VALIDATING A FREQUENTLY USED POKER-PLAYING HEURISTIC BY MEANS OF MONTE CARLO SIMULATION

Charles J. Teplitz, School of Business Administration, University of San Diego, 5998 Alcala Park, San Diego CA 92064, 619-260-4867, Teplitz@sandiego.edu
Steven F. Kling, MBA Candidate, School of Business Administration, University of San Diego, 1740 S. Westgate Ave, Unit H, Los Angeles CA 90025, 412-414-2596, Steven.Kling@gmail.com

ABSTRACT

When a decision must be made, the decision maker is often faced with a lack of time or a lack of data necessary to perform a proper problem analysis. In this situation, the decision maker might choose to use heuristics as a means of selecting between several decision alternatives. This paper will validate, through Monte Carlo simulation, a heuristic commonly employed in the playing of a popular poker game. The simulation will demonstrate that the heuristic tested does indeed provide acceptable approximations for predicting the likelihood of specific outcomes.

INTRODUCTION

The poker game of Texas Hold'em has become an extremely popular form of poker in recent years, especially after the invention of the *hole* camera used to expose the player's hidden hole cards to television viewers. Texas Hold'em is usually played by nine or ten players seated at a table with cards distributed by a casino dealer. All players in the game use five shared *community* cards in conjunction with their own two hidden hole cards to create their best possible five-card poker hand from the seven cards available.

Doyle Brunson [3, p. 419], the reputed godfather of poker, has explained that the game of Texas Hold'em "has more variety to it than any other form of Poker. And more complexity." Dan Harrington and Bill Robertie [8, pp. 12-13] explain the game's moniker of *The Cadillac of Poker* as being due to two factors: "the [limited] amount of information available to the players, and the ability to control the … odds offered to your opponent". They believe the attractiveness of the game is because the "two hidden cards allow plenty of room for deceptive maneuvering, while five exposed cards allow a good player to make plenty of deductions about the opposing hands".

While playing the game, players must frequently attempt to estimate the likelihood that the card needed to complete their hand will appear before the end of the game. On four separate occasions in a typical game of Texas Hold'em, each player remaining in the game needs to attempt to determine this probability. Performing the calculation is complicated by several factors: (i) the opponents' cards remain undisclosed throughout the game, (ii) there will be additional, as yet unknown, cards made available for use by all players, and (iii) the opponent is doing everything they can through their actions and words to convince the other players that they hold a superior hand. Not only is the poker player forced to operate within an environment of incomplete and misleading information, but in this setting it is unlikely that the decision maker will be provided sufficient time to perform the calculations necessary to determine the likelihood of having a winning hand.

Making the best decision when time and information are limited is a daunting task which often results in the decision maker settling for less than perfect guidance when attempting to increase the likelihood of a positive outcome. Simon [16] suggested calling this *satisficing* and recommended that decision makers find alternatives that are *good enough* rather than to try to find those that maximize their payoffs. Heuristics, or rules of thumb, fall into this category. It is generally understood that, by definition, the results derived through a heuristic approach can be inferior to those attained through an exhaustive algorithmic approach but are usually obtained quicker and at a lesser expense.

In this paper we will use Monte Carlo simulation to evaluate the accuracy of a common heuristic used to estimate the likelihood of obtaining a specific hand in a single game of poker. To empirically test the validity of the heuristic in question would require observing thousands of hands of poker and the systematic recording of the cards in play and the actions taken by the players. In lieu of undertaking this arduous task, simulating the many hands needed to properly test the heuristic requires far less time and is much easier to accomplish.

DECISION MAKING IN POKER

Much research has appeared in the literature over the years in an attempt to provide guidance to the decision maker in various types of poker games. Discussion has ranged from game theory strategies for simple two-person zero-sum games [4] [5] [7] [12] [14] [18] to strategies for playing in full table games [1] [2] [19]. In the literature reviewed, virtually all the researchers relied on the *Law of Large Numbers* to validate their estimation of the likelihood that a player's needed card will appear. Simply put, the Law of Large Numbers indicates that the probability of success converges to the expected value as the number of independent repetitions increases [17].

A Common Poker Heuristic

Obviously the ability to determine the likelihood of winning is a key factor in a player's overall success in the game. But since there is no way to accurately estimate this probability, how would a poker player determine such a value during the playing of a game? There are several options: (i) memorize the theoretical probabilities for various poker hands or draws to poker hands, (ii) do a complex probability analysis at the table, or (iii) use a simple heuristic for determining the probability. Option 1 is realistic but requires a diligence to memorize the odds for all possible scenarios. It also requires the ability to recall the memorized values as needed. Option 2 is not very realistic given the limited time available to the player in the midst of the game. Option 3 seems like the best option assuming the appropriate heuristics exist.

Frequently cited in reference texts designed to aid poker players is a two-step heuristic which supposedly can aid in estimating the likelihood of attaining a particular poker hand when on a draw to such a hand [6] [9] [10] [11] [13] [15]. The heuristic claims to accurately predict the likelihood of a player on a draw to a complete hand (straight, flush, etc.) after the first three community cards are dealt ultimately making the desired hand by the end of the game. The published procedure is:

Step 1. Subtract the number of winning cards previously seen from the total number of potential winning cards in a standard deck of cards. The difference represents the number of *outs* remaining for the player (i.e., the number of cards that could complete the player's hand).

Step 2. With two community cards yet to be dealt, multiply the number of outs as computed in Step 1 by four and report the result as a percentage.

VALIDATING THE HEURISTIC

Design

Monte Carlo simulation was chosen as the instrument for testing this heuristic. It can replicate the dealing of a ten-player game of Texas Hold'em for as many hands as needed to generate a statistically significant sample size in order to determine the long run expected probability of attaining a specific hand when having a post flop draw to that hand. The experiment performed will simulate a player's hand in a 10-handed game of Texas Hold'em after all players are dealt their two down cards and the three community cards have been exposed to all players. The focus of the simulation will be on those times when the subject player had four cards that represented a draw to a flush or a straight.

This simulation will be a Bernoulli experiment in which, over a series of independent trials (poker hands dealt), we will be able to determine the true underlying probability to which the observed results converge. The observed probability will be compared to the probability predicted by the heuristic in an effort to validate or discredit the heuristic in question.

To test the heuristic, three scenarios were considered: (i) the player is on a flush draw, (ii) the player is on an open-ended straight draw, and (iii) the player is on a gut-shot straight draw. In the flush draw scenario, the only relevant hands were those in which the subject player had two cards of one suit in their hole cards and there were exactly two cards of that suit in the first three community cards dealt. In the open-ended straight draw scenario, the only hands tallied were those in which the subject player had four consecutive cards between the player's two hole cards and two of the first three community cards dealt (e.g., 3,4,5,6) regardless of suit. In the gut-shot straight draw scenario, the only hands tallied were those in which the subject player had four cards between the two hole cards and the first three community cards dealt which needed a card whose rank fell within those four cards to complete a straight (e.g., 3,4,**?**,6,7,) again regardless of suit.

Procedure

The simulation for this study was created on a personal computer using Microsoft Excel® and its Visual Basic® tool. The simulation began with a list of 52 cards each depicted as a three digit number, with the first digit representing the suit of the card designated as a "1", "2", "3" or "4" and the last two digits representing the rank of the card, with "01" representing an ace, "02" through "10" representing the equivalent rank, with "11" representing a jack, "12" representing a queen, and "13" representing a king.

A simulation of five million hands of poker was performed to ensure that an adequate number of random drawing situations were sampled. For every simulated poker game, the following metrics were tracked for those runs in which the subject player was one card away from completing the hand they were drawing to post flop:

1. The number of times the subject player had a one-card draw post flop,
2. The number of times the desired hand was made with the turn card (community card number 4),
3. The number of times the desired hand was made with the river card (community card number 5),
4. The average number of actual outs existing post flop.

The sum of the second and third metrics will indicate the total number of times the desired hand was attained after the subject player needed one card to reach the desired hand post flop. The ratio of the sum of these metrics to the first metric would indicate the simulated probability of being successful in this situation. The fourth metric will be used simply to compare the simulated number of outs to the number of outs predicted by the heuristic. Ultimately, the probability of success determined through the simulation will be compared to (i) the probability derived through theoretical analysis to demonstrate the validity of the simulation, and (ii) the probabilities defined by the heuristic to demonstrate the accuracy of the heuristic.

Results

Of the five million runs of the simulation, the three scenarios being studied each occurred more than 100,000 times. An initial significant finding is the validation of the construction of the simulation and the sampling process used in the simulation as evidenced by comparing the simulated long-run probability to the value derived through a traditional theoretical approach. Table 1 shows that for all three scenarios, the simulated long-run probability and the long-run theoretical probability derived through traditional analysis were the identical.

Table 1 also compares the projections of the heuristic to the findings from the simulation. The two key values to be compared are: (i) the average number of post flop outs, and (ii) the probability of making the desired hand when the player is on a draw to that hand post flop. For each of the three scenarios, the relevant conditional cases (having a specific type of draw post flop) were reviewed to see how accurately the heuristics prescribed these values.

The first notable observation regarding these values is that in all scenarios the estimated number of outs derived by the first step of the heuristic was nearly double the average number of outs observed in the simulation. It should come as no surprise that this parameter is being dramatically overestimated by the heuristic. The heuristic naively assumes that <u>all</u> potentially beneficial cards were still available to the subject player after the flop (9 for a flush draw, 8 for an open-ended straight draw and 4 for a gut-shut straight draw). In fact, there are many possible reasons that some or all of these cards would actually be unavailable to the player: (i) there was one burn card prior to the flop and there will be two more before

TABLE 1: COMPARISON OF THEORETICAL, SIMULATED, AND HEURISTIC ESTIMATES

Post-Flop Scenario	Theoretical Value	Simulation (5,000,000 runs)	Heuristic (P=#outs x 4)
Flush Draw			
Average Number of Outs	5.36	5.36	9
P(Flush attained when on draw)	35%	35%	36%
Open Ended Straight Draw			
Average Number of Outs	4.77	4.77	8
P(Straight attained when on draw)	31%	31%	32%
Gut Shot Straight Draw			
Average Number of Outs	2.38	2.38	4
P(Straight attained when on draw)	16%	16%	16%

the hand is over; (ii) there are 9 other players each with two hole cards, so there are 18 other cards which could be the desired card; (iii) there will be 24 cards of the deck that will not come into play during the hand, any of which could be the desired card. It would seem that a properly formatted model (heuristic) would have as its dependent variable a number closer to the average number of possible outs as opposed to the maximum number of possible outs. Table 1 shows this to be the case for the simulated and the theoretical models.

Given that the heuristic being studied uses the number of outs as its sole independent variable, it would appear that the predictive power of the heuristic just isn't there. It would seem totally unreasonable to expect this heuristic to provide an accurate estimate of the likely chance of converting a drawing hand to a completed hand. However, the simulation demonstrates that this is not the case.

As seen in Table 1, the heuristic's predictions turned out to be surprisingly close to the theoretical probability, to within less than a single percentage point of the underlying probability in all scenarios. How is this possible given the inaccuracy of the heuristic's sole independent variable? It appears that the multiplier used in the second step of the heuristic corrects for the inaccurate value of the dependent variable. In fact, that's exactly how the multiplier was established. In determining the appropriate value for the multiplier in the heuristic, it was found that the heuristic always overestimated the number of outs by 67.8%, regardless of the type of draw. It was simply a matter of removing this bias that would result in this easy-to-use heuristic being capable of closely approximating the unknown probability. Setting the multiplier's value to four was enough to cause the heuristic to predict the desired probability within one percent of the theoretical probability without having to perform all the calculations necessary to obtain the true value.

CONCLUSIONS

The simulation demonstrated that the heuristic certainly provides the decision maker, the player, with a fairly accurate estimation in the face of the total lack of information available when a play/no play decision is necessary post flop while holding a drawing hand. For the poker player forced to make a rapid decision, the quickly derived percentages of success prescribed by the heuristic would likely be considered "close enough" to provide the player with a reasonable estimate of the needed information in order to make a fairly accurate assessment of the expected value of continuing play.

On the other hand, regardless of the method used to determine the probability of drawing a needed card, it should be remembered that derived probabilities are long-run averages based on a multitude of future poker hands. Unfortunately, the poker player is attempting to determine the likelihood of making the desired hand on this singular occurrence. Computing this probability is simply not possible without the aid of a crystal ball, tea leaves, or other such devices. This impossibility does not preclude a poker player from occasionally needing to know such a number. This paper provides evidence to the poker player that, in the face of this impossibility, this well-tested heuristic is at least capable of quickly and accurately deriving the long-run average likelihood of filling the draw. This knowledge will help the player make a proper play/no play decision more times than not thereby ensuring a positive experience, and bankroll.

REFERENCES

[1] Barr, G. D. I., & Dowie, K. Swiveling Probabilities - The Winning Hand at Poker. *The American Statistician*, 2011, 65 (3), 170–176.

[2] Bina, B., Chen, H.W., & Milgram, P. A Model of Expert Decision Making in Post-flop Betting in Poker. *Proceedings of the Human Factors and Ergonomics Society*, 2008, 433-436.

[3] Brunson, D. *Doyle Brunson's Super System*. New York, NY: Cardoza Publishing, 1978.

[4] Ferguson, C., & Ferguson, T. On the Borel and von Neumann Poker Models. *Game Theory and Applications*, 2003, 9, 17-32.

[5] Ferguson, C., Ferguson, T., & Gawargy, C. Uniform (0, 1) Two-person Poker Models. *Game Theory and Applications*, 2007, 12, 17-37.

[6] Flynn, M., Mehta, S., & Miller, E. *Professional No-limit Hold'em*. Las Vegas, NV: Two Plus Two Publishing, 2007.

[7] Friedman, L. Optimal Bluffing Strategies in Poker. *Management Science*, 1971, 17, B764-B771.

[8] Harrington, D., & Robertie, B. *Harrington on Hold'em*, (Vol. 1). Las Vegas, NV: Two Plus Two Publishing, 2004.

[9] Hilger, M. *Internet Texas Hold'em*. Atlanta, GA: Dimat Enterprises, 2003.

[10] Jones, L. *Winning Low Limit Hold'em*. Pittsburgh, PA: ConJelCo, 2000.

[11] Moshman, C. *Sit 'n Go Strategy*. Las Vegas, NV: Two Plus Two Publishing, 2007.

[12] Newman, D. *A Model for 'Real' Poker*. Operations Research, 1959, 7, 557-560.

[13] Rodman, B., & Nelson, L. *Kill Phil*. Las Vegas, NV: Huntington Press, 2005.

[14] Seale, D. A., & Phelan, S. E. Bluffing and Betting Behavior in a Simplified Poker Game." *Journal of Behavioral Decision Making*, 2010, 23, 335-352.

[15] Shulman, B. *52 tips for Texas Hold'em Poker*. Las Vegas, NV: Card Player Press, 2005.

[16] Simon, H. A. A Behavioral Model of Rational Choice. *Quarterly Journal of Economics*, 1955, 69, 99-118.

[17] Tijms, Henk. *Understanding Probability* (3rd ed.). Cambridge, UK: Cambridge University Press, 2012.

[18] Von Neumann, J., & Morgenstern, O. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 1944.

[19] Zadeh, N. (1977). Computation of Optimal Poker Strategies. *Operations Research*, 1977, 24, 541-562.