Aaron Brown

Time Enough for Counting

The introduction of the steamboat to the Mississippi River in the early 19th century stimulated an extraordinary period of financial innovation, which culminated in the development of the modern derivative exchange by the end of the century. At the same time, and by the same people, and for the same reason, and hampered by the same moralists, the game of poker was invented.

The history of poker is controversial. As usual, the controversy stems from a failure to define terms clearly. Poker was not created fresh, nor did it evolve from older games. The essential innovation that made poker a different species than any game that came before was a refinement.

The game of poker withstands attempts to accurately simulate it. Perhaps some of the methodology of derivative accounting holds some clues as to how to accurately model its strategies

The vying game

Poker, and poker-like features in other games, is also the only theoretically interesting popular game. There is interesting work done in games like chess and go, but the basic theory is known. Existing algorithms, with enough computing power, could play perfectly. The challenge is to come up with clever approximations that play well using available computer power. But we still have no accepted theory for poker, or bidding in bridge or Monopoly.

Game theory is the standard mathematical approach. The basic principle of game theory is to pick the strategy that has the best expected outcome against the worst-case counter-strategy your opponents might select. While this criterion gives satisfying answers for many games, there are simple examples (the most famous being the Prisoner's Dilemma), in which the game theoretic answer is clearly wrong[1].

The problem with applying game theory to poker (and, as I pointed out in an earlier article, Monopoly[2]) is there are easy collusive strategies[3] for your opponents that virtually guarantee your defeat. The only way to get a game theoretic answer with a chance of winning is to assume independence among your opponents. But ignoring the effect of your strategy on joint behavior of your opponents misses a crucial element of the game.

It will be an important achievement when someone comes up with a complete game-theoretic solution for a realistic poker game. But it is unlikely to be a good strategy in the sense of providing a good return against human players[4]. Moreover, it is likely to be complex and highly dependent on small details of the rules. I doubt it will give much insight into actual poker. At the moment, the best computer programs can beat most casual players, but lose to serious human players.

I think the proper way to analyze poker is through derivative accounting. This article sketches a beginning to that analysis.

Most gambling games are either games of pure chance (like roulette or craps) or games of varying degrees of skill (such as whist or backgammon) with preset stakes or stake computation methods. Any skill involved relates to the play. Chance games traditionally appealed to lower classes, but were seen as a sign of dissipation in an aristocrat. Skill games were considered more refined, even if played for large amounts of money.

A minority of gambling games have "vying" aspects, chances to adjust the stake with impact on the play of the game[5]. The doubling cube in backgammon, for example, or bets like insurance or doubling down in blackjack. Gambling is older than human history and vying elements may well be just as old. In any case, they can be documented hundreds of years earlier than poker.

Poker is unique as a pure vying game. The outcome is entirely luck[6], yet it is most definitely a game of skill. The skill cannot be exercised on the play of the cards, only on the setting of the stake. The other essential element is that stake contributions are equal, unlike all known older vying. This combination of pure luck and pure skill made poker possibly the first democratic game, with equal appeal to all social classes. It would be absurd to claim that no one had ever played a game like this before steam came to the Mississippi, but any earlier attempts died without record. All modern pure vying games can trace a lineage back to the riverboat gamblers. The game that was forged in the 1820s incorporated elements of earlier games from China, Persia, Germany, Spain, England and France (and probably other places as well). No one game or country can claim primary influence. But the essential innovation was pure American.

The future is futures

The Mississippi River drains almost all the land between the Allegheny and Rocky Mountains. This includes roughly 60 per cent of the natural resources of the US and Canada. Alternative transportation routes in and out of this region are very expensive. Before the steamboat, that statement applied to upriver transportation on the Mississippi. The few settlers in the region had to be largely self-sufficient, and could not import much capital equipment. That made them inefficient and sparse. They could ship goods easily downriver to New Orleans (and therefore the world) but without traffic in the reverse direction the region could not develop.

A self-contained plantation that spends five per cent of its economic resources on a cash crop to be exchanged for expensive luxury items and capital goods can be run without much information from the outside world. But a specialty business, say a lumbering operation, which produces only export goods and imports necessities and capital goods, is another matter. Such a business is impractical without sophisticated financial contracts to link the prices of inputs and outputs. Management decisions require frequent detailed price information.

Cities and city-wannabes along the river system began frantic competition to provide the necessary financial services. Chicago triumphed eventually and became the second-largest city in the US[7]. Along with financial services, cities invested in transshipment and processing infrastructure, and amenities for travelers. Competition in entertainment, including prostitution, drugs, alcohol and gambling, was explicit and fierce.

The professional does have an edge, but it does not come from knowledge of game mechanics or probability. It comes from simple application of strategy and psychology

Possibly for the first time in history, this created an important economic need for an honest gambling game. You need professionals for smooth operation of games. Professionals will only play pure luck games in which they have an edge (either disclosed or undisclosed), and will only play games of skill against people of lesser skill. Either of these situations create an aura of dishonesty or exploitation, which drives away customers. Poker has no built-in edge, and it is much harder to cheat at vying games than pure luck or skill games. The professional does have an edge, but it does not come from knowledge of game mechanics or probability. It comes from simple application of strategy and psychology, which any novice can figure out in five minutes. How to play good poker is obvious to anyone, being able to play good poker is a rare attainment.

If you doubt the importance of this, consider the stereotype of a professional gambler. It's very negative. He probably cheats and, if he doesn't, he exploits people to earn a living. We respect that it takes skill, he must play or cheat well, and defend himself against unhappy losers; but we don't admire him.

Compare that to the image of a top-notch poker player. She combines a razor-sharp brain, nerves of steel, shrewd strategic sense, expert knowledge of applied psychology and is a sportsman. She's someone you pick as a leader, not someone you run out of town. She excels in an arena that is a good analogy to life, where the game is simple and fair. In fact, poker may be the only activity popularly considered to evidence both masculinity and intelligence.

Strip poker

The simplest game that combines all essential elements of poker has two players, A and B, each of whom are given random numbers, X and Y respectively, drawn independently from a uniform distribution between 0 and 1. There is a \$1 stake. A has two choices. She can check, in which case both players reveal their numbers, and the one with the higher number takes the \$1. Or she can raise by a fixed amount, \$$R$. If A raises, B has the choice to fold, conceding the pot to A, or call. If B calls, he also contributes \$$R$ to the pot. Then both players reveal their hands and the one with the higher number collects the \$1 + 2 \$$R$ pot.

Suppose A decides to raise with any number greater than P. B will call if his number Y is greater than $P + (1-P)R/(1+2R)$. This makes A's expected value $(1-(1-P)^2R^2/(1+2R))/2$. This quantity is obviously maximized for $P = 1$, meaning that A never raises and collects half the \$1 original stake on average.

Careful readers will have noticed an inefficiency in A's strategy. If her number X is between P and $P + (1-P)R/(1+2R)$, she never wins a showdown. From B's point of view, it doesn't matter if A instead raises with $X < (1-P)R/(1+2R)$ or $X > P + (1-P)R/(1+2R)$. B will still call with the same hands, and win the same percentage of

>

The only edge you get in poker comes from bluffing. Unfortunately, your opponents also have the opportunity to raise you

showdowns. However, A now has stronger average hands when she checks, by $(1-P)R/(1+2R)$. Now A's maximum expected return comes from setting $P=(1+R)/(2+R)$, which gives her an expected value of one-half times $1+R/[(2+R)(1+2R)]$. This expression is maximized when $R=1$, that is when the raise is the size of the existing pot. That's why pot-limit poker games are considered the best tests of skill.

Although simple, this example illustrates some important points about poker.

1. If you raise, you expect to lose money if you are called.
2. Raising only makes sense if you also bluff. Bluffing allows you to more than recoup your showdown losses when your opponent folds against your weak hands.
3. You bluff with your worst hands. Many players prefer to bluff with a marginal hand. There are two ways to look at a bluff. You can say you bluff to encourage opponents to call your raises with marginal hands, bluffing with marginal hands does the reverse. Or you can say you bluff to improve the average quality of the hands you call with, bluffing with your strongest calling hands does the reverse.
4. The only edge you get in poker comes from bluffing. Unfortunately, in real poker, your opponents also have the opportunity to raise you. Without this feature, everyone could play breakeven poker by never raising. Success requires you to exploit your option to raise better than your opponents exploit their option against you.

Poker is square

The following diagram illustrates the situation geometrically for $R=1$. A wins in the yellow areas, B wins in the blue. The darker colors mean the player wins the bet $2R$ as well as the original \$1 pot. In region I, A bluffs and is called. In regions II and VII, A raises and B folds. In region III, A checks and loses, in region IV, A checks and wins. In region V, A raises, B calls and B wins; in region VI, A raises, B calls, and A wins. Note that in regions III, IV and VII, and also II below the diagonal dotted line, the outcome is the same as if there is no betting, as if the original pot is simply awarded to the player with the higher number. In regions V and VI the outcome is the same, but the stake is increased. Since the regions are the same size, the increase offsets. This is where all the drama of the game is, and all the volatility, but it is unsystematic risk that diversifies away to zero in the long run.

In region I, player B wins as he would if there were no betting, but he wins an increased stake. The top of region II, down to the horizontal dotted line represent some of the hands on which B folds the winning hand. It is just large enough to offset the additional losses from region I. A's profit, therefore, is represented by the trapezoid between the two dotted lines in region II. The essential part of the poker game, the only source

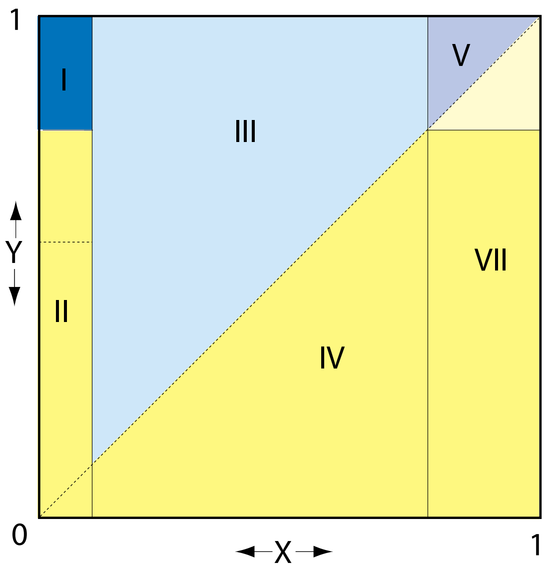

of systematic winnings, is played out with the weakest hands, and hands that are never revealed. This is the deep mystery of poker.

The center, regions III and IV have width 0.5 at $R=0$ which increases to 1 at $R=\infty$. This is the probability at A will not raise. The width of the righthand slice, regions V, VI and VII, also starts at 0.5 at $R=0$, but declines to 0 at $R=\infty$. This is the probability that A will raise with a strong hand, it's also the height of the top slice, regions I, V and VI, the probability that B will call a raise. The probability that A will bluff is represented by the left slice, regions I and II. It has width 0 at $R=0$ and $R=\infty$, it reaches maximum width of 1/9 at $R=1$.

Showdown

I asserted that this model captures the essential elements of poker, now it's time to test that. We need some empirical data. From 1995 to 2001, ten million poker hands were played using an Internet Relay Chat protocol, and all the actions were recorded. One of the frequent players won the 2000 World Series of Poker championship, and many others were successful tournament players. Still other players were serious mathematical poker investigators or the computer programs they created.

This makes a convenient database, but there are two problems. First is that the games were not played for real money. We can get around this to some extent by restricting ourselves to high buy-in games. Anyone could join and get a \$1,000 stake, but some games required bankrolls of up to \$5,000 to play. That means the players were all successful in the lower buy-in games, and they cared about keeping their winnings up. These were long-time active players with a lot of ego invested.

In another way, this population is a better study than real money games. We're not trying to design a system for beating real players, we're trying to understand the game. Real people playing for real money make predictable errors, for example big winners and big losers almost always call far too often on large pots. Emotion distorts strategy. For understanding the game, it might be better to observe mathematicians, computer geeks and serious professionals in a low-pressure situation. Also, the broader audience (including the computer programs) means a broader range of strategies, which should make our results more robust.

Another objection to the database is the hole cards are not recorded unless there was a showdown[8]. I have no problem with this. I want a model based on observable information. Folded and uncalled hole cards are not observable in poker. If two models explain the observable data equally well, but one could be shown to be false based on the unobservable, I don't care. This is one of the differences between applied fields, like finance, and theoretical ones, like economics.

It may seem like a foolish stretch to apply such a simple model to real poker data. In the real game there are many players, not just two, and many rounds of betting. The probabilities are not uniform or symmetrical, except before the first card is dealt, and the probabilities change over the course of the hand. The ratio of bet to pot is not constant. In the square, a raiser never folds. In real poker, a raiser on one round may fold on the next. The square assumes only one call is possible, in real poker every player at the table might call.

But my finance training included the Capital Asset Pricing Model. That took the problem of asset pricing and threw away what everyone thought were the key parameters of the problem. Complex and disparate investor goals were crammed into a two-dimensional grid. Time was ignored, all decisions were single-period. All investors were assumed to agree, and to care only about expected return and standard deviation. Yet the model succeeded spectacularly, clearing away decades of nonsense and error. Of course anomalies were found, but only by people who started with the model. The anomalies focused attention on the complex aspects of asset pricing, the ones that cannot be explained by simple mean-variance optimization. If two dimensions are enough for asset pricing, maybe they're enough for poker as well.

Using the 500,000 hands of $5,000 buy-in Texas Hold'em in the IRC poker database, we find that players raised on fraction 0.4493 of their opportunities. That is consistent with an R of 0.2256, which is roughly the ratio of the average raise to the average pot at the time of raise. A raise probability of 0.4493 implies a win percentage of 0.3452 when the player doesn't raise, the actual percentage is pretty close, 0.3710.

The square then predicts a raise should be called with frequency 0.3795. The actual frequency was 0.6148. That is inconsistent with any positive R. Given the raise and call frequency, the raiser should win 0.5766 of the showdowns, the actual frequency was 0.5187.

A better approach is to recognize each poker bet as what it is, an option on the pot

Finally, the square predicts that a player with an opportunity to raise should get on average 0.5462 of the pot, while a player with an opportunity to call should get 0.4538. The actual numbers were pretty close, 0.5699 and 0.4397. These add up to slightly more than one due to the net pot contribution by players making the blinds. With optimal play, the square predicts a raiser advantage of 0.5349. So raisers appeared to have twice their theoretical advantage, due to players calling too often.

Accounting troubles

The square obviously fails this empirical test, unless we assume the level of play in the database was very low. The trouble is with our accounting. Each raise/call opportunity is treated as a separate decision. The profit and loss from a single pot is allocated to many individual decisions, resulting in multiple counting. There were more call opportunities than raise opportunities, and the frequencies from the population of raise opportunities are not the same as the frequencies from the population of call opportunities. This may have distorted our test, it certainly makes it impossible to drill down to discover if the discrepancies between prediction and data are exploitable lapses in strategy by real players or defects in the model.

A better approach is to recognize each poker bet as what it is, an option on the pot. Calling a bet buys you a contingent claim. Raising a bet knocks out existing options and sets a price at which new options are offered to the old holders. If we could value these options, we could treat each decision on a stand-alone basis. After a player makes a decision, she holds an option. If there are no more decisions, the option either pays off (everyone else folds or she wins a showdown) or expires worthless (she loses a showdown). If she has another decision she either has a new option, or nothing (she folds). In each case, her first decision had a profit or loss, the difference between the option value after her decision is made, and the value after her next decision or the end of the hand.

This gives us a much clearer view into what is really going on in poker. It is easy to fall into error, in poker and with derivatives, by inconsistent counting. Is the purpose of a bluff to encourage raises, in which case it succeeds when it is called, or to win pots with weak hands, in which case it succeeds when it is not called? Both views are true, but holding both at once leads to inconsistent analysis. In general, whether a player wins or loses money in a hand, it is important to apportion the result properly to the decisions that led to it. It's not the final call/fold decision that determines the outcome, it's the cumulative effect of all the steps along the way.

If the goal is to learn how to play good poker, we need a front-office derivative pricing model, one that takes all the available inputs and gives a precise value. Such models are complicated. They give precise prices until something changes in the market, and then they stop working. Traders don't mind replacing models frequently, but accountants and risk managers need consistency. To understand markets or poker, we prefer a simple, robust model. A front office model must be correct for every decision, because the market will punish errors in either direction. An accounting model need only give good average values since we're aggregating it over a large number of decisions.

I chose the simplest reasonable pricing

>

Poker should fit naturally into a tree (although not a binomial one) and we have many good techniques for tree-based derivative pricing

model, after each decision a player remaining in the pot has an option worth the pot divided by the number of players still in the hand. If a raise is outstanding to some of those players, I assume the average fraction of players call it before making the computation. The most important pricing factor I ignore is the strength of the player's hand. One reason is I don't always have the data. Another reason is that strength depends on the conditional expectation of the opponent's hands, which is a complex modeling problem of its own. It makes more sense to start with a simple model, then use the results to refine the data to the point a more complicated model can be devised. We can't work on the anomalies until we learn what they are, what features of play cannot be explained by the simplest strategic principles.

Model performance

To verify the utility of this pricing model, I compared the variance of each decision outcome to the total variance of the hand. The standard deviation of a player's outcome in a hand was $615. There were an average of 10.67 player-decisions per hand, using my pricing model they had a profit/loss standard deviation of $85. A perfect forward-looking model that divided the eventual outcome among each of the decisions would have a profit/loss standard deviation of $615/10.67 = $58. A model with no predictive value would have a profit and loss standard deviation of $615 divided by the square root of 10.67, or $188. So my simple model reduces the standard deviation of decision outcome by 80 per cent of the possible amount[9]. Instead of big surprises at showdown time, this model breaks each hand down into a series of much smaller surprises throughout the play of the hand.

Now I can compute statistics by dollar amount rather than count. Although players raised on 0.4493 of their opportunities, they raised more with smaller option values. Weighted by holding, players raised only 0.3791.

From the standpoint of a player with a raise opportunity, here are the actual outcomes and their predicted frequencies using R = 0.5517, which gave the best overall fit to the data[10].

With decent accounting, the actual play seems to conform pretty closely to a simple strategic model. There are discrepancies. Players raise less often than the model suggests, and get called more often when they do raise. That should mean players with the option to raise win more showdowns, both when they check and when they raise. However the opposite is the case. That suggests a consistent cause, possibly that players sometimes forgo a raise on their strongest hands, particularly in early betting rounds, in hopes of keeping more players in the pot. This is the kind of thing we can investigate once we have numbers that add up.

Outcome	Regions	Actual Frequency	Predicted Frequency
Raise and lose showdown	I and V	0.0834	0.0715
Raise, opponents fold	II and VII	0.2539	0.2786
Don't raise, lose	III	0.3878	0.3607
Don't raise, win	IV	0.2331	0.2474
Raise and win showdown	VI	0.0418	0.0418

Of course, the model needs further testing. One test would be to see if it gives more accurate estimates of player skill than simple alternatives like average profit per hand. Another would be to see if actual and predicted frequencies move in the same way over time or different situations.

This is obviously only a sketchy start to analyzing the game of poker. It uses standard financial tools in the hope of coming up with more meaningful answers than previous approaches have produced. A basic understanding of the strategy, and a clear accounting, should lead to useful data for developing a pricing model good enough for play. Poker should fit naturally into a tree (although not a binomial one) and we have many good techniques for tree-based derivative pricing.

The other intent here is to illustrate how transparent accounting for derivatives can transform a confusing mass of apparently conflicting data into a simple picture. Whether, in fact, that picture is correct is a question for further investigation. But the pricing model need not be perfect for the accounting statements to be useful.

Aaron Brown is a Vice President in the Credit Risk department of Morgan Stanley. This article expresses his personal views, which are not necessarily those of his employer or any other entity.

REFERENCES

1 Not mathematically wrong, game theory is internally consistent, but not a satisfying answer. That suggests that game theory asks the wrong question about these games.
2 *Wilmott* issue 3, January 2003 and subsequent issue
3 Collusion is the whole point of Monopoly. Explicit collusion is against the rules of poker, but implicit collusion is an important strategic factor.
4 It will not lose in two-player, zero-sum versions of poker, but it will probably not win quickly against average human players; and I don't think it will win at all with more than two players or if non-zero-sum aspects of poker are included.
5 The usual definition of "vying" games is games where players compete to win a central pot. But the key element from a theoretical viewpoint is the competition consists of adjusting the stake.
6 The draw decision in draw poker notwithstanding
7 But Minneapolis won the purely financial part of the race.
8 This serious respect for privacy is refreshing. It used to be the norm for the Internet, it's very sad how quickly we have moved to the opposite extreme.
9 The more common measure is R2, the fraction of variance explained by the model.
10 This is considerably higher than the average ratio of bet to pot size. However, given the high degree of abstraction of the square, that is not a concern. In some strategic sense, it can represent the average gain or loss multiplier that occurs when a player raises.