, blackjack and craps, all offer wagers with a house edge under 1.5%. There are very few table games that make it into a casino with a house edge under 1.5% and those that do are mostly blackjack variants. Well known examples are Spanish 21, Super Fun 21, and Blackjack Switch. For example, with liberal rules and perfect play, Blackjack Switch has a house edge under 0.20%. In practice, due to player errors, Blackjack Switch often returns in excess of 3%. Spanish 21 and Super Fun 21 also have huge error rates. Blackjack variants also move very quickly, generating many more decisions per hour than other games, making the games affordable for the casino.

Likewise, for table games, the edge can be over 5%, but the game better move very slowly and have a compelling reason to play. Otherwise, a house edge over 5% is likely to burn players out too fast. The game Mississippi Stud has a house edge close to 5% and a huge error rate. It is relatively new and popular in some markets, but does not seem to be catching hold. Caribbean Stud has a house edge over 5.2% and has been steadily losing market share for years.

Roulette maintains its popularity, even though every wager at roulette has a fixed house edge of 5.26%. The house edge can easily be computed by hand as follows. First note that without the 0 and 00 spots on the wheel, the payouts for roulette are designed to be absolutely fair based on a wheel with 36 spots: a 0.00% house edge. By putting both 0 and 00 on the wheel, there are two bad spots on the wheel out of 38 total spots. So, for double-0 roulette, the house edge is 2/38 = 5.26%. For single-0 roulette, the edge is 1/37 = 2.70%.

It is not quite true that every roulette wager has the same house edge; there is an exception that we note. Double 0-roulette sometimes offers a wager called "first 5" that pays 6-to-1. This wins if any of 0, 00, 1, 2, or 3 hits. It loses if any other number is hit. The house edge for this wager is 3/38 = 7.89%. The great majority of roulette players know better than to take this bet, just as most blackjack players know to never take insurance.

For games with side bets, it is usually the case that the main game has a lower house edge than the side bet. Moreover, the lower the edge on the main game, the higher the house edge on the side bet. Side bets do not have the same practical limitations in the size of their payouts or house edge as the main bet. A 3% edge on the main game can easily go with a 7% house edge on the side bet, for example.

As an example, consider the Pair Plus bet on Three Card Poker. The main game in Three Card Poker (with the standard 1-4-5 bonus table) has house edge 3.37%. Shuffle Master includes pay tables for the Pair Plus with house edges varying from 2.32% to 7.28%. When Three Card Poker was first introduced, the 2.32% version of Pair Plus was most common. Now the 7.28% version is the main one used. There is logic to this increase; from the casino's side they earn a lot more from a well established and loyal player base; from the player's side the change is just "cosmetic" and they don't notice it.

For our lottery examples, it is easy to compute the house edge by running a cycle. That is, we see what happens after every ticket is sold for each lottery. Whether it is lottery 1 or lottery 2, there are 1 million tickets, and each cost $1. So a full cycle corresponds to an investment of $1,000,000 (1 million dollars) from the players. For each lottery, $900,000 is returned in prize money: there are either 9 tickets worth $100,000 each, or 900,000 tickets worth $1 each. For each lottery, the state keeps the remaining $100,000. Thus, for each lottery, the house edge is ($100,000) / ($1,000,000) = 10%.

Theoretical Win (Theo)

Of all the numbers that describe a casino player, the one that is most often sought out by marketing is the theoretical win (theo). This number describes the expected house win for a player, based on the total amount wagered by the player and the game played. To compute the theo, only two numbers are needed: the house edge of the game and the total amount wagered. There may be technical problems with getting accurate values for these numbers, but at its core, the formula for theo is simple:

$$\text{Theo} = (\text{house edge}) \times (\text{total amount wagered}).$$

The theoretical win for players who play Casino War is (0.0290) × (total amount wagered). The theoretical win from players who place a wager on the pass line in craps is (0.0141) × (total amount wagered).

As a specific example, if players wager $1 per spin on roulette, for one million spins, the house expects to earn about 0.0526 × $1,000,000 = $52,600. Some players will win big in a session, other will lose big. Some will have mediocre results up or down. Whatever happens in practice, the theoretical win for the casino can always be computed using the house edge and total amount wagered.

A player's theo is the fundamental number that is used as a basis for assessing the worth of the player. This number has nothing to do with the actual win or loss of the player in a session or over a period of time. Promotions, comps, hosting and other reinvestment strategies are often based primarily on a player's theo measured over some period of time. Theo is based on the history of a player's action, and is used, along with demographic and metrical information, as a forecast of future action. It is the key dollar value that drives reinvestment in the player.

For example, if a player plays roulette for 4 hours, at $20 per spin, and there are 30 spins per hour, then the player has wagered a total of 4 × $20 × 30 = $2,400. Regardless of what actually happened at the table with the player, the casino uses the house edge to compute the player's theo. In this case, the player's theo is $2,400 × 0.0526 = $126.24. If the casino is reinvesting at a rate of 20% of theo, then the casino has 20% x $126.24 = $25.24 set aside to reinvest in the player. The hypothetical roulette player above may be given a free buffet lunch. Management knows that in theory the casino has won a lot more than the cost of a lunch from this player, even if the player won $1,000 in the session.

For slot machines and keno, a player's theo can be precisely determined. For those games that require strategic decisions, like blackjack or video poker, a player will rarely play well enough to get the theoretically optimal house edge. In craps, the pass line wager has a house edge of 1.41%. However, most players will make additional bets (like "place" or "buy" or "proposition" wagers) that have a higher house edge. In combination, most players play craps giving up a house edge in excess of 2%. For this reason, there is a bit of "art" to computing the practical house edge for many table games as well as video poker. Without an exact value for the house edge, it is impossible to correctly compute a player's theo. Figure 3 gives values that are used by some casinos for the house edge for popular table games.

House Edge Used for Computing Theo	
Game	**Edge**
Blackjack	2.00%
Craps	2.00%
Roulette	5.26%
Three Card Poker	3.50%
Roulette	5.26%

Figure 3.

If the casino uses values for the house edge for table games that are too high, then marketing will get a value of a player's theo that may be larger than it actually is. This may lead to over investment in a player and in some cases the casino can become upside down to the player. Over investment gives players a sense of entitlement. It is very difficult to repair a marketing program that has over invested in players for an extended period. Loyal players may wonder why they suddenly are worth less, causing animosity and possibly leading to defection.

If the casino uses values for the house edge for table games that are too low, then marketing will underestimate a player's theo. This may lead to under-investment in a player. Players who are underrated will feel like the casino is cheap and may go in search of a casino that treats them better.

The problem with either over or under-estimating a player's theo is that players will seek out the casino that treats them best. This creates market pressure on a competitive group of casinos to move towards the casino with the most generous marketing model. This pressure to keep up with a competitor is a challenge that must be accurately analyzed. Don't blindly follow the leader. It's best to create a model that gets things right, then stick to a plan based on those numbers.

<u>Return to Player (RTP)</u>

The "Return to Player" or RTP for a casino game is defined by:

RTP = (Total amount won) / (Total amount wagered).

Note that this number is based on actual wagers and actual winnings. It is not a theoretical number. It is not computed by a spread sheet computation, by running a cycle, or by extensive game simulation. The RTP is based on real people placing wagers on games in a live casino environment.

For games that have no strategy (like Casino War, craps or roulette) it is usually the case that

$$\text{RTP} \approx 100\% - (\text{House Edge}).$$

Moreover, the greater the number of times the game is played, the more the RTP represents information about the actual return of the game and the more "exact" we expect this equation to become. "Exact" does not mean that the two sides get closer and closer to each other. What it does mean is that the ratio of the two sides of this equation gets closer and closer to 1 as the number of rounds gets larger and larger.

$$[100\% - (\text{House Edge})] / \text{RTP} \rightarrow 1.$$

The number 100% – (House Edge) is referred to as the game's theoretical RTP. Sometimes the language gets loose here, and the value 100% – (House Edge) is referred to as the game's RTP.

Let's look some examples.

For roulette, the RTP is 100% - 5.26% = 94.74%

Consider a player who is shooting craps and betting the pass line with no odds for a house edge of 1.41%. Because there is no strategy in craps, the RTP for the pass line bet is 100% - 1.41% = 98.59%.

For Jacks or Better with the "9/6" pay table, the RTP is 100% = 0.46% = 99.54%. With the "8/5" pay table, the RTP is 100% - 2.70% = 97.30%.

Loose slots have a RTP above 96%. Tight slots have a RTP below 92%.

Let's go back to those lotteries and investigate the RTP. In each lottery there were 1 million tickets. Let's assume that each of those 1 million tickets was purchased. The total amount wagered by the players is $1,000,000. The total amount won by the players is $900,000: either 900,000 tickets that paid $1 each or 9 tickets that paid $100,000 each. Thus the RTP if we run a full cycle is:

$$\begin{aligned} \text{RTP} &= (\text{Total amount won}) / (\text{Total amount wagered}) \\ &= (\$900{,}000) / (\$1{,}000{,}000) \\ &= 90.00\%. \end{aligned}$$

Assume we offer lottery 2 and sell 500,000 tickets. Thus, the total amount wagered is $500,000. We know that exactly 90% of the tickets are winning tickets. So, based on knowing the house edge, we expect the players to win approximately 90% of $500,000, or about $450,000. The state has a theoretical win of $50,000 to purchase shoes for the kids. For lottery 2, this income is pretty certain because of the high hit frequency (hit frequency = 90%) and the low payout of the winning tickets. The theoretical RTP of 90% transfers to a fairly accurate expectation about the real world.

Now, assume we offer lottery 1 and sell 500,000 tickets. Again, the total amount wagered is $500,000. There are only 9 winning tickets in the lot of 1 million tickets, so somewhere between 0 and 9 of the winning tickets will be sold. It is most likely that either 4 or 5 winning tickets were sold. Slightly less likely is that either 3 or 6 winning tickets were sold. And so on, with the most improbable case being that either 0 or 9 winning tickets were sold.

Based on the house edge of 10%, we expect the players to win approximately 90% of $500,000 or $450,000. Because the winning tickets are each worth $100,000 and there are only 9 winning tickets, the players actually won some unknown hundreds of thousands of dollars, from $0 to $900,000. The most likely case will be that the players won either $400,000 or $500,000 (4 or 5 winning tickets). There is simply no way to win 90% of $500,000 ($450,000). If we sold 4 winning tickets, the RTP is ($400,000) / ($500,000) = 80%. If we sold 5 winning tickets, the RTP is ($500,000) / ($500,000) = 100%. We can never have an RTP of 90% if we sell exactly 500,000 tickets.

We now consider some actual casino examples. If the player shoots one million rounds of craps betting the pass line, and he wagers $1 on each round, then at the end of those one million rounds the player has wagered $1,000,000. The player has "theoretically" won back somewhere in the neighborhood of $985,900 of those wagers. Because winning and losing on the pass line happen about equally as often (like lottery 2), it's a safe bet that the player's actual win and theoretical win are going to be pretty close, so the RTP should be pretty close to 98.59%. The more the game is played, the closer the RTP will get to 98.59% on the pass line.

Suppose a player is playing roulette and is betting $1 straight up on red-7. In this situation, the house edge is 5.26%. Therefore we expect the RTP to be approximately 94.74%. If this player bets $1 on red-7 on one million spins of the roulette wheel then at the end of those spins the player has wagered $1,000,000 and has "theoretically" won back approximately $947,400. Because hitting red-7 is an event that doesn't happen very often (1-in-38), it is reasonable to expect the player's actual win and theoretical win are substantially different (like lottery 1). Now suppose a player wagers that same $1,000,000 one wager at a time on “even” instead of red-7. In this case, because winning and losing happen about equally often, the RTP will quickly converge to 94.74% for this player.

The point is that the theoretical RTP should be pretty to the actual RTP for games that require no strategy and have a high hit frequency. The more rare a winning event (the lower the hit frequency), the more play is required for these values to converge to each other.

For games that have a strategy, mistakes will be made by the players. Because of this, the value of the RTP is usually less than 100% – (House Edge). For games with relatively trivial strategy (like Three Card Poker), the RTP is usually only slightly less than 100% – (House Edge). For extremely complicated games, like Blackjack Switch or Super Fun 21, the RTP can be several percentage points below the theoretical value.

Many casinos (especially online casinos where every player interaction is logged) create monthly RTP reports so they can track the exact performance of each game. Figure 4 gives examples of some RTPs for a month from an online casino based on live play. Looking carefully at this table, we can make some observations.

Return to Player				
Game	**Rounds**	**Wagered**	**Returned**	**RTP**
Blackjack	242,002	$6,625,064.50	$6,528,754.00	98.55%
Casino War	74,503	$561,650.00	$555,023.00	98.82%
Craps	345,566	$2,795,163.00	$2,692,950.40	96.34%
Let it Ride	209,548	$1,753,779.00	$1,636,824.00	93.33%
Pai Gow Poker	13,910	$214,870.00	$207,951.00	96.78%
Roulette	148,826	$2,164,293.00	$2,039,874.00	94.25%
Three Card Poker	290,825	$4,129,708.00	$4,006,709.00	97.02%

Figure 4.

Note that blackjack, Let it Ride, Pai Gow and Three Card Poker are all returning below the theoretical RTP for perfect strategy. Based on the results given in Figure 1, blackjack has a theoretical RTP of 99.54%, Let it Ride has a theoretical RTP of 96.49%, Pai Gow Poker has a theoretical RTP of 97.34%, and Three Card Poker has a theoretical RTP of 97.15%. On the other hand, Casino War is above the theoretical RTP of 97.10% and roulette is close to the theoretical RTP of 94.74%. The RTP for craps cannot be analyzed because of the variety of wagers available.

Be careful of a common fallacy here. The cards, dice, reels and wheels don't know what theoretical RTP they are supposed to be and try and get there. If the RTP is 94.25% at Roulette after 148,826 spins, there is no expectation that the game will try and make that up in the next 148,826 spins, so that the RTP moves closer to the theoretical value of 94.74%. If the pass line in craps is returning 102%, that does not mean the dice are due to go cold. The actual RTP is not an indication of what's bound to happen in the future to make up the difference. Or, as they say in mutual funds and other investments, past performance is not an indicator of expected future performance.

Drop, Win and Hold

In many casinos, cash wagers are not permitted for table games. Instead, the circle of life is to give the casino some cash, get chips, play for a while, win or lose some chips, and finally exchange the chips in your possession for cash at the cashier's cage. Or, as a friend of mine would say to the pit boss, with her cute deer-caught-in-the-headlights look: "you mean I can actually get cash for these things, I never did that before?"

When a player puts his money on the felt in exchange for chips, that money is placed into a metal box next to the table. The cash that goes in doesn't come out until much later, when it is counted. The total money at the table that is exchanged for chips is called the "drop" for that game. If more than one table of a game is offered, the drops are often totaled together to get an overall drop for that game in the casino. Casinos rarely consider drop for a single player. Instead, the drop is totaled over all players of a game (or all players for all tables of a game), for a shift, a day, a week, a month, a quarter, or a year.

The "win" is the amount the casino won on that game. This amount is from the casino's point of view, not the player's. So, if a player exchanged $500 in cash for chips at roulette, and left the table with $325, then the win is $175. If the player exchanged $500 in cash and left with $850, then the win is -$350. A negative win means the casino lost and the player won.

As with the drop, casinos rarely consider win for a single client. The win is totaled over all players of a game (or all players for all tables of a game), for a shift, a day, a week, a month, a quarter or a year.

Finally, the "hold" is defined to be a percentage, and is given by the equation:

$$\text{Hold} = (\text{win} \times 100) / \text{drop}.$$

For example, if a casino dropped \$54,000 at roulette for a week, and won \$8,800 during that same week, then the hold for the week is:

$$\text{Hold} = (\$8{,}800 \times 100) / \$54{,}000 = 16.30\%$$

Each month the Gaming Control Board of the State of Nevada publishes a "Gaming Revenue Report." This document contains information by county (and major gaming area within that county, if appropriate) on the number of tables, the total win and the hold for each major game category at all non-restricted locations. Full statewide results are reported as well.

Figure 5 gives the Nevada statewide drop, win and hold statistics for the fiscal year July 1, 2008 to June 30, 2009. For example, blackjack dropped \$9,839,603,000. The casino win was \$1,141,394,000. The hold was:

$$\text{Hold} = (\$9{,}839{,}603{,}000 \times 100) / \$1{,}141{,}394{,}000 = 11.60\%.$$

Nevada Drop, Win, Hold -- Full Year				
Game	**Tables**	**$ Drop (x 1000)**	**$ Win (x 1000)**	**Hold**
Baccarat	207	$ 7,318,264	$ 746,463	10.20%
Blackjack	2,997	$ 9,839,603	$ 1,141,394	11.60%
Caribbean Stud	18	$ 36,926	$ 9,564	25.90%
Craps	395	$ 3,071,038	$ 402,306	13.10%
Let it Ride	120	$ 248,829	$ 57,181	22.98%
Mini-Baccarat	135	$ 908,568	$ 96,490	10.62%
Pai Gow Poker	298	$ 653,926	$ 114,568	17.52%
Pai Gow Tiles	48	$ 137,118	$ 22,076	16.10%
Roulette	465	$ 1,585,181	$ 310,537	19.59%
Three Card Poker	247	$ 546,182	$ 141,079	25.83%
Other Games	297	$ 534,614	$ 130,339	24.38%

Figure 5.

Drop indicates how much activity is happening at the table. The higher the drop, the more players are handing money to the dealer and asking for chips. Having a high drop is a good thing; it indicates player interest. Of course, having a high win is also good; the casino wants to make as much as possible from each game. The higher the win, the more the casino is earning from actual game play. The tenuous balance between the desire to have both a high drop and a high win is measured by the hold. So, the first lesson about the hold is that there is no theoretical hold percentage for a game that can be computed. More precisely, it is defined to be a ratio of two quantities, neither one of which can be theoretically computed.

If the win gets too high compared to the drop, then people won't want to play, so the drop will get smaller as people move away from the game. As a result, the hold is likely to be a larger percentage. A game with a hold over 30% is usually winning too much too fast and burning out players. Conversely, if the win is too low, then people will flock to the game because it will be viewed as loose. It is likely to have a higher drop, and hence the hold is likely to be a smaller percentage. A game with a hold under 10% is winning too little too slowly and the situation will become intolerable to the casino. A reasonable hold for a casino game is between 10% and 30%, which is what Figure 5 indicates.

Hold is measured by examining the dollar value amounts computed by the normal accounting procedures used in table games. It is also possible to compute the theoretical hold, which is defined by the following equation:

Theoretical hold = (theo x 100) / drop.

Example 1 (Theoretical hold): Suppose a casino game has house edge of 5%. Suppose a player walks up to the game and buys $100 in chips (a drop of $100). Assume the player wagers a consistent $5 per hand. If the game moves at a pace of 40 rounds per hour and the player plays for two hours, then he will play 80 rounds of the game on his original $100 buy-in (assuming he doesn't bust out). Another way of viewing this is that playing 80 rounds means wagering a total of $5 x 80 = $400. It follows that the player churned his original $100 buy-in four times.

The theo (see Math-101) for the player is:

theo = 0.05 x $400 = $20.

It follows that the player theoretically leaves the table with $80 in chips. The theoretical hold is:

Theoretical Hold = ($20 × 100) / $100 = 20%.

The reason this hold is theoretical is that it is not based on the player's actual results. In practice, this player may have lost his entire buy-in, in which case the hold would be ($100 x 100)/$100 = 100%. Or maybe the player won $300, in which case the hold would be (-$300 x 100)/$100 = -300%.

This example underscores the need to measure hold over large amounts of times and many players. If the focus is too narrow, there is nothing to be gained by this metric. Unlike slots, "hold" is about the actual results of the game in practice. It does not measure the house edge and it does not measure the results of one player.

Now, let's rip apart the blackjack hold data in Figure 5. As indicated by Figure 5, the statewide hold was 11.60%. Figure 6 breaks down the statewide hold by county (and major gaming area within that county, if appropriate) for the same fiscal year, July 1, 2008 to June 30, 2009.

The theoretical house edge for blackjack is roughly the same everywhere it is offered, yet blackjack on the strip held 10.94%, in Reno it held 14.86%, in Laughlin it held 15.78%, and blackjack in Wendover held a mighty 20.48%. This challenges the basic premise that the value of the hold says something about the house edge of the game. The game and its edge do have something to do with determining the hold, but many intangible elements also come into play.

The easiest way to understand these numbers is in terms of playing time. In Las Vegas there are shows to see, meals to eat, conventions to attend, clubs to visit, and a host of other distractions that keep people from sitting for very long at the tables. A remote area with a small selection of casinos and a medium level of other activities will usually correspond to somewhat longer playing times. Thus places like Laughlin or Reno are going to have a marginally higher hold. An extremely remote area with almost nothing else to do, like Wendover in Elko County, is going to have the highest hold. Once there, stuck there, nothing else to do. The results for Humboldt County and the balance of Washoe County are most likely statistical anomalies due to the extremely low volume of play. As illustrated in Example 1, low play volume can result in aberrant values for the hold.

Nevada Blackjack Hold by County -- Full Year	
County	**Hold**
Statewide	11.60%
Carson Valley Area	16.27%
Clark County (Las Vegas: Downtown)	11.18%
Clark County (Las Vegas: Strip)	10.94%
Clark County (Las Vegas: North)	16.27%
Clark County (Laughlin)	15.78%
Clark County (Boulder Strip)	14.06%
Clark County (Mesquite)	13.47%
Clark County (Balance of County)	11.51%
Douglas County (South Lake Tahoe)	14.00%
Elko County (Wendover Area)	20.48%
Elko County (Balance of County)	22.24%
Humboldt County	43.05%
Lyon County	12.06%
Nye County	17.10%
Washoe County (Reno area)	14.86%
Washoe County (Sparks area)	15.56%
Washoe County (North Lake Tahoe)	15.60%
Washoe County (Balance of County)	-10.52%
White Pine County	14.58%
Balance of Counties	14.34%

Figure 6.

One of the tragedies of casino operations is that hold is used to measure performance. Drop is important because it measures the total activity the game is generating. Win is important because it indicates the net profitability of the game. The ratio of these two seems as though it should be an important and precise value. But the hold for table games is subject to many more forces than house edge and total amount wagered.

To see how vulnerable the hold is to extraneous forces, we consider a director of table games (DTG) at a Las Vegas strip property who is given a mandate by upper management. Suppose this casino has a 10% hold on its blackjack games and the DTG is given the order to increase the hold to 14%. As a practical matter, there are a few things the DTG can try.

1. Get the players to play longer. To do this, offer free drinks, party pits, friendlier dealers and non-smoking tables. Make the experience at the table one that is tough to walk away from.
2. Get the dealers to deal faster. The hold is increased by the rate at which hands are dealt, otherwise known as the "hands-per-hour" or "decisions-per-hour."
3. Decrease the number of players at each table. If the dealer can deal 360 individual hands per hour, then at a full table of 6 people, each person is going to see 60 hands in an hour. If the number of spots at the table is reduced to 5, then each person is going to see 72 hands per hour. More hands per hour equates to a higher hold (but not necessarily a higher win; there is also the cost of additional dealers).
4. Open more tables and drive players to them using appropriate table minimums. Again, the objective is to deal more hands per player per hour. The fewer number of players at a table, the more hands each player at that table will see in an hour.
5. Get the players to play more money per hand. Given a specific playing duration, the more money that gets on the felt, the higher the overall win. This is often done by increasing the table minimum as the day goes on. This comes with the cost of driving some small bettors away.
6. Make minor modifications in the game rules to increase the house edge. With a higher house edge, the house will win more per dollar wagered at the table. Rules like having the dealer hit on soft 17 or not allowing the player to double down after a split are typical adjustments. These rule changes comes with a cost: even small rule changes may drive some players away. Having blackjack pay 6-to-5 is not "minor."
7. Set up procedures so that chip fills, cleaning spills at the table, and resolving player disputes are handled as quickly as possible.
8. Change the shuffling procedure to speed it up. As part of this, reposition the cut card in the shoe to allow more hands to be dealt from each shoe and hence less shuffling. Have the dealer shuffle as infrequently as possible.

Clearly, some of the techniques for increasing the hold work against each other. For example, by dealing more hands per shoe the game becomes more vulnerable to card counting. Making the rules too house-favorable will drive players away, decreasing the win. Many players like to play at full tables and may walk away if the tables are sitting empty or have one or two players at them. There is a balance point that can only be determined based on understanding the local market.

Example 2. Assume that the average blackjack player at a property buys in for \$200 and plays a strategy that corresponds to a 1% house edge. He likes to play at a full table dealt from a shoe with six spots.

Assume the dealer deals 60 hands per hour to this player. If the player makes \$10 wagers and plays for three hours, then the player sees a total of 180 hands. The house theoretically wins \$10 × 180 × 0.01 = \$18. This corresponds to a pretty dismal theoretical hold of:

Hold = (\$18 × 100) / \$200 = 9.0%.

Now suppose the DTG decides to increase the game speed to 72 hands per hour by reducing the number of spots from six to five. This same player will then see 216 hands in his three hours of play. The house theoretically wins \$10 × 216 × 0.01 = \$21.60. This corresponds to a theoretical hold of:

Hold = (\$21.60 × 100) / \$200 = 10.8%.

Now suppose this player loves the company of the dealer and other players, gets lots of free drinks, and thus settles in for an extra hour, so that he plays four hours instead of three. This same player is now playing 72 × 4 = 288 hands. The house theoretically wins \$10 × 288 × 0.01 = \$28.80. This corresponds to a theoretical hold of:

Hold = (\$28.80 × 100) / \$200 = 14.4%.

Now suppose that after two hours of play, the director raises the table minimum from \$10 to \$15, forcing this player to raise his bet to \$15 if he wants to stay playing with his friends at this great table. The player now plays 2 hours (144 hands) at \$10 and 2 hours (144 hands) at \$15. The house theoretically wins \$10 × 144 × 0.01 = \$14.40 for the first two hours and \$15 × 144 × 0.01 = \$21.60 for the second two hours, for a total theoretical win of \$36 from this player. This corresponds to a theoretical hold of:

Hold = (\$36.00 × 100) / \$200 = 18.0%.

Now suppose that the DTG decides that he can change the rules to make the game more profitable for the casino. He modifies the game so that when the player gets a blackjack it pays 6-to-5 instead of the standard 3-to-2. This adds 1.4% to the house edge. The player is now playing against a house edge of about 2.4%. Obviously, if all things were equal and this player played for the full four hours as before, this would greatly increase the hold. But this player is quickly disgusted with the game: the 2.4% edge seems unfair, so he leaves after an hour and tells the other players to leave as well. During the hour, he was wagering \$10 per hand (he didn't stay long enough to have the minimum raised to \$15) and saw 72 hands. In this case the theoretical win for the house is \$10 × 72 × 0.024 = \$17.28 from the player. This corresponds to a theoretical hold of:

Hold = (\$17.28 × 100) / \$200 = 8.6%.

Another game that sticks out in Figure 5 for its low hold is baccarat (and mini-baccarat). Baccarat is an exception to the usual notions of the importance of hold. The game moves very slowly and players typically buy in for large amounts of money. Players will sit out a few hands just to let the "patterns reemerge." These things drive down hold. Most of the items on the list given to increase hold don't apply. Baccarat players insist on a slow game. The game is already dealt down to the last few cards of the shoe. The game tends to have a serious tone, so friendly dealers don't help keep players at the table. The rules can't be modified. The players are highly superstitious so nothing can be changed. Of course great casino ambiance and service always help the hold, but aside from that, there is little that can be done: the game of baccarat will forever have a very low hold. The only trick for increasing hold is to include a side bet with a high house edge and hope the players will bet some of their money on it. The Dragon side bet (house edge 9.37%) is one of the most popular. On the plus side for the casino, baccarat tends to bring in upscale players who are dedicated to the game and are often repeat and loyal customers for years. On the down side, the Dragon side bet might slow down the main game with the resolution of small wagers. There's always a balance.

A caveat: the hold is not about what one player does or one player's results. The hold is averaged over the results of all players at all tables playing the game over a period of a shift, a day, a week, a month, a quarter, or a year. This statistic is often broken down by table, by pit, by casino, by geographic area, by county and statewide. The hold is at once about all players everywhere and about one player's experience. It is a highly misused and misunderstood number that is affected by so many variables that it ends up meaning pretty much nothing. Nevertheless, the hold is the holy grail of numbers for a table game.

<u>The Danger of Confusing Hold (slots) with Hold (table games)</u>

To see what can go wrong, let's look at an example from blackjack. Assume marketing has not learned that the word "hold" has two meanings. Because they deal with slots all the time, the only meaning of the word "hold" they know corresponds to:

Hold = "house edge"

Marketing learns in the meeting that blackjack is holding 14%. They mistakenly believe that:

House edge for blackjack = 14%.

Because marketing doesn't know the actual numbers, they are not aware that the house edge for blackjack is actually under 1%. This lack of basic education allows them to proceed with their analytics under the remarkably incorrect assumption that the house edge for blackjack is 14%.

At the meeting, marketing is given a directive to run a promotion for blackjack to try and draw in new players. After long deliberation, they arrive at a 2-to-1 blackjack special. The gist of the promotion is that a player who receives a blackjack will be paid at the rate of 2-to-1 instead of the normal 3-to-2.

After consulting with a mathematician, marketing is informed that by changing the payout on blackjack to 2-to-1, the house edge will be reduced by 2.38%. Because marketing believes that hold = "house edge" and that the hold is 14%, they think that the new house edge under this promotion will be 14% - 2.38% = 11.62%. In other words, marketing feels like they can offer this promotion with plenty of room to spare. A hold of 11.62%, they reason, is still a pretty tight game.

The truth is that the house edge for blackjack is about 0.5%. By giving back 2.38%, the house has created a rule that has turned blackjack into the player's favor by about 1.88% playing basic strategy.

Soon after marketing begins advertising the promotion, players from around the country have arrived with fistfuls of cash, armed to place table-maximum wagers.

This may seem like a farfetched example. But the truth is that this exact scenario has played out over and over, in casinos across the states and around the world.

It is therefore worth remembering, every single day, that:

$$\text{Hold (slots)} \neq \text{Hold (table games)}$$

The following article, written by Ken Smith (www.bjinfo.com) first appeared in the "Blackjack Insider Newsletter," July 2003, #42. Soon after publication, it was removed from the website that hosted it and was no longer easily available for public viewing. I can't say for certain, but I will speculate that the author and/or publisher received feedback from advantage players that they shouldn't freely educate casinos. This article was later included in the e-book *"How to Win More Blackjack Tournaments"* by Kenneth R. Smith.

Ken has been more than generous in permitting me to use this account in my public lectures. It is therefore a great honor to be able to present this article in full.

Example 3. (© 2003 Ken Smith, reprinted with permission):

> The Imperial Palace in Biloxi was a busy place on Father's Day June 15th. The casino had widely publicized a promotion available for that Sunday. Blackjacks would be paid at 2-to-1 instead of the normal 3-to-2 odds for 24 hours starting at 12:01 AM Sunday. The promotion applied to all 6-deck tables of blackjack in the casino, and the promotion was good up to the normal table limits.
>
> Paying an extra half-bet for blackjack may sound like nothing special, but it makes a big difference in the game. Blackjack is nearly a breakeven proposition at the outset, so anything that increases the payback can create a lucrative player opportunity. Players get a blackjack once every 21 hands or so on average. By receiving an extra half-bet every 21 hands, the player gets the benefit of an extra 2.3% return. The normal Imperial Palace six-deck game allows resplitting of Aces, and a basic strategy player will average a loss of about 0.4% of his bets. But subtract our 2.3% extra return from that 0.4% house edge, and we end up with a very attractive 1.9% player advantage for a basic strategy player.
>
> That's a pretty tempting deal. Play basic strategy, and you can expect to win 1.9% of your total action. Figuring a leisurely pace of 60 hands per hour, each player stands to make more than one bet an hour in expected profit. I put my family on notice that I would be absent on Father's Day. I planned to spend as much time as I could stand at the tables at the Imperial Palace, and I hoped to win a nice Father's Day profit for my efforts.

I didn't realize then just how widely word of this promotion had spread. I brought plenty of cash to Biloxi, and figured I should show up early in case it was tough to get a seat. Well, I was right about that! I wandered into the casino at 8:30 PM Saturday night, 3 ½ hours before the promotion was to start. It took only a minute to see that I had a problem. There were 14 six-deck blackjack tables on the casino floor, and there wasn't a single seat available at any of them! I hoped that perhaps it was just a busy Saturday night, but it didn't take long to see that wasn't the case. There were plenty of seats available at the handful of two-deck games. Those games weren't eligible for the promotion, so when every six-deck table had seven players crowded around it, I knew I was too late. I kept hoping for someone to leave, and I hung around the casino circling like a vulture.

While waiting, I talked to a lot of bystanders, who were in the same boat as me. I talked to people who had driven down from Memphis, from Jackson, and others who had flown in from other parts of the country. There were ninety-eight seats available for the promotion, and I would guess there were at least that many extra players hoping for a seat. As midnight approached, the situation didn't improve. I tried buying a seat, offering $300. I found no takers. I learned that the last seats had filled up around 5:30 PM, more than 6 hours before the scheduled start. And these players weren't budging.

I decided I'd hang around for the show at midnight, when I expected to see a sea of cash appear on the tables. It was quite a spectacle. Starting about 10 minutes before midnight, the buy-ins began, with thousands of dollars spread and black $100 chips and purple $500 chips stacked and counted. As one of the dealers came by me headed out of the pit for her break, she just shook her head, saying "Well, the craziness has started." She was right.

At midnight, the average bet size on the tables made a dramatic jump. Almost all of the players had been betting the table minimum while they waited. Now, with the 2-to-1 payoffs in effect, the situation changed immediately. Most of the table maximums were $2000, and of the 98 spots in play, I estimate that 50 of them were filled with table-max $2000 bets. At table after table, there were stacks of 4 purple chips filling the betting circles. I stopped to watch a table near the end of one pit. Of the seven spots at the table, the first six players bet $2000 on the first hand after midnight. The last player made a $200 bet instead. After the cards were dealt, one player had split and another doubled. When the dealer busted, there were nowhere near enough chips in the tray to pay the more than $16,000 won on hand #1 of the promotion.

This same scene played out at many of the tables, and play ground to a halt while the tables waited for fills. Players at a few tables wised up and stopped betting multiple purple chips if it was obvious that the dealer couldn't pay them if they won. They began betting black $100 chips instead, until the tray was emptied of black. Then they switched to $25 bets and waited for more chips to arrive. By 12:45, there were only 3 tables that still had purple chips in the tray. Those were the tables where the dealer had been hot, and the players were losing. Finally, just before 1 AM, the security crews had managed to refill most of the chip trays. The $2000 bets again took the place of the $25 bets. At the table where hand one wasn't payable, they finally got more chips, but the first fill was only about $30,000. Three hands later they had emptied the tray again. This time, the fill came much quicker and the action resumed.

It was amazing to just walk around the floor and look at the amount of money flowing across the tables. There were several players who quickly amassed stacks of chips worth $50 or $60

> thousand. This couldn't last. At 3 AM, the pit personnel hurriedly hand-wrote new betting limit signs at each table. The new table maximum in the casino was $500. Around 75 of the spots in play were now flat bets at the new max of $500. I went to get a few hours of sleep, since it was quite apparent that I wouldn't be getting a seat anytime soon.
>
> I came back down to the casino around 7 AM, hoping that a few of the players would be too tired to continue. No such luck. Most of them didn't even look tired after playing all night! I stood around chatting with other wanna-be players. At 8 AM, the table limits were dropped again, this time to $200. Still nobody moved. Then at 10 AM, the scene changed again. The casino personnel began closing tables. By 10:30 AM, they had eliminated all the tables but one. The promotion still ran, with a maximum bet of $200, but only one six-deck table remained open in the entire casino. When asked about their plans for the rest of the day, they indicated that the plan was to reopen the other tables as double-deck games.
>
> At that point, I had seen enough, so I headed back home to enjoy the rest of Father's Day. Later on that day, the table max at the lone surviving table was lowered to $50, where it stayed for the rest of the day. Word from the casino was that they had lost over a million dollars overnight. I don't think they had any idea what they were getting into. It was an amazing night. I just wish I'd gotten there earlier!

Similar accounts abound, though none have been written with as much journalistic precision. The most recent one (July, 2010) was a "Triple Down" promotion for blackjack at the Mohegan sun. This promotion allowed players to triple down instead of doubling down in blackjack, leading to a hefty player advantage. The reporting I have is that this promotion occurred because of a misunderstanding of the meaning of the word "hold".

So remember,

$$\text{Hold (slots)} \neq \text{Hold (table games)}.$$

Probability

The backbone of casino game mathematics is "probability." Informally, we understand probability as a number that describes the chance that something will occur. It is usually given as a fraction or decimal with a value between 0 and 1, or as a percent with value between 0% and 100%. A probability of 0 means the event can never occur. A probability of 1 means the event always occurs. For example, toss two dice and have the sum come up 13; that's impossible, so the probability is 0. Toss a coin and have it come up either heads or tails; that's a certainty, so the probability is 1. Dice and coins never land on edge in our mathematically perfect world.

The formal theory of probability begins by understanding what's known as the "sample space." This is simply a description of all possible outcomes – everything that can possibly happen. Some examples:

- There are 2 outcomes when a coin is tossed; the sample space is {Heads, Tails}.
- There are 6 outcomes when a single dice is rolled; the sample space is {1, 2, 3, 4, 5, 6}.

- There are 36 outcomes when two dice are rolled (the first dice and the second dice each produce a value from 1 to 6, so there are $6 \times 6 = 36$ outcomes). The sample space is {[1,1], [1,2], [1,3], [1,4], [1,5], [1,6], [2,1], … and so on}, or more simply, using set builder notation, it is $\{[x,y] \mid 1 \le x \le 6 , 1 \le y \le 6 \}$.

- There are 38 outcomes when a roulette wheel is spun. The sample space consists of the numbers 1 through 36, together with zero and double-zero.

- There are 52 ways a single card can be dealt from a deck. The sample space consists of the set of card values (rank and suit).

- There are 1,326 two-card starting hands in Texas Hold' em. The sample space consists of the set of pairs of cards.

Many experiments have a sample space that is easily understood from the nature of the experiment but may not be explicitly described. Sample spaces are often very large for casino games, reflecting the intuitive feeling that there are too many things that can happen to count them all. Mathematicians have to count them all.

- There are 66,300 ways a player can be dealt two cards against a dealer up-card in single deck blackjack.

- There are 2,598,960 ways that 5 cards can be dealt to a player in poker.

- There are 407,170,400 ways that the player and dealer can be dealt their respective three cards hands in Three Card Poker.

- There are 55,627,620,048,000 ways two players can play out a hand of Texas Hold' em, including their initial two cards, the three flop cards, the turn and the river, playing head's up against each other.

An "event" consists of some of the things that can happen in the experiment. An event is a way of describing some of the things out of everything that might happen in a game. Here are some examples of events that correspond to some of the games in the list above:

- Tossing a coin and getting heads.
- Rolling two dice and getting a sum of 7.
- Being dealt a blackjack against a dealer up-card of an Ace.
- Being dealt a pat full house as a five card poker hand.
- Being dealt a straight that loses to a higher straight in Three Card Poker.
- Being dealt a pocket pair in Texas Hold' em.

To compute the probability of an event we need to know two pieces of information. First, we need a full count of the number of individual elements in the sample space. Second, we need to know how many individual elements are in the collection that corresponds to the event. Simply put, we need to know the size of the sample space and the size of the event.

Knowing these values, we define the probability of the event by the equation:

P(event) = (Size of event) / (Size of sample space).

The word "Probability" is cumbersome to write out; it is customary to use the letter "P" when referring to the probability of an event.

What is not evident by the equation for probability is how to count the size of various collections. Unfortunately, in gaming there are very few simple problems and these counting problems can be exceedingly complex. For those cases when counting is easy, probabilities can be quickly computed. We'll go through several examples to demonstrate some of the techniques. Hopefully these examples will help clarify the concept of probability for casino games and some of the methods that are used to come up with these values.

Example 4. When tossing a coin, let H = "Heads" and T = "Tailes." For a single coin toss, the sample space is {H, T} and it has two elements in it. To get a "heads" corresponds to the event {H}, which has a single element in it. So:

P(heads) = 1 / 2 = 0.5000.

If a coin is tossed three times, then the Sample space is {HHH, HHT, HTH, THH, HTT, THT, TTH, TTT}. This sample space has 8 elements in it. The event of getting two heads is {HHT, HTH, THH}. The probability of getting two heads out of three tosses is:

P(two heads in three coin tosses) = 3 / 8 = 0.375.

Example 5. In Texas Hold' em, there are 52 possible first cards you can be dealt and 51 possible second cards you can be dealt. Because the order of the cards doesn't matter, we divide by 2 to take account of symmetry. The size of the sample space is:

$$(52 \times 51) / 2 = 1326.$$

We consider the event of being dealt a pair. If we look at pairs of deuces, there are six possible pairs (again, order doesn't matter): {[2C,2D], [2C,2H], [2C,2S], [2D,2H], [2D,2S], [2H,2S]}. For any rank of card, there are six possible pairs of that rank. There are 13 possible ranks for the pair, and six ways of making the pair at that rank. This gives $13 \times 6 = 78$ pairs. So, the size of the event of being dealt a pocket pair is $13 \times 6 = 78$. So:

P(pocket pair) = 78 / 1326 = 0.0588.

Example 6. When rolling two dice, the sample space is $\{[x,y] \mid 1 \le x \le 6 , 1 \le y \le 6 \}$; the sample space has 36 elements in it. The event "getting a sum of 7" corresponds to the subset of the sample space {[1,6], [2,5], [3,4], [4,3], [5,2], [6,1]}. This subset has 6 elements in it. Therefore, P(sum of 7) = 6 / 36 = 0.1667.

Figure 7 gives the probabilities for the various sums of two dice.

Craps	
Total (X)	**Prob (P)**
2	0.0278
3	0.0556
4	0.0833
5	0.1111
6	0.1389
7	0.1667
8	0.1389
9	0.1111
10	0.0833
11	0.0556
12	0.0278

Figure 7.

Example 7. In single deck blackjack, there are (52 × 51) / 2 = 1,326 possible two card starting hands for the player. For each such hand, there are 50 cards left in the deck that can be dealt as the up-card for the dealer. Thus there are 1,326 × 50 = 66,300 possible starting situations. So the sample space has size 66,300.

The event we're interested in is being dealt a blackjack against a dealer Ace. A player blackjack consists of a face card (16 of them) and an Ace (4 of them). There are 16 × 4 = 64 ways for the player to be dealt a blackjack. There are 3 remaining aces out of 50 cards for the dealer up-card. So there are 16 × 4 × 3 = 192 ways for the player to have a blackjack against a dealer up-card of an Ace. The event of a player blackjack against a dealer Ace has size 192. Therefore:

P(Blackjack vs. Dealer Ace) = 192 / 66,300 = 0.002896.

This means that the situation where a player can consider taking "even money" comes up about 29 times in every 10,000 hands, or about once every 345 hands. At a full table and at standard dealing speeds, this amounts to about one player "even money" decision per table per hour in single deck.

Doing the same computations for a six-deck shoe game, we get a sample space with size 15,039,960 and an event with size 52,992, so that:

P(Blackjack vs. Dealer Ace) = 52,992 / 15,039,960 = 0.003523.

This means that even money situations come up about 35 times per 10,000 hands, or about once every 284 hands. Even money is far more common, and therefore potentially far more profitable for the casino, in a shoe game.

Example 8. In this example we are going to go through the mathematical details of the Pair Plus bet in Three Card Poker and compute all the probabilities.

First of all, the sample space consists of all the distinct three card hands that can be dealt from a deck of 52 playing cards. There are 52 possible values for the first card, 51 possible values for the second card, and 50 possible values for the third card. Three cards can be rearranged in six possible ways, but still be the same three cards: ABC, ACB, BAC, BCA, CAB and CBA. Therefore, the number of distinct three card hands is:

$$(52 \times 51 \times 50) / 6 = 22{,}100.$$

Event: Straight Flush. These are easy to count. In order to have a straight, we have to get one of the hands A23, 234, 345, 456, 567, 678, 789, 89T, 9TJ, TJQ, JQK, QKA in one of the four suits. There are 12 straights and each straight can happen in one of the four suits, so there are 12 × 4 = 48 ways to get a straight flush. So,

$$P(\text{Straight Flush}) = 48 / 22{,}100 = 0.002172.$$

Event: Three of a Kind. Again, we can count these directly. For deuces, we can get trips with any of the hands [2C,2D,2H], [2C,2D,2S], [2C,2H,2S] and [2D,2H,2S]. There are four trips for any rank and there are thirteen ranks, so there are 13 × 4 = 52 ways to get three of a kind. So,

$$P(\text{Three of a Kind}) = 52 / 22{,}100 = 0.002353.$$

Event: Straight (not straight flush). We begin by counting all straights, including the straight flushes. As above, a straight consists of one of the hand types A23, 234, 345, 456, 567, 678, 789, 89T, 9TJ, TJQ, JQK, QKA. Let's look at A23 in particular. We can have any suits for the A, 2 and 3; there are four choices for the suit of each (clubs, diamonds, hearts and spades). Therefore the A23 straight consists of picking a card from each of {AC, AD, AH, AS}, {2C, 2D, 2H, 2S} and {3C, 3D, 3H, 3S}. There are 4 × 4 × 4 = 64 ways of choosing these three cards. So there are 64 straights of type A23. There are 12 types of straights and each straight can happen in 64 ways. So the total number of straights is 12 × 64 = 768. But this number includes the 44 straight flush hands that we already counted, so we have to subtract those. This leaves 768 – 48 = 720 straights that are not a straight flush. So,

$$P(\text{Straight}) = 720 / 22{,}100 = 0.032579.$$

Event: Flush (not straight flush). We need to count all the hands XYZ where X, Y and Z have the same suit and the hand is not a straight flush. We'll pick on the suit Clubs. There are 13 cards. We can choose any of these 13 for the first card, any of the remaining 12 for the second card, and any of the remaining 11 for the third card. Because three cards can be rearranged in six possible ways and the order the cards were dealt doesn't matter, it follows that the number of ways of getting a flush in clubs is (13 × 12 × 11) / 6 = 286. Of these flushes, 12 of them are a straight flush. This leaves 286 – 12 = 274 flushes in clubs that aren't straight flushes. Because there are four suits, it follows that the total number of flushes is 274 × 4 = 1,096. So, P(Flush) = 1,096 / 22,100 = 0.049593.

Event: Pair. There are 6 ways of getting a pair of any rank. For example, to get a pair of deuces the player must hold one of the pairs [2C,2D], [2C,2H], [2C,2S], [2D,2H], [2D,2S] or [2H,2S]. Once the pair is dealt, the third card has to be a card of a different rank (otherwise the hand would be trips). There are 48 cards that are not the same rank as the pair. Putting this together, there are 13 possible ranks for the pair, 6 pairs of that rank, and 48 possible third cards of a different rank. Multiplying these together, we get a total of 13 × 6 × 48 = 3,744 pairs. So,

P(Pair) = 3,744 / 22,100 = 0.169412.

Event: Nothing. This final case is the easiest. To get nothing simply means that the hand isn't any of the above. We simply subtract all of the results above from the total number of hands. This gives 22,100 – 48 – 52 – 720 – 1,096 – 3,744 = 16,440 hands that are losing hands for the Pair Plus wager. So,

P(Nothing) = 16,440 / 22,100 = 0.743891.

Hit Frequency. We now illustrate how to compute the hit frequency. For the Pair Plus wager, a player wins any time he has a pair or better. The hit frequency corresponds to the probability of being dealt a pair or better. To find the number of elements in this event, we simply add up all the possible winning hands. Doing this, we get 48 + 52 + 720 + 1,096 + 3,744 = 5,660. So,

Hit Frequency = P(Winning) = 5,660 / 22,100 = 0.256109.

It follows that the hit frequency is 1-in-3.9. Rounding off, the player wins about once every four times he makes a Pair Plus wager.

Finally, note that if we add all the probabilities together, we get 0.002172 + 0.002353 + 0.032579 + 0.049593 + 0.169412 + 0.743891 = 1.000000. This simply means that the probability that something is going to happen is 1, a certainty.

Figure 8 summarizes the probabilities we just computed and the payouts for the various events using the most generous pay table Shuffle Master offers for this wager. It is customary to indicate a payout of -1 when the player loses; after all, he lost his wager. In all other cases, the player's wager is returned and he is paid the amount indicated in the table.

Three Card Poker: Pair Plus		
Event	**Prob**	**Payout**
Straight Flush	0.0022	40
Three of a Kind	0.0024	30
Straight	0.0326	6
Flush	0.0496	4
Pair	0.1694	1
Nothing	0.7439	-1

Figure 8.

Variance and Standard Deviation

A casino game has a house edge: if a person wagers \$100, then the house edge predicts the return on that \$100 the player will receive, on average. For example, playing roulette, the player will receive about \$94.74. Playing blackjack, the player gets back about \$99.54. Playing Three Card Poker (just the Ante/Play), the player gets back about \$96.63. These values are the "average return" the player gets when he makes an initial wager of \$100, and they correspond to the house edge for each game.

Of course, a player never gets exactly \$94.74, or \$99.54, or \$96.63 back at the end of a round. He may lose his wager of \$100, he might push, or he might win \$100, \$200, or something else can happen. For example, if the player placed \$100 on 34-red in roulette and won, he would win \$3,500. If a blackjack player split to four hands, doubled on each, and then watched as the dealer pulled a six-card 21 beating him on all hands, then the player would lose \$800. Each time a round completes, the amount the player actually wins or loses will be a value that is usually different than the average return he expects. Over the long run, these different returns will average out. The *average return* is the ratio of the sum of all the returns to the sum of all the initial bets.

We need a measurement of how wildly the various returns on individual rounds of a casino game are from the average return. That's where the standard deviation and variance come in.

The *standard deviation* of some data can be intuitively thought of as "the average distance of a data value from the average of the data." It is customary to use the Greek letter σ to denote the standard deviation. A data value means one of the pieces of data collected. Some of the data values will be less than the average of the data; some of the data values will be greater than the average of the data. Some may equal the average. Each data value is a certain distance from the average of the data. The point here is to find the "average" of the distance from the average.

There is a technical aspect to the standard deviation that comes from what it means to find the "average." We illustrate this technicality by first considering two methods of finding the average of values X and Y.

One notion of an average of two values X and Y is to compute the value (X + Y)/2. For example, with X = 3 and Y = 5, the average is (3 + 5)/2 = 8/2 = 4. This method of computing the average is known as computing the "arithmetic mean."

Another way of taking the average of two values X and Y is to take the square root of the average of the sum of the squares. That is, the average is sqrt((X^2 + Y^2)/2). For example, with X = 3 and Y = 5, the second way gives sqrt((3^2 + 5^2)/2) = sqrt(34/2) = 4.12. Note this is slightly larger than 4, the arithmetic mean. This method of computing the average is known as computing the "root mean square."

With more values, a similar method is used to compute these averages. For example, with the values 3, 7, 8, 2, the arithmetic mean is (3 + 7 + 8 + 2)/4 = 5. Similarly, the root mean square is sqrt((3^2 + 7^2 + 8^2 + 2^2)/4) = sqrt((9 + 49 + 64 + 4)/4) = sqrt(126)/4 = 5.61. Again, note that the root mean square is slightly larger than the arithmetic mean.

For the standard deviation, the average is computed using the second of these methodologies. The reason is that some of the deeper theorems rely on subtle properties the root mean square has that the arithmetic mean does not have.

To compute the standard deviation of some data:

1. Find the arithmetic mean of the data.
2. Find the distance of each value from the arithmetic mean.
3. Square each of these distances.
4. Compute the arithmetic mean of the squares.
5. Take the square root of the value in step 4.

For example, the standard deviation of the values 3, 7, 8 and 2 is found as follows:

1. The arithmetic mean is (3 + 7 + 8 + 2)/4 = 5.
2. The distance of each number from 5 is: -2, 2, 3, and -3.
3. The square of these distances gives the values 4, 4, 9 and 9.
4. The arithmetic mean of the squares gives (4 + 4 + 9 + 9)/4 = 26/4.
5. The square root of the value obtained from step 4 is sqrt(26/4) = 2.55.

So, the standard deviation of the data 3, 7, 8 and 2 is 2.55.

When this definition is applied to a casino game, the standard deviation is "the average distance between the result of a round of a casino game and the average return (theoretical RTP) for that game." In other words, we look at the various returns a casino game has, the different outcomes that are possible for the player, and ask: "on average how far away are those returns from the expected theoretical return?" The answer is the "standard deviation," As previously mentioned, we denote this value σ.

Along with the "standard deviation," the concept of "variance" and "volatility" are usually described. *Variance* is defined to be the square of the standard deviation, that is, variance = σ^2. It is the value obtained at step #4 in the computation of the standard deviation. The purpose of introducing the variance is that many mathematical computations are easier, for very technical reasons, when applied to the variance.

We will use the word "variance" casually, with the understanding that it is really measuring the same thing as the standard deviation: just take the square root of the variance to get the standard deviation. Among the many reasons that authors, players, casino management, and game developers prefer the word "variance" is that it is shorter, easier to write, less technical sounding, and has fewer syllables than "standard deviation." It is more intuitive to say a game has "high variance" than to say it has "high standard deviation."

Finally, *volatility* is often used as a synonym for either the standard deviation or variance, depending on the context. Volatility is also used informally to describe the degree of unpredictability in a measurement or in a financial exchange, like gambling. Phrases like "high volatility" are usually applied to financial instruments with big up and down swings. Similarly, "low volatility" suggests gradual movement.

Intuitively, the larger the value of the standard deviation (variance, volatility), the more the data is spread out, or deviates from the average. A standard deviation (variance, volatility) of zero means that every piece of data has the same value as the average for all the data. We are going to explore these issues by comparing a day-job to an experience playing a casino game.

Suppose you work at a job and your paycheck is $1,500 per week. You quickly come to rely on the steadiness of that income. You make choices about purchases, vacations, health care and other matters based on knowing an exact amount of cash that will be paid to you on a certain date. Any time the data points are all the same, the standard deviation is zero. Because the only data point is $1,500, your average salary each Friday is $1,500. The standard deviation is σ = $0.00.

We're going to experiment with a few more ways of getting paid. All of these methods have the same expected salary: getting paid $1,500 per week over the long run.

1. Suppose each Friday you roll a single six-sided dice to determine your weekly paycheck. Your pay is $0 if a "1" is rolled; $500 if a "2" is rolled; $1,000 if a "3" is rolled; $2,000 if a "4" is rolled; $2,500 if a "5" is rolled; and $3,000 if a "6" is rolled. Then your average salary is: ($0 + $500 + $1,000 + $2,000 + $2,500 +$3,000) / 6 = $1,500. Using steps 1 - 5 given above, the standard deviation is σ = $1,080.23.

2. Suppose each Friday you flipped a coin. Heads, you get $3,000 and tails you get nothing. Clearly your average salary is $1,500, so over the long run you are earning the same amount. The standard deviation is σ = $1,500.00.

3. Suppose each Friday you roll a single six-sided dice to determine your weekly paycheck. Your pay is $0 if the dice shows a 1, 2, 3, 4, or 5. You get paid $9,000 if the dice shows a 6. Your average salary is ($0 + $0 + $0 + $0 + $0 + $9 000) / 6 = $1,500. The standard deviation is σ = $3,354.10.

4. Finally, suppose each Friday you roll two dice. If the sum is 12 you get paid $54,000. For any other sum of the two dice, you get paid $0. You can do the math or not, but on average your pay is $1,500 per week. The standard deviation is σ = $8,874.12.

I would be very uncomfortable having a day job and not knowing come Friday if I will get paid or not. I want to know that fifteen hundred dollars will be in my hand. If I wanted to gamble with the least volatility, then using method 1 would be the choice. If I wanted to gamble on my salary with the greatest volatility, then method 4 would be the choice.

Most of us want no variance at all when it comes to our day job. Turning this around, when people play a casino game, they are willing to pay the casino a certain fraction of their wagers in order to get have an experience of variance that suits their comfort level. The product the customer is purchasing is the variance and the adrenaline rush that comes with it. No one would play a zero variance slot where every single time they pulled the handle, they lost 5 cents on their $1 wager. Nor would players enjoy wagering $1 in the machine that says "Change" with its 0% house edge and zero variance. Both of these zero variance games sound pretty dumb. However, players have been knows to play with zero variance on roulette – they play the same wager on all 38 numbers; this scheme loses just over 5 cents for every dollar wagered. Players do this, for example, as a way to earn free drinks or other comps. Likewise, simultaneously playing the Banker and the Player at baccarat, or both the Pass and Don't Pass wagers in craps are situations with very low volatility: most rounds a specific amount corresponding to the house edge is lost.

Let's look at the standard deviation for the two lotteries.

> Lottery 1. For 9 data points, we have a value \$100,000 as the data value, and for 999,991 data points we have a value \$0 as the data value. The standard deviation is σ = \$300.00 (the variance is σ^2 = \$89,999.19).
>
> Lottery 2. For 900,000 data points, we have a value \$1 as the data value, and for 100,000 data points we have a value \$0 as the data value. The standard deviation is σ = \$0.30 (the variance is σ^2 = \$0.09).

For most players, the variance for lottery 1 is too high and the variance for lottery 2 is too low. Players like a range of variance in their casino games. Sometimes high variance will sell, but low variance almost never sells. Just as customers don't want variance when it's part of their income, customers do want variance when it's part of gaming entertainment. For most of us, our comfort level for our weekly paycheck is a variance of \$0. But when a player sits down at a casino game, it's all about variance. Variance creates excitement. Variance drives the adrenaline rush. Variance is the product.

Common Casino Games: Mean, Variance and Standard Deviation

Figure 9 (following page) gives a list of many well known table casino games, their mean result for a single round wagering 1 unit, their standard deviation and variance.

The wagers in Figure 9 with the highest volatility are, in this order, Mississippi Stud, the Dragon 7 side bet in EZ Baccarat, playing a number straight-up in roulette, Let it Ride, and a 2 or 12 proposition bet in craps. There is a lot of similarity to playing a number straight-up in roulette (odds are 1-in-38 and pays 35-to-1) and playing a proposition bet on 2 or 12 in craps (odds are 1-in-36 and pays 30-to-1). Hence these two wagers have very similar values for their standard deviation.

High variance games and wagers are for those who like a lot of volatility. The results are going to be all over the place in a session. Wild things are likely to happen. There will be big jackpot hands and long losing streaks. There won't be much, if any, back and forth, winning and losing one bet at a time.

The high value for the variance in Let it Ride has been problematic for the game. It is almost as wild a ride as playing a number straight-up in roulette. Moreover, a player cannot choose any method of playing the game that reduces his variance. Let it Ride has dropped from a high of 221 tables in Nevada in 1996 to 121 tables in 2009 (see Appendix C). Let it Ride continues to lose market share because of its unshakeable volatility. Players lose too quickly while hoping for that jackpot hand.

The game Mississippi Stud has a pay profile similar to Let it Ride; I've heard it referred to as "Let it Ride in reverse." However, Mississippi Stud also has nearly twice the standard deviation as Let it Ride, making it by far the most volatile non-traditional table game in wide distribution. Mississippi Stud is demographic specific. It works best in casinos with a lot of different carnival games; players can drift to it when they want a certain experience, but few sit at it for a long time. It quickly drains players who must play 4 wagering units per round to chase a possible win. Any time a player is dealt a high pair, he instantly wins 10 units, possibly a lot more. The single highest volatility hands for this game are four-of-a-kind and a royal flush. The pay table is highly skewed towards the top end; players lose a lot while pursuing high paying rare hands. It is impossible for a player to win less than four bets at a time; losing streaks can be long and tortuous.

Casino Games: Mean (μ), Std. Dev. (σ), Variance (σ^2)			
Casino Game	**μ**	**σ**	**σ^2**
Baccarat (Banker)	0.9896	0.93	0.86
Baccarat (Player)	0.9876	0.95	0.90
Blackjack (Las Vegas rules)	0.9972	1.15	1.32
Caribbean Stud	0.9478	2.24	5.02
Casino War	0.9712	1.05	1.10
Craps (Pass/Come)	0.9859	1.00	1.00
Craps (Don't Pass/Don't Come)	0.9864	0.99	0.98
Craps (Any Craps)	0.8889	2.51	6.30
Craps (Proposition 2, 12)	0.8611	5.09	25.91
EZ Baccarat (Banker)	0.9898	0.94	0.88
EZ Baccarat (Dragon 7 Bet)	0.9239	6.09	37.10
Let it Ride	0.9649	5.17	26.73
Mississippi Stud	0.9509	9.84	96.76
Pai Gow Tiles	0.9850	0.75	0.56
Pai Gow Poker	0.9854	0.75	0.56
Roulette (Even/Odd)	0.9474	1.00	1.00
Roulette (Column of 12)	0.9474	1.39	1.94
Roulette (Number straight-up)	0.9474	5.76	33.21
Three Card Poker (Ante/Play)	0.9663	1.64	2.69
Three Card Poker (Pair Plus)	0.9768	2.91	8.47

Figure 9.

The games with the lowest volatility are, in this order, Pai Gow Tiles, Pai Gow Poker, baccarat (Banker), EZ Baccarat (Banker), baccarat (Player), roulette (even/odd) and craps (don't pass/don't come). These games are for those who shun variance and want the lowest volatility. Players' results will move slowly up and down. These games will be safe: players don't have to worry that big events are going to suddenly change their session results.

At the extreme of volatility for blackjack side bets is the Lucky Ladies side bet (not included in the table). For a six deck game with pay table 4-9-19-125-1000, the Lucky Ladies side bet has $\mu = 0.7529$ and $\sigma = 4.96$. It has a very high house edge of $100 \times (1 - 0.7529)\% = 24.71\%$ (the Nevada state legal maximum house edge is 25%). I am not aware of any table game or side bet (or slot machine, or keno pay table or video poker game) that has a higher house edge.

A side bet that is even more volatile than Lucky Ladies is the Dragon 7 side bet in EZ Baccarat. A player wins this bet by making a wager that a Banker three-card total of seven will beat the player's hand. The odds are about 1-in 44.38 and this side bet pays 40-to-1. The Dragon side bet has $\mu = 0.9239$ (house edge of 7.61%) and $\sigma = 6.09$.

Most other side bets, no matter the game, are much less volatile than Lucky Ladies in blackjack or the Dragon side bet in EZ Pai Gow. In general, only progressive and streak side bets have a higher level of volatility.

Take note of the standard deviation for the roulette wagers listed in Figure 9. These illustrate the wide range of variance the player can choose when playing roulette. The player can wager on even/odd, a column, a number straight-up, or he can mix and match as he decides. The house edge never changes.

The four wagers listed for craps in Figure 9 illustrate that a player can pick and choose his variance. However, in craps, in order to get high variance without taking "odds," the player will have to play high house edge wagers. This is why odds are so popular in craps; odds are pure variance, nothing more. The house edge for all odds wagers in craps is 0.00%, the theoretical RTP is 100%. Over the long run, the player expects to break even on his odds wagers. Most casinos offer odds; some use it as a marketing tool, allowing up to 100 × times odds to be purchased. The primary purpose of offering odds is to allow the player to have the experience of volatility he wants at no additional cost to the player or casino.

<u>The Long Run</u>

The opening scene in the play "Rosencrantz and Guildenstern Are Dead" shows the main characters involved in an experiment in probability theory. Rosencrantz and Guildenstern are on a mountain path when Guildenstern sees a coin on the trail. He picks up the coin, and flips it, then flips it again and again... What follows is unlikely; heads comes up 157 times in a row. Here is the concluding dialogue from that scene:

> GUILDENSTERN: It must be the law of diminishing returns. I feel the spell about to be broken. (He flips a coin high into the air, catches it, and looks at it. He shakes his head.) Well, an even chance.
>
> ROSENCRANTZ: Seventy-Eight in a row. A new record, I imagine.
>
> GUILDENSTERN: Is that what you imagine? A new record?
>
> ROSENCRANTZ: Well.
>
> GUILDENSTERN: No questions, not a flicker of doubt?

ROSENCRANTZ: I could be wrong.

GUILDENSTERN: No fear?

ROSENCRANTZ: Fear?

GUILDENSTERN: (He hurls a coin at Rosencrantz.) Fear!

ROSENCRANTZ: (looking at the coin) Seventy-Nine.

GUILDENSTERN: I think I have it. Time has stopped dead. The single experience of one coin being spun once is being repeated …

ROSENCRANTZ: Hundred and fifty-six.

GUILDENSTERN: … a hundred and fifty-six times! On the whole, doubtful. Or, a spectacular indication of the principle that each individual coin spun individually is as likely to come down heads as tails, and therefore should cause no surprise each individual time it does.

ROSENCRANTZ: Heads. I've never seen anything like it.

GUILDENSTERN: We have been spinning coins together since … I don't know when. And in all that time, if it is all that time, one hundred and fifty-seven coins, spun consecutively, have come down heads one hundred and fifty-seven consecutive times, and all you can do is play with your food!

Getting heads 157 times in a row is an event with a probability so small it can just barely be determined by an Excel spreadsheet computation:

P(157 consecutive heads) ≈ 0.00000000000000000000000000000000000000000000000547382.

Let's put the experience of these characters into the context of the "long run." Suppose we are conducting an experiment tossing a coin over and over. Let's count the number of times heads and tails each come up. We will need three counting variables for this. We let "H" denote the number of times that heads comes up, "T" will denote the number of times tails comes up, and "N" will denote the total number of times the coin has been tossed.

Many people mistakenly believe that "in the long run, the number of heads and tails balances out, so that they are roughly the same." This is false; that is not what the "long run" says at all. The most common fallacious thinking about the long run is the belief that eventually we must have the number of heads equal to the number of tails, or H = T. Based on this belief, people wager more on heads when a lot of tails have appeared, or more on tails when a lot of heads have appeared, using the logic that the other variable has to catch up. That is not the law of averages. Unequivocally, this is not what "the long run" means. It is not a statement that things come out exactly right if enough rounds of the experiment are conducted.

The "long run" simply means that the ratio of reality to theory gets closer and closer to 1 as the number of times we conduct the experiment gets large. To be precise in the case of the coin experiment, the probability of heads is 0.5. The number of times we expect heads to come up in N tosses is (0.5) × N. That's theory. Reality is the variable H that counts the actual number of times heads came up. So, the "long run" says that the fraction:

$$H / [(0.5) \times N]$$

gets closer and closer to 1 as N (the number of tosses) gets large.

Suppose we are back in Tom Stoppard's play, and heads has come up 157 times in a row. Then H = 157 and N = 157, so the fraction in question is:

$$157 / [(0.5) \times 157] = 2.0000$$

That's not very close to 1. Now what happens if the coin is tossed 2,000 more times and it reverts back to a normal 50-50 for those 2000 times, despite the beginning drama? Then after N = 2,157 tosses we have H = 1,000 + 157 = 1,157 and T = 1,000. The values of H and T have not gotten any closer together. So what about the long run? Now the fraction is:

$$1{,}157 / [(0.5) \times 2{,}157] = 1.0728.$$

We see that the initial issue with the coin is dwarfed by the steadiness of the events to follow. So, long term normal behavior overwhelms short-term aberrations. That's the long run!

Let's reset and run another example from scratch. In this case I am going to show you that the values of H and T can get further and further apart, but still have "the long run" work out right. Consider the following possible outcomes for tossing the coin:

- After N = 20 we have H = 11 and T = 9. The fraction is 1.10000.
- After N = 200 we have H = 102 and T = 98. The fraction is 1.02000.
- After N = 2000 we have H = 1003 and T = 997. The fraction is 1.00300.
- After N = 20,000 we have H = 10004 and T = 9996. The fraction is 1.00040.
- After N = 200,000 we have H = 100005 and T = 99995. The fraction is 1.00005.

And so on. As you can see, the values of H and T are getting further apart. After N = 20, the values of H and T differ by 2. After N = 200,000 these values differ by 10. Continuing this pattern we see that H and T are never equal to each other and that they continue to get further apart. Nevertheless, the fraction describing the long run gets closer and closer to 1. That is, the fraction H / N is getting closer and closer to 1 / 2, the theoretical probability of getting heads.

The long run is the statement that:

$$(\#\text{Counted in reality}) / (\#\text{Expected by theory}) \to 1$$

This ratio gets closer to 1 as the number of times the experiment is conducted gets larger and larger. Having this fraction get close to 1 is not the same as having the numerator and denominator get close to each other in their values. That's the tough mathematical point that it's hard for many people to get. Because this is a ratio, it does not say anything about the exact values of the variables. Fractions can get closer and closer to a value, even as the numerator and denominator get further apart.

For example, the following series of fractions gets closer and closer to 1, even as the numerator gets further and further from the denominator:

11/10, 102/100, 1,003/1,000, 10,004/10,000, 100,005/100,000 …

The long run is not a mystery saying that things have to even out, that hot dealers must go cold, that red and black have to occur the same number of times in roulette, that slots are more likely to pay out if they haven't hit in a while, that the player is due for a blackjack. There is no "evening out" implied by the long run. There is simply the slow march of fractions converging to the expected value of 1.

The purpose of the variance is to give a sense for how quickly the fractions get close to 1. In a low variance game, the ratio of reality to theory very quickly converges to 1. The higher the variance, the more rounds it takes for this fraction to settle down. Low variance corresponds to few short-term aberrations, in other words, very few large payouts. High variance corresponds to a greater frequency of short-term aberrations, in other words, large payouts are more common.

The fractions are going to get closer and closer to 1, that's the law (the so-called "Law of Large Numbers"). It's the journey to "1" that gives the player the experience he wants.