Is poker different from flipping coins?

Tony Henke
January 19th, 2007
Washington University in Saint Louis

I. Introduction

Most poker players believe that "big stacks," those players with relatively more chips than their opponents have a strategic advantage while players with relatively fewer chips, "short stacks," are at a strategic disadvantage. This paper first develops a definition of 'strategic advantage' in the context of a poker tournament. Specifically, it will be shown that in a non-strategic environment a player's probability of victory is exactly equal to the proportion of their chips to the table as a whole. The paper then examines data from four years of World Poker Tour tournaments and finds weak evidence for the existence of a strategic disadvantage of being a "short stack" and a corresponding strategic advantage of being a "big stack." The paper continues by examining which features of a poker game fail to give rise to these strategic implications of relative chip position. Both Borel's and Von Neumann's second game-theoretic models of poker are analyzed and found to be identical to a non-strategic game (such as flipping coins) when optimal play is used. Finally, the paper develops a theory of what causes relative chip position to impact final outcomes in a poker tournament. Several testable hypotheses are produced.

Until now, this analysis has never been done. As poker, and poker tournaments in particular, become more popular, this analysis will continue to grow in importance. At least one million dollar poker tournament takes place every month. In 2006 a poker tournament awarded $12 million to first place. A great deal of money changes hands. This analysis will help us understand how that transfer of money is taking place. This analysis may also have applications outside of the realm of poker. The analysis in this paper applies to a class of strategic interactions – those with winner take all payoffs and a stochastic process that determines outcomes.

II. Flipping Coins: a non-strategic game

What is meant by a "strategic" advantage or disadvantage? To answer this question let us first examine what we can expect in a purely non-strategic setting. Imagine a game between two players where each player starts with some positive integer number of chips (not necessarily the same number for each player) and in every round of play each player tosses a single chip into the pot and then a coin is flipped and whatever is in the pot is given to the player who wins the coin toss. This process is repeated until one player has all of the chips and is declared the winner. There is no strategic play of any kind. In this game, and, as we will soon show, in many other games like it, a player's probability of victory [hereafter referred to as P_i(win)] is exactly equal to the proportion of their chips to the total number of chips in play (hereafter referred to as p_i). A proof of this proposition follows[1].

> We shall calculate $u_i = \pi_i(C_0)$ and $v_i = \pi_i(C_n)$, the probabilities that a player starting with i chips ultimately enters the absorbing states 0 (no chips) and n (all n chips), respectively. The system of equations becomes
>
> $$u_1 = q + pu_2\,,$$
> $$u_i = qu_{i-1} + pu_{i+1}\,, \qquad 2 \le i \le n-2.$$
> $$u_{n-1} = qu_{n-2}\,,$$
>
> We try a solution of the form $u_r = x^r$. Substituting in the middle equations and removing common factors leads to
>
> $$px^2 + q = x.$$

[1] The proof given is adapted from one in *A First Course in Stochastic Processes* by Samuel Karlin and Howard M. Taylor, second edition, p. 92-3.

There are two solutions, $x = 1$ and $x = q/p$. Thus the quantities $u_r = A + B(q/p)^r$, $r = 1, 2, \ldots, n-1$, satisfy the middle equations above for any values of A and B. We now determine A and B so that the first and last equations are fulfilled. (If $q = p$, the solution $x = 1$ is a double root of $px^2 + q = x$, and one then has to replace $(q/p)^r$ by r.) In the case $q \neq p$ this leads to the conditions

$$A + B(q/p) = q + p(A + B(q^2/p^2))$$

or, simplifying,

$$A = 1 - B$$

and

$$A + B(q/p)^{n-1} = q(A + B(q/p)^{n-2}) \quad \text{or} \quad p^n A + q^n B = 0.$$

Solving, we get

$$A = q^n/(q^n - p^n), \qquad B = -p^n/(q^n - p^n)$$

Combining, we have

$$u_r = [(q/p)^n - (q/p)^r]/[(q/p)^n - 1] \quad \text{if } q/p \neq 1.$$

If we have a perfectly weighted coin $q = p$. If $q = p$, we find similarly that $A = 1$, $B = -1/n$ so that

$$u_r = (n-r)/n \qquad \text{when } p = q.$$

Since $v_i = 1 - u_i$ (because one of the absorbing states must eventually be reached), we get

$$v_r = r/n.$$

The left side of this last equation is $P_i(\text{win})$ and the right side is p_i. ■

Thus, the reader has aggressively been made to believe that $P_i(\text{win})$ truly equals p_i in the case of two players with a per-turn expected value of chips equal to zero. However, we have not shown whether this result holds in the case of three or more players. Due to insufficient mathematical prowess on the part of the author, no proof will be given to support this result. However, a computer simulation of the three player version of the flipping coins game does in fact confirm that $P_i(win) = p_i$ in the three player case.[2] Cases with four or more players are assumed to conform to this rule.

A "strategic" advantage then is one that deviates from this "non-strategic" standard. Specifically, if we find that poker tournaments significantly deviate from $P_i(win) = p_i$, a deviation occurring due to the nature of the strategic interactions in play, we will have found a "strategic" advantage.

III. Poker: are we still flipping coins?

To determine if poker tournaments deviate from our newfound rule we examine data from 58 World Poker Tour tournaments taking place between June 2002 and April 2006[3]. For each tournament two cross-sections of time are examined: first, when there are six players left, constituting the "final table," and second, when there is only one player left, with all of the chips. If poker tournaments follow the $P_i(win) = p_i$ rule, we expect, ceteris paribus, a player with five percent of the chips to win five percent of the time, and a player with fifty percent of the chips to win fifty percent of the time. We break up the players into seven groups, by their proportion of the total chips when the final table begins. For each group, we expect the average proportion of chips of that group to equal the proportion of the time that a member from that group wins their tournament.

Table 1: Descriptive Statistics

Group	Win %	Average p_i	Win % - avg. p_i
0 – 4.99%	0 (0 of 24)	3.94%	-3.94%

[2] The code for this program is included in the Appendix. The main function is not included.
[3] The data discussed can be found at www.worldpokertour.com.

5 – 9.99%	6.41 (5 of 78)	7.62%	-1.21%
10 – 14.99%	10.13 (8 of 79)	12.34%	-2.21%
15 – 19.99%	17.91 (12 of 67)	17.51%	0.40%
20 – 24.99%	25.64 (10 of 39)	22.16%	3.48%
25 – 29.99%	27.27% (6 of 22)	27.00%	0.27%
30+%	43.59% (17 of 39)	38.65%	4.94%

These data conform precisely to how we would expect the players to fare if big stacks have a strategic advantage and short stacks are at a strategic disadvantage: the small stacks win less often and the big stacks more often than they should, with the effects generally becoming larger as one moves to the extremes. (In the case of six person games an average chip stack is 16.66%. So a player with less than 16.66% is a small stack, and a player with more than 16.66% a big stack.) These effects are economically significant. A quick examination of the 2006 World Series of Poker Main Event (the aforementioned $12 million tournament) final table shows why. The chip leader entered the final table with 34.69% of the chips in play. The anecdotal evidence here puts his chance of victory at somewhere near 39.63%. That difference would be worth $592,800.

Of course, there is an alternative explanation that just as ably explains what we are seeing. Namely, that skill is endogenous in this system. More specifically, it is reasonable to believe that highly skilled players will arrive at the final table with more chips than their less skilled counterparts and then, because they are more skilled than their opposition, will win more often than the $P_i(\text{win}) = p_i$ rule predicts. To determine if the endogeneity of skill explains all of the deviation from the rule in the data we construct a proxy variable for skill. We use whether or not a player had previously reached a World Poker Tour tournament final table. This new proxy variable correlates to p_i with strong statistical significance.[4] To test if big stack advantages persist after controlling for skill, we run a probit model on a player's proportion of chips at the start of play at the final table (*startprop*), a dummy variable for if the player has the fewest chips at the table (*sixth*), and our proxy variable for skill (*experience*). We estimate:

$$P_i(\text{win}) = \Phi(\beta_0 + \beta_1 \textit{startprop} + \beta_2 \textit{sixth} + \beta_3 \textit{experience})$$

If, controlling for experience, our $P_i(\text{win}) = p_i$ rule holds we expect *startprop* to have a positive coefficient and *sixth* to have no significant coefficient whatsoever. We find that both *experience* and *sixth* are both economically and (weakly) statistically significant.

Table 2: Estimate of Probit Equation[5]

Independent Variable	**Coefficient (standard error)**
intercept	-1.796747 (.20879) p-value: 0.000
startprop	3.976571 (.85811) p-value: 0.000
sixth	-.698127 (.43638) p-value: 0.110
experience	.2749549 (.17376) p-value: 0.114

Of course, this analysis is far from conclusive. We have examined only one particular proxy for skill in conjunction with one particular measure of the disadvantage of short stacks. There is some concern that those variables may not have been

[4] The exact relationship is detailed in the Appendix.
[5] These results are nearly identical when using total number of final tables reached as the proxy for skill. Those results are reported in Table 3 in the Appendix.

appropriate choices. First, let us consider the *sixth* variable. It was chosen because it is the player with the smallest stack who is supposed to be at the greatest disadvantage if a strategic disadvantage of short stacks exists. Of course, not all players in sixth are at a terrible disadvantage (one player had over 12.6% of the chips and was still the smallest stack at the table.) Nevertheless, sixth seems to be a good choice because it should have no explanatory power if the $P_i(\text{win}) = p_i$ rule holds. The more problematic choice is the proxy for skill. It is a virtual certainty that skill does not perfectly correlate with previously reaching a World Poker Tour final table. In fact, since skill is variable over time and there is currently no effective method of ranking poker players, it seems that no proxy for skill would be immune to criticism. What can be done? We must construct a more complete model of the "strategic" advantages of big stacks in poker tournaments. The next section seeks to accomplish that goal.

IV. Causes and Testable Implications

Before we discuss what gives rise to deviations from the $P_i(\text{win}) = p_i$ rule, let us first examine what fails to do so. In this regard computer simulations of both Borel and Von Neumann poker have been useful.[6] The computer is programmed to simulate optimal play between two players until one player wins. This process is then repeated millions of times for every possible combination of starting chip stacks. In both these versions of poker, one player is at a disadvantage to the other (neither player's per turn expected value is zero). To solve this problem, the computer simulations randomize which player goes first in each round. We find that both Borel poker and Von Neumann

[6] The model's discussed are drawn from *Fun and Games: A Text on Game Theory* by Ken Binmore. In this discussion optimal play refers to the optimal play discussed by Binmore and always constitutes a Nash equilibrium.

poker follow the $P_i(win) = p_i$ rule. This eliminates bluffing and variable bet sizes as possible sufficient conditions for deviations from our rule to take place. It should be noted that in the computer simulations the players are maximizing their expected value of chips per round rather than their probability of victory, this may or may not change optimal play strategies.

What then might give rise to deviations from the rule? We find two theoretically compelling reasons why the advantage of "big stacks" exists in poker tournaments. First, and most compellingly, in poker tournaments, short stacks simply aren't trying to maximize their probability of winning. The standard payout structure (how much money a player gets for finishing first, second, third, and so on) often rewards players with relatively few chips to hold on for as long as possible, hoping that someone else gets eliminated before they do, and thereby move up the pay ladder. This strategy does not maximize P(win) and thus gives "big stacks" a better chance at victory.

Second, is what we call the "asymmetric restriction of choice." When one player's optimal move would be to raise 500 chips, but only has 300, his choice set is restricted, and this hurts his chances of victory. Of course, this is only a problem if this restriction of choice occurs asymmetrically: if all players are equally restricted then no player has an advantage over another. Which players will be asymmetrically restricted? Those with the fewest chips.

Of course, there are probably other factors at play as well. Some players may not try to maximize their probability of victory or their expected cash payout. Instead they might try to maximize the gutsiness of their TV image (since many of these tournaments are televised) or simply avoid being embarrassed on national television and become overly cautious. Psychological factors, such as intimidation, may also play a significant role.

This discussion of causes gives rise to a number of testable hypotheses. First, our model predicts that in tournaments with a steeper pay scale (more of the payoff going to first place) the strategic advantage of big stacks will be less than those with a flatter pay scale. Second, our model predicts that forms of poker games with more open ended betting systems (no-limit betting) will have larger advantages for big stacks than those with more restrictive betting systems (pot-limit or strict limit betting). Third, and most surprisingly, our model predicts no strategic advantage for big stacks in two person poker games. In a two person game, maximizing P_i(win) is equivalent to maximizing expected payout. Also, in a two person game, a restriction of one player's choice set necessarily restricts the other player's choice set, so that no restriction of choice is asymmetric. Unfortunately, the data set used in this study is too limited to analyze these predictions. On the other hand, with the number and variety of poker tournaments that take place every day, the possibility for future analysis seems very strong.

V. Conclusion

We have shown that in a purely non-strategic environment a player's probability of victory in a tournament setting is exactly equal to the proportion of their chips to all the chips in play. We have provided evidence that poker tournaments deviate from this rule, although there is ample concern that the endogeneity of skill in the data we examined may swamp these results. Most importantly, we have constructed a rubric for analyzing this problem in the future. We have established a well defined rule for analysis and, by delineating a series of causes that may give rise to deviations from this rule, we have produced several testable hypotheses.

Sources

Binmore, Ken. *Fun and Games: A Text on Game Theory*, D.C. Heath and Company, Lexington, Massachusetts, 1992.

Karlin and Taylor. *A First Course in Stochastic Processes*, second ed., Academic Press, New York, 1975.

Appendix

Table 3: Estimated Probit Equation
with *experience* as total final tables played

Independent Variable	**Coefficient (standard error)**
intercept	-1.824199 (.21915) p-value: 0.000
startprop	4.000127 (.85866) p-value: 0.000
experience	.0670308 (.04770) p-value: 0.160
sixth	-.74849 (.43156) p-value: 0.122

Relationship between number of final tables reached and starting p_i

$$startprop = .1398423 + .0124465*experience$$
$$(.0089) \qquad (.0032)$$

Relationship between having reached a final table previously and starting p_i

$$startprop = .1512806 + .0352259*expdum$$
$$(.0074) \qquad (.0112)$$

Three player flipping coins program

```
/* Three person "poker" */
int three(int x_chips, int y_chips, int z_chips, int N)
{
  int i, n, r, m, p, x_wins=0, y_wins=0, z_wins=0, store_x, store_y, store_z;
```

```
  store_x = x_chips;
  store_y = y_chips;
  store_z = z_chips;

  srand(time(NULL));

 for(i=0; i<N; i++)
 {
  x_chips = store_x; y_chips = store_y; z_chips = store_z;

  while(x_chips > 0 && y_chips > 0 && z_chips > 0)
  {
   r = rand()/(RAND_MAX + 1.0) * 3;

   if(r == 2)      {x_chips += 2; y_chips--; z_chips--;}
   else if(r == 0) {x_chips--; y_chips += 2; z_chips--;}
   else if(r == 1) {x_chips--; y_chips--; z_chips += 2;}
  }

  while((x_chips > 0 && y_chips > 0) || (x_chips > 0 && z_chips > 0) || (y_chips > 0 && z_chips > 0))
  {
  if(x_chips == 0)
  {
   while(y_chips > 0 && z_chips > 0)
   {
   n = rand()/(RAND_MAX + 1.0) * 2;

   if(n == 0) {y_chips++; z_chips--;}
   if(n == 1) {y_chips--; z_chips++;}
   }
  }

  else if(y_chips == 0)
  {
   while(x_chips > 0 && z_chips > 0)
   {
   m = rand()/(RAND_MAX + 1.0) * 2;

   if(m == 0) {x_chips++; z_chips--;}
   if(m == 1) {x_chips--; z_chips++;}
   }
  }

  else if(z_chips == 0)
  {
   while(y_chips > 0 && x_chips > 0)
   {
   p = rand()/(RAND_MAX + 1.0) * 2;

   if(p == 0) {x_chips++; y_chips--;}
   if(p == 1) {x_chips--; y_chips++;}
   }
  }
  }

  if(x_chips == 0 && y_chips == 0) z_wins++;
  if(y_chips == 0 && z_chips == 0) x_wins++;
  if(x_chips == 0 && z_chips == 0) y_wins++;
 }

  return(y_wins);
 }
```