A Simulation of Bayesian Impact on Heads Up Poker

Trevor Squires

Last Updated: July 30, 2023

1 Introduction

Poker has traditionally been a game about human psychology. In the early 2000s, popular players such as Phil Hellmuth would be adorned for their incredible reads of their opponents. There have been many books written on the "tells" of an opponent whether it be the tone of their voice, the stillness of their eyes, or the way they handle their chips. However, as the game has traversed well into the 21st century, mathematics has been helping players identify another more obvious tell. This new tell would have to be completely independent of human behavior as the rise of online tournaments has more or less removed that aspect of the game. Players would have to rely on the only information they have at their disposal - the actions of their opponents.

If we assume that each poker player is interested in winning money, then we can assume that each player is trying to select their optimal action given the information available to them. By observing their actions and evaluating the information they used to make said actions, we can backsolve for the private information of each player. That is, the way that each player bets is itself a tell of their hand.

Bayesian statistics attempts to quantify how we can use our opponents' tells effectively. Here in this document, we will explore exactly that: how Bayesian statistics can help in the context of heads-up Texas Hold'em poker (THP).

2 Rules and Assumptions

As mentioned above, we will specifically be analyzing the game of heads-up Texas Hold'em poker. THP is a style of gambling in which n players are dealt 2 hole cards face-down from a standard deck of 52. These are private to each individual player. There is an initial buy-in called the big blind from one of the players. The game begins with pre-flop action in which each player is given the opportunity to bet into the pot, but must at least match the maximum bet from all the other players before proceeding. If a player does not wish to continue in the hand, they can fold and relinquish all bets made into the pot but do not have to bet anymore for this hand. If, at any point, there is only 1 remaining player in the pot, they are declared the hand winner.

After each player has matched the maximum bet, we move into the post-flop street which is usually just called the flop. At this point, three community cards are dealt face-up. These community cards combine with each player's hole cards to form a 5-card hand. Again, players may bet to try to claim the pot for themselves or fold and not continue betting any money. After the betting stage is over and there are more than 2 players remaining in the hand, we move to the turn in which another community card is dealt. Another round of betting ensues before moving to the final street, the river. A fifth and final community card is dealt followed by a final round of betting. If there are more than two players remaining in the pot after the river betting, then we proceed to a showdown in which the best 5-card hand of each player is evaluated against each other. The best 5-card hand is deemed the winner of the hand. The hand winner receives all the money in the bet into the pot.

Hand evaluation for THP is fairly straightforward logic, with the strength of a hand being closely associated with the probability that the hand happens at random. For example, 2-of-a-kind is more common than 3-of-a-kind and is therefore weaker. An ordering of all hand types is shown in Figure 2. For simplicity of our analysis, we assume that there are only two players (heads-up), there is no small blind, and ties are broken randomly.

2.1 A Simple Strategy: Pot odds

Given the rules of poker explained above, how should we as a player decide to take our action (bet)? It is clear that if we think that our hand has good showdown value, i.e. is a strong hand, then we should bet our own money into the pot. This either makes the other player also bet their chips which would increase the total amount of

Figure 1: Texas Hold'em Hand Ordering

money we can win if our hand is the best, or forces them to fold in which we win all of the money currently in the pot. On the flip side, if we think our hand is weak, we should try to "defend" our chips already in the pot while also trying not to add too much more to it.

A very simple (and by no means optimal) strategy can be developed off of this mindset. Suppose that we knew our probability of winning p, usually called equity, and it is our turn to act. The pot is currently worth M chips, we are faced with a bet b from our opponent. Should we make the call or not? If we call b, then our expected reward is our equity times the total amount in the pot less the amount of the call i.e. $r = (M + b)p - b$. If we fold, our expected reward is obviously 0. Thus, we call if and only if

$$\frac{b}{M+b} < p.$$

The left-hand side of the above equation is typically referred to as our pot odds. The right-hand side is the aforementioned equity, win probability, or also card odds. That is, we call if our card odds exceed the pot odds.

Now suppose that we are the first player to act. The pot is M, but now we must decide how much to bet. We can bet 0 (check) in this position, or if we like our hand, we can force our opponent into a call decision with a raise. This decision is tremendously more nuanced than the decision to call or not. In fact, this is arguably the fundamental question of poker: how much should a player bet? We won't attempt to get into game theory optimal strategies here, so we propose a simple one. When allowed to bet, we can simply bet the amount that we were willing to call with if we had been faced with a bet from our opponent.

This strategy is far from perfect. In fact, it frequently has the player raising bets unnecessarily. But it does maintain the natural property of "if our hand is strong, we raise the pot". To help stem the negative side-effects of this strategy, we will also add in a few bits of randomness and pot control. Our full decision-making process works as follows:

1. Compute our desired bet $b = \frac{Mp}{1-p}$
2. If b is less than what we must bet to call, we fold.
3. If b is less than twice the amount to call, we call rather than raise
4. If b is greater than twice the amount to call, we call with probability $\frac{1}{4}$ to disguise our hand. Otherwise, we bet b

Again, this strategy is not perfect, but it is somewhat intuitive and satisfies many basic principles. We will use it as a starting point for our analysis and a platform for which we can employ Bayesian techniques.

3 The Bayesian Update

Our strategy mentioned above is only as good as the inputs used to perform it. In particular, we made a very massive assumption that we knew the win probability of our hand. Since this is very often massively dependent on hidden information, our opponent's cards, this is not something that we know in practice. However, we can form an estimate of our win percentage if we were to estimate the distribution of hands that our opponent may have. This distribution of our opponent's hands - their range - can then be used to determine how often we are to win against them.

When the hole cards are dealt and before there is any action, we know after removing our hole cards from the deck, there are ${}_{50}C_2$ remaining hands that our opponent could have which are all equally likely. By the same logic, there are then ${}_{48}C_5$ possible runouts of the community board. Only after all these cards have been dealt will we know the winner. Thus, in theory, all it takes to compute our expected win percentage is to play out all such boards against all such opponent's hands and sum up the number of times we win. Dividing out by the number of total possible runouts we evaluated would give us our win percentage. Although this process is very computationally expensive, the point here is that none of these estimates rely on any hidden information. A highly intelligent supercomputer could, in theory, employ this exact process of computing win probabilities and use it to their advantage.

To speed things up computationally, however, we will use another simulation to estimate our win probabilities. Instead of simulating all possible runouts of the community board, we can use a Monte-Carlo simulation to simulate a much smaller number of runouts for computational feasibility. This certainly adds variance into the mix, but for our testing purposes, it will affect both players equally.

3.1 Bayesian Theorem

In the previous section, we discussed a simple strategy for playing THP and a way to theoretically gather all the information necessary to implement the strategy. However, we did make a simplifying assumption: our opponent's hands are equally distributed. When the game first begins, we have no information to say otherwise. After a few betting rounds, however, we do have our opponent's tells - their bets - to update that distribution. Mentally updating our opponent's range of hands will improve our estimate of our win probability. As we argued previously, our win probability is crucial in effectively employing our poker strategy. A more accurate win probability will enable us to make better decisions. Exactly how we should use our opponent's actions to update our view of their range brings us Bayesian statistics.

To study this topic, let us stray away from the poker table and speak more abstractly. Suppose that A and B are events, things that can happen. We can use events to talk about circumstances in the future that haven't materialized yet. For example, we could let A be the event a particular stoplight turns red within 5 seconds. Or B could be the event that more than four babies are born in the next twelve minutes. These events are associated with a bit of randomness. We can't say whether they are true or false, but we can ask questions about the probability of the event happening. Given how large the world's population is, we can estimate that $P(B)$ is quite high.

Events can also be chained together to form more interesting ones. The event $A \cap B$ is the event that both A and B are true. A less trivial type of event is one in the conditional sense. We can write $A \mid B$ to mean the event that A is true, given that we know B is true. In this particular example, we probably don't gain any insight into whether or not a stoplight turns red simply by knowing if a sufficient amount of babies are born, but these conditional events and their respective probabilities can help us more accurately describe the world around us.

Take for example a vending machine in an abandoned warehouse that we will use twice. Let A represent the probability that the vending machine works properly on the first attempt and B the same for the second attempt. Given that this vending machine is very worn out and doesn't appear to have power, we would estimate the probabilities $P(A)$ and $P(B)$ to be very low. However, suppose after inserting the appropriate funds, we find that the vending machine works perfectly on our first attempt. This gives us great optimism about the second attempt, much more so than we did previously. This intuitive line of thinking is captured perfectly by conditional probability. Before knowing any information about A, we would have said that $P(B)$ is not very large. After learning the materialization of A, however, we have a change of heart about this abandoned vending machine. That is, the event $B \mid A$ is much more likely than just B. The event $B \mid A$ represents our updated view of the environment after an event has occurred.

Not only is this natural behavior modeled by conditional probability, but it can also be quantified by it too. When $P(B) \neq 0$, we can view the conditional probability of A given B as the ratio of the probability of both events being true to the probability of only B being true, i.e.

$$P(A \mid B) = \frac{P(A \cap B)}{P(B)}.$$

That is, $P(A \mid B)$ is simply $P(A \cap B)$ when we restrict our attention to only looking where B is true. This relation is symmetric in A and B and the two can be interchanged. Solving for $P(A \cap B)$ and substituting into the same equation for $P(B \mid A)$, we obtain the famous Bayes Theorem

$$P(B \mid A) = \frac{P(A \mid B)P(B)}{P(A)}.$$

This equation may seem pointless at first, but it can be leveraged to compute an updated view of our surroundings simply by observing the environment.

3.2 Application to Texas Hold'em

We now return to our poker table where we are again faced with the task of updating our opponent's range. Suppose that we've just been dealt two 8s and chose to raise the big blind to put pressure on our opponent. Our opponent takes a look at their cards and immediately goes all-in and bets all of their chips. A pair of 8s is a fairly strong hand in heads-up poker, especially pre-flop, but our opponent has apparently decided that their hand is strong too. In a vacuum, a pair of 8s has 50% equity and we would have higher equity than pot odds to warrant a call, but does our opponent's overly large sizing give us any additional information?

Let's analyze with a particular opposing hand in mind, 2-7 offsuit. This is statistically the worst pre-flop hand in poker. To begin, let B denote the probability that our opponent has this hand. Before any chips were spent, we assumed that our opponent had this hand equally as likely as any other: $P(B) = 1/{}_{50}C_2$. What we now care about is the probability that our opponent has this hand after seeing their large bet size. Let A denote the probability that our opponent jams (goes all-in) after seeing our large bet. We are interested in $P(B \mid A)$.

Using Bayes theorem, we know that

$$P(B \mid A) = \frac{P(A \mid B)}{P(A)_{50}C_2}.$$

We can use our assumptions around $P(A \mid B)$ to update our working knowledge of $P(B \mid A)$. Examining the former, the likelihood of our opponent jamming with the worst hand in the game is probably very small for any reasonable player. Sure, it can happen if they are bluffing, but most of the time this is a very non-profitable strategy. In fact, if we assume our opponent is employing the aforementioned strategy, we have very good reason to believe that this is not a bluff.

All that remains is estimating $P(A)$. However, note that $P(A)$ is entirely independent of B. That is, this value does not change based on which hand of our opponent we are trying to get a better update for. Thus, since we know that the probability that our opponent holds any of the ${}_{50}C_2$ remaining hands must sum to 1, we can ignore this constant factor and normalize the distribution of hands afterward. This leads us to an incredibly simple strategy for updating our opponent's range.

Let H_i be the event that our opponent has hand indexed by i. We initialize $P(H_i) = 1/{}_{50}C_2$. After seeing a bet b, we loop over all i and update via $P(H_i) = P(H_i) \cdot P(B \mid H_i)$ where B is the event that our opponent bets b. Finally, we normalize the distribution by dividing $P(H_i)$ by $\sum_i P(H_i)$ for all i. In this way, we are simply raising the probability of each hand relative to how likely it was that the opponent made the bet that they did given a particular hand. Applying this strategy to the above example, we can be almost certain that our opponent does not have the worst starting hand in THP.

With this Bayesian update strategy in hand, we now have a reasonable way of updating our opponent's hand throughout the game.

4 Computational Results

In this section, we will put all the pieces of our simulation together in order to test the effectiveness of our Bayesian update strategy.

We begin by initializing our game with two players. They begin with equal starting chips and we fix a big blind amount. Both players also have the same betting strategy that we described previously. The only difference between the two players is that one of them, which we'll call the Bayesian, updates their opponent's range estimate after each bet. The other player, the Frequentist, only uses the initialized uniform distribution. We play hands until one player does not have the necessary chips to bet the big blind and the other player is determined the winner.

In order to update the opponent's range, recall that we did need an estimate of the probability that the opponent would make the observed bet given a certain hand. To avoid simulating the unlikely scenario that we know exactly how our opponent plays their hand, we implement the following strategy:

1. Fix an opponent hand i
2. Compute the win percentage for i assuming a uniform distribution of opposing hands
3. Compute the ratio of the observed bet b to the perceived desired bet b' using the same $b' = \frac{Mp}{1-p}$ equation
4. If the ratio was greater than 2, it was unlikely the player had this hand and we only assume bluffs would make this play. Return 1/10
5. If the ratio was smaller than 1/2, it is unlikely the player would bet this passively, but they could be disguising a strong hand. Return 3/10
6. Otherwise, return 6/10

It is important that the assumption of how the opponent plays is not too similar to their actual strategy. If the Bayesian principle is successful even when the strategies do not perfectly line up, we are able to argue that the Bayesian update is fairly robust.

4.1 Bayesian Impact

After simulating many rounds of THP between the two players, we can very easily see a difference in the results. With a starting stack of 2000 and a big blind of 10, the Bayesian player doubled up on the Frequentist in roughly 75% of all simulated runs. Somewhat unsurprisingly, the most impactful hands ended up being where the Frequentist had an above-average win percentage, but the Bayesian was stronger. On the flip side, the Bayesian was able to get out of above-average hands more often than the Frequentist after observing large bet sizes.

Perhaps more interesting is the effect that some of the tuning parameters had on the results. When the big blind was increased, the gap between the two players was narrowed. This should make intuitive sense - a larger big blind implies that more of the bets in each hand are forced and fewer are decided upon by the players themselves. When we take away the ability to influence the outcome through decision-making, the winner becomes more of a coin flip than anything else.

Similarly, the further that the perceived strategy and the actual strategy of the Frequentist grew from each other, the less of an advantage the Bayesian player had. In these simulations, however, the Bayesian never relinquished an overall victory, just a narrower margin of victory.

5 Conclusion

Overall, this experiment was an interesting one that I felt was well worth the time given my admittedly lack of experience in all things statistics. In many aspects, this was a rather significant undertaking to showcase an idea that could have been displayed in a drastically simpler setting. However, I am also an avid fan of poker and so this was well worth the time invested.

As far as the results go, this simulation showed rather straightforwardly that ignoring the opponent's tells of their bets puts one at a significant disadvantage as well as how to leverage Bayes theorem to do so. Although the base strategy is far from game theory optimal, it's not hard to see how much of an impact using all available information can be. In fact, this goes to show that even in an environment absolutely devoid of human interaction, we can still extract tells from mathematical analysis alone.