

Poker Superstars: Skill or Luck?

Similarities between golf — thought to be a game of skill — and poker

Rachel Croson, Peter Fishman, and Devin G. Pope

“Why do you think the same five guys make it to the final table of the World Series of Poker every year? What are they, the luckiest guys in Las Vegas?”

—Mike McDermott (Matt Damon in the 1998 film “Rounders”)

The popularity of poker has exploded in recent years. The premier event, the World Series of Poker Main Event, which costs \$10,000 to enter, has increased from a field of six in 1971 to 839 in 2003 and 5,619 in 2005. Broadcasts of poker tournaments can frequently be found on television stations such as ESPN, Fox Sports, the Travel Channel, Bravo, and the Game Show Network. These tournaments consistently receive high television ratings.

Poker also has garnered the attention of many influential academics. It served as a key inspiration in the historical development of game theory. John Von Neumann and Oskar Morgenstern claim that their 1944 classic, *Theory of Games and Economic Behavior*, was motivated by poker. In the text, they described and solved a simplified game of poker. Other famous mathematicians/economists such as Harold Kuhn and John Nash also studied and wrote about poker.

For all its popularity and academic interest, the legality of poker playing is in question. In particular, most regulations of gambling in the United States (and other countries) include poker. In the United States, each state has the authority to decide whether it is legal to play poker for money, and the regulations vary significantly. In Indiana, poker for money is legal only at regulated casinos. In Texas, poker for money is legal only in private residences. In Utah, poker for money is not legal at all. The popularity of online poker for money has raised further questions about the right (or ability) of states to regulate this activity. At the national level, the U.S. Department of Justice recently stated that the Federal Wire Act (the Interstate Wire Act) makes online casino games illegal (in addition to sports wagering), although the U.S. Fifth Court of Appeals subsequently ruled that interpretation incorrect.

That said, there are heated arguments on both sides of the regulation debate. Those in favor of regulating argue that poker is primarily a game of luck, such as roulette or baccarat, and that it should be regulated in a manner similar to those games. Those in favor of lifting regulations argue that it is primarily a game of skill—a sport such as tennis or golf—and it should not be regulated at all. So, is professional poker a game of luck or skill?

Several ‘star’ poker players have repeatedly performed well in high-stakes poker tournaments. While this suggests skill differentials, it is far from conclusive. In how

many poker tournaments have these stars participated in which they did not do well? Furthermore, even if poker competition among top players were random, we would expect a few players to get lucky and do well in multiple tournaments.

We use data from high-stakes poker and golf tournaments and identify the rates at which highly skilled players are likely to place highly. We use golf as a comparison group, as it is an example of a game thought to be primarily skill-based. If the data from golf and poker have many similarities, especially in terms of repeat winners, those data could suggest poker is equivalently a game of skill.

Data

In a large poker tournament, individuals pay an entry fee and receive a fixed number of chips in exchange. These chips are valuable only in the context of the tournament; they cannot be used elsewhere in the casino or exchanged for money. Players are randomly assigned to tables, typically including nine players and one professional dealer. Players remain in the tournament until they lose all their chips, at which point they are eliminated. Some tournaments include a “rebuy” option, where players can pay a second entry fee and receive more tournament chips. Others include an “add-on” option, where they can pay a small extra fee (often used to tip the dealers) and receive more tournament chips. At some point during the tournament, these options disappear. As players lose their chips, they are merged to create a roughly equal distribution of players per table.

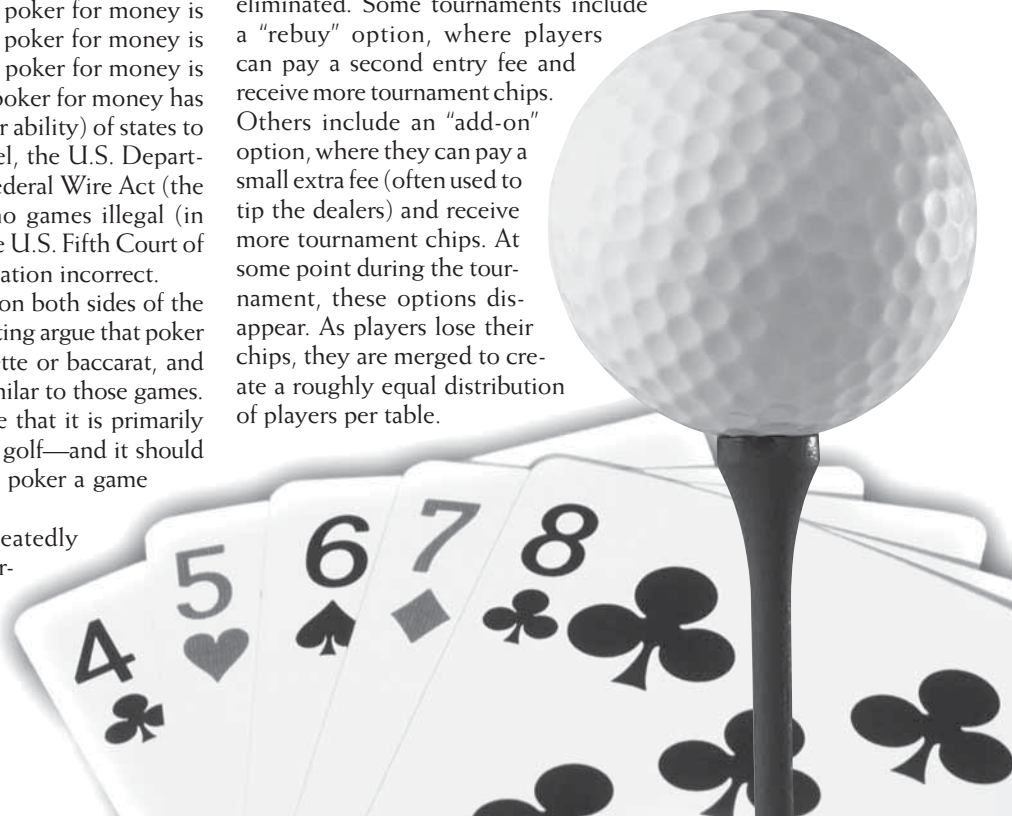


Table 1—Descriptive Statistics

	Poker	Golf
% with 1 top 18 finish	70.1	24.3
% with 2 top 18 finishes	14.7	14.7
% with 3 top 18 finishes	6.9	18.9
% with 4 or more top 18 finishes	8.3	42.2
Number of tournaments	81.0	48.0
Number of individuals	899.0	218.0

Note: Poker summary statistics represent data from all high-stakes (\$3,000 or greater buy-in) limit and no-limit Texas Hold'em tournaments between 2001 and 2005 from the World Series of Poker, the World Poker Tour, or World Poker Open. The 899 players represent those who finished in the top 18 of at least one of these 81 tournaments. The golf summary statistics represent data from all Professional Golfers' Association tournaments in 2005. The 218 players represent those who finished in the top 18 of at least one of these 48 tournaments.

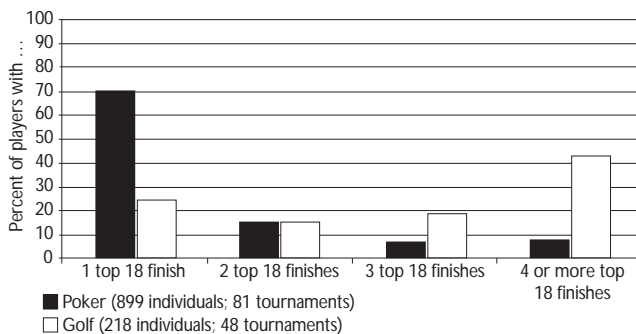


Figure 1. Descriptive statistics

Identifying skill discrepancies among top poker players is complicated by the lack of precise tournament data. The lists of entrants for large poker tournaments are not available, and outcomes are typically only recorded for players who finish in the final two or three tables. Thus, it is not possible to know the total number of tournaments for which a given player has participated. In our data, we have 899 poker players who finish in the top 18 of a high-stakes tournament at least once. The average tournament has between 100 and 150 entrants. Thus, a given person has an 11%–17% chance of entering a given tournament. Due to the lack of data on tournament attendance, it is impossible to know if players who frequently show up at final tables are more skilled than other players, or if they simply play in more tournaments.

To circumvent this selection issue, we employ a strategy that focuses on individuals who finished in the top 18 in high-stakes tournaments (the two final tables). As data are typically available for all players who finish in the top 18 of a given tournament, we can overcome the selection issue by focusing on just these individuals. Thus, while we are unable to identify the number of tournaments an individual has played in, we are able to identify the number of times a player has played in a tournament of 18 players. We can analyze whether certain players consistently outperform other players conditional on being in the top 18, or whether the outcomes appear to be random.

We use data from limit or no-limit Texas Hold'em tournaments that are part of the World Series of Poker, World Poker Tour, or World Poker Open. Texas Hold'em is a variant of poker in which all players are given two personal cards and there are five community cards that apply to all players' hands.

The goal is to make the best five-card hand from the two personal cards and the five community cards. Betting occurs after each player receives his or her cards, again after three of the five community cards are revealed, again after the fourth community card, and finally after the fifth community card. In limit Texas Hold'em, the bet amounts each round are fixed; whereas, in no-limit Texas Hold'em, a player can wager as many chips as he or she wants above a set minimum wager.

Using information gleaned from *pokerpages.com*, we record outcomes for the top 18 finishers of tournaments since 2001 that had at least a \$3,000 buy-in. For a small number of tournaments after 2001 (and for all tournaments prior to 2001), the top 18 finishers were not recorded or not available and, thus, were not included in the analysis. A total of 81 separate poker tournaments fit these criteria. Table 1 presents summary statistics for the poker players in these tournaments.

We similarly collect data for all 48 Professional Golfers' Association (PGA) tournaments in 2005. We record the name and final rank of each player who finished in the top 18 in each tournament. In golf, there are often ties. We record an average rank for these situations (i.e., if two players tie for third place, each player is given a rank of 3.5). Table 1 provides summary statistics for the golf players in these tournaments.

Empirically, we are interested in using information about past performance to predict the outcome of individuals in a given tournament, conditional on them being among the final 18 contestants. Our main outcome variable will be the individual's rank in this tournament of 18 (1 through 18), with lower ranks being better. If we are able to predict an individual's rank in this tournament of 18 based on their past performance, this implies that outcomes are not random. We also will compare our predictive ability between golf and poker (see Figure 1).

Methods

We fit the data using ordinary least squares regression. Thus, given standard notation, the coefficients (β) are estimated such that $\hat{\beta} = (X'X)^{-1}X'y$. The variance of this estimator is $(X'X)^{-1}X'\sum X(X'X)^{-1}$, where $\sum = E[(y - Ey)(y - Ey)']$. Typical OLS estimation assumes homoscedasticity and the independence of error terms across observations. These assumptions imply that $\sum = \sigma^2I$, thus the variance of the OLS estimator can be represented as $(X'X)^{-1}\sigma^2$.

One might worry that one or more of these assumptions will fail in our case. For example, in many cases, we have observations for the same player across different tournaments in our data set. Thus, the error terms on these observations may not be independent. While we present typical OLS coefficient estimates, we adjust the standard errors in our model to account for the possibility of heteroskedasticity and that the error terms on observations from the same player may not be independent of each other. In other words, the standard errors we present are "robust" and "clustered" at the player level. Mathematically, this implies that, instead of assuming $\sum = \sigma^2I$, we allow \sum to have off-diagonal terms that are not zero and to have diagonal terms that are different from each other. These terms are simply represented by the appropriate products of the residuals $((y - Ey)(y - Ey)')$ when calculating the standard errors on our OLS coefficients. The classic 2002 econometric text by Jeffrey Wooldridge supplies an even more detailed description of this process.

Table 2—OLS Regressions with Robust Standard Errors: Rank (1st–18th)

	Poker			Golf		
	(1)	(2)	(3)	(4)	(5)	(6)
Experience	-0.781 [.278]*			-1.420 [.382]*		
Finishes		-0.225 [.098]*			-0.222 [.089]*	
Previous Rank			0.203 [.050]*			0.033 [.056]
Constant	9.810 [.173]*	9.707 [.166]*	7.189 [.490]*	10.270 [.331]*	9.743 [.253]*	8.566 [.595]*
R-Squared	0.5%	0.9%	2.8%	1.6%	1.2%	0.1%
Observations	1494	1494	595	811	811	586

Note: Columns (1)–(6) present coefficients and robust standard errors clustered at the player level from regressions with finishing rank (1st–18th) as the dependent variable. Experience is an indicator that equals one if the player had previously finished in the top 18 of a tournament in our sample (0 or 1). Finishes is the number of times the individual has previously appeared in the top 18 of a tournament in our sample (ranges 0 to 10 for poker and 0 to 14 for golf). Previous Rank indicates the average rank for all previous tournaments in which the player finished in the top 18 in our sample (ranges from 1 to 18).

* significant at 5%

Our baseline econometric specification is $Rank_i = \alpha + \beta X_i + \varepsilon_i$, where $Rank_i$ is the rank at the end of a tournament for player i and X_i is a measure of previous tournament performance for player i . We will examine three measures of previous tournament performance to see how well they explain current rank. Our first measure is called “experience,” and it records whether a player has previously finished in the top 18 of another tournament prior to the one whose rank we are predicting (thus, it takes the value of either 0 or 1). Our second measure is called “finishes,” and it records the number of times a player has previously finished in the top 18 of another tournament prior to the one whose rank we are predicting (this variable ranges 0 to 10 for poker and 0 to 14 for golf). Our third measure is called “previous rank,” and it records the average rank of a player in all previous tournaments in which the player finished in the top 18.

To assess the sensitivity of results to the chosen model, the analysis is repeated using an ordered probit model, a regression format designed to handle situations where the dependent variable has several discrete categories ordered in some way (such as rank). In comparison with least squares, the ordered probit is more robust, but also more computationally intensive. Results from the ordered probit are the same as those we find using OLS. We present OLS coefficients in this paper for purposes of clarity and ease of interpretation. (Results from the ordered probit are available from the authors.)

We will conduct two types of statistical tests. The first focuses on only the poker data. If there are no skill differentials among poker players, we would expect the coefficient on experience, finishes, and previous rank to be statistically insignificant. This would indicate that, conditional on making it to the final 18, one’s final rank is not influenced by previous tournament performance. However, if some players are more skilled than others, we would expect to find statistically significant and negative coefficients for experience and finishes in the above specifications (past experience and success should be associated with a reduction in rank [e.g., from 7th place to 6th place]) and a positive coefficient for previous rank (a higher rank in previous tournaments of 18 should be associated with a higher rank in this one).

A second test we use is a comparison of the results between golf and poker tournaments. We compare the size of the coefficients of interest. If golf has statistically larger coefficients than poker (in absolute value), then there is more skill in golf than in poker. If the coefficients in golf are not statistically different than those in poker, we will conclude that poker has similar amounts of skill (and luck) as golf.

Results

Table 2 presents the results. Robust standard errors are presented in brackets below the coefficient values. Our first analysis involves simply looking at the poker data and identifying whether previous success predicted current success. Clearly it does. The coefficient on experience (whether a player has previously finished in the top 18) is significantly and negatively correlated with a player’s rank in the given tournament, suggesting an increase in finishing (–.78 ranks, $p < .01$). The coefficient on finishes (the number of times a player has previously finished in the top 18) is significantly and negatively correlated with a player’s rank in the given tournament, suggesting an increase in finishing as well (–.22 ranks, $p < .05$). The coefficient on previous rank (the average rank for the player in previous tournament finishes) is significantly and positively correlated with a player’s rank in the given tournament (.20 ranks, $p < .01$). These results clearly suggest poker is, at least somewhat, a game of skill.

But, how much skill? A comparison with golf can illuminate this question. If we compare the estimated coefficients on the experience variable, we find that these coefficients are not statistically different from each other ($t = 1.35$, $p > .05$). Similarly, there are no statistically significant differences between the estimated coefficients on finishes ($t = 0.10$, $p > .05$). For the final measure of previous performance, previous rank, the coefficient for poker is statistically larger than the coefficient for golf ($t = 2.24$, $p < .05$).

Figures 2a and 2b show two of these relationships graphically. Figure 2a depicts the average rank in a given tournament as a function of finishes. Figure 2b depicts the average rank in a given tournament as a function of previous rank. Both show the average rank, as well as a linear fit of the data. These figures

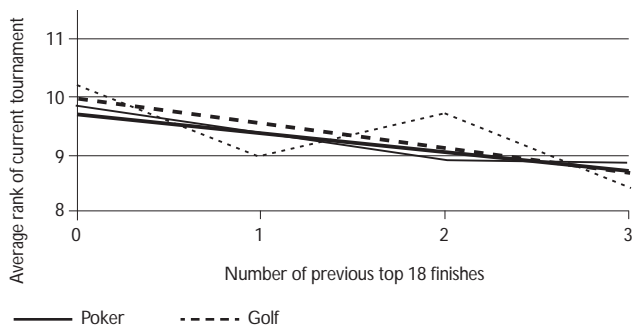


Figure 2a. Relationship between rank and number of previous top 18 finishes

Note: This figure depicts the average rank for poker and golf players who finish in the top 18 for a given poker or golf tournament. The number of previous top 18 finishes (finishes in the analyses above) is the total number of previous top 18 tournament finishes for each player in our sample (0, 1, 2, 3, or 4 or more). The straight lines indicate linear fits of the data. Note that the slope of these lines is not exactly the same as the slope from the regressions, as we have simplified the variable finishes for ease of display.

visually depict our regression results from Table 2: Both poker and golf show a significant negative relationship between current rank and finishes. Poker, but not golf, shows a significant positive relationship between current rank and previous rank.

That said, the R-squared values for the regressions we report for both poker and golf are extremely low (ranging from .1%–2.8%). This suggests that, in general, it is very difficult to predict the ordering of a given set of poker or golf players who finish in the top 18 of a given tournament. Although our measures of previous performance are statistically significant predictors of current performance, they still only explain a small amount of the overall variation that exists in poker and golf, as one might expect to be the case in many sports and games, especially those with explicit randomization such as poker.

Discussion and Conclusion

We present evidence of skill differentials among poker players finishing in one of the final two tables in high-stakes poker tournaments. We show two main results. First, there appears to be a significant skill component to poker: Previous finishes in tournaments predict current finishes. Second, we find the skill differences among top poker players are similar to skill differences across top golfers.

While our analysis provides evidence for skill being a factor in poker (significant regression coefficients), the current evidence needs further support from other analyses (primarily because of the small R-squared). Thus, this analysis should be considered a first attempt to answer this question, and we hope this article will stimulate further efforts.

A second limitation of the present study is that models do not specifically account for repeated observations from some players in the analyses and that results within a tournament for different players are correlated. These aspects of the data would impact standard errors in analyses, but perhaps not too strongly. First, most players appear in just a few tournaments, so they are used not many times. In poker, this is especially true. Second, few pairs of players appear in the same pairs of tournaments. Thus, the amount of information that could be learned by modeling ranks for pairs of players is quite limited. This is especially true in poker.

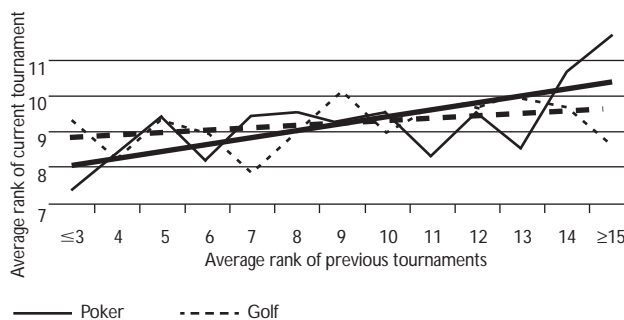



Figure 2b. Relationship between rank and average previous rank

Note: This figure depicts the average rank for poker and golf players who finish in the top 18 for a given poker or golf tournament. The average rank of previous tournaments (previous rank in the analyses above) is the average rank the individuals achieved in previous tournaments in which they made the top 18. The straight lines indicate linear fits of the data.

While we provide evidence for the impact of skill on poker outcomes, we cannot provide insight regarding the cause of this result. We do not know, for example, if poker players are skilled because they are good at calculating pot odds and probabilities, good at reading their opponents' tells (subtle physical cues that signal the strength of a player's hand), or simply better at bluffing or intimidating the rest of the table. Similarly, we cannot identify the source of skill differentials at golf. Are these due to better driving skills, better putting skills, or better strategies? Further research (with more data) is clearly needed to identify which skills are at play. However, our evidence argues that at least some portion of poker outcomes are due to skill, and we hope this will illuminate the raging regulatory debate in the United States and elsewhere. 

Further Reading

- Emert, John and Umbach, Dale (1996) "Inconsistencies of 'Wild-Card' Poker." *CHANCE*, 9(3):17–22.
- Haigh, John (2002) "Optimal Strategy in Casino Stud Poker." *Journal of the Royal Statistical Society, Series D: The Statistician*, 51:203–213
- Heiny, Eric. (2008) "Today's PGA Tour Pro: Long but Not so Straight." *CHANCE*, 21(1):10–21.
- Kuhn, H.W. (1950) "A Simplified Two-Person Poker." *Contributions to the Theory of Games*, I. H.W. Kuhn and A.W. Tucker (eds.). *Annals of Mathematics Studies*, Number 24. Princeton, New Jersey: Princeton University Press.
- Nash, J.F. and Shapley, L.S. (1950) "A Simple Three-Person Poker Game." *Contributions to the Theory of Games*, I. H.W. Kuhn and A.W. Tucker (eds.). *Annals of Mathematics Studies*, Number 24. Princeton, New Jersey: Princeton University Press.
- Von Neumann, J. and Morgenstern, O. (1944) *Theory of Games and Economic Behavior*. Princeton, New Jersey: Princeton University Press.
- Wooldridge, J.M. (2002) *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.