Exploitability and Game Theory Optimal Play in Poker

Jen (Jingyu) Li[1,a]

Abstract. When first learning to play poker, players are told to avoid betting outside the range of half pot to full pot, to consider the pot odds, implied odds, fold equity from bluffing, and the key concept of balance. Any play outside of what is seen as standard can quickly give away a novice player. But where did these standards come from and what happens when a player strays from standard play? This paper will explore the key considerations of making game theory optimal (GTO) plays in heads-up (two player) no limit Texas hold'em. To those new to the game, it involves dealing two cards that are revealed only to the player they are dealt to (hole cards), and five community cards that are revealed with rounds of betting in between. Hands are compared by looking at the highest five card poker hand that can be made with a player's hole cards combined with the community cards. This paper will focus on exploitative strategies and game theory optimal play in heads-up poker based on examples of game scenarios from [1].

Keywords: Discrete math, probability, poker theory, game theory.

Mathematics Subject Classification: Mathematics Subject Classification According to AMS..

Recibido: de 0 Aceptado: de 0

1. Introduction

Poker is a game that has been extensively studied from a mathematical standpoint, as it is interesting from a game theory standpoint and highlights considerations that must be made when making decisions under uncertainty and

[1]Department or Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, US
[a]jingyuli@mit.edu

deals with expected value of strategies over time. It is a game with strategies that are not immediately intuitive and the value of those strategies are only seen over a large number of hands. To reduce complexity, this paper will focus on heads-up (two player) poker. To those new to the game, the game begins with each player being dealt two cards which are hidden from the other player. A round of betting takes place, where there are four actions available to the players: check, bet, call, raise. A player can check or bet if no amount has yet been made in the current round of betting and a player can call (match the amount bet by the opponent) or raise (bet an additional amount on top of opponent's bet) if the opponent bets. After the initial round of betting (pre-flop), the first three community cards (visible to both players) come out (flop). Another round of betting proceeds before the fourth card comes out and likewise before the fifth and final card. After all cards are out, there is one last round of betting before the players' hands are compared (showdown). The complexity of poker arises from inferring probabilities through the many rounds of betting and making decisions that consider events in the future. To understand the mathematics behind playing optimally, we dissect the game into constrained sub-problems, but the concepts derived through these examples are relevant in real play.

2. Pot Odds

Definition 2.1. We refer to a ***made hand*** as a poker hand that is already guaranteed given a player's hole cards and currently revealed community cards.

Definition 2.2. We refer to a ***draw*** as a hand that can be made given certain community cards come out.

Example 2.1.
Suppose Alice has $A\diamondsuit A\spadesuit$ and Bob has $5\heartsuit 6\heartsuit$. The community cards on the ***turn*** (stage of game where 4 community cards have been revealed) are $K\diamondsuit 9\heartsuit 2\clubsuit Q\heartsuit$. Alice has a made hand of a pair of aces and Alice has a draw to a straight. Now if both players knew each other's cards, they would agree that if the last card is a 3 or 8 of any suit, Bob wins, otherwise Alice wins. In this world of perfect information, neither Alice nor Bob would bet on the river (when the last card comes out), because the winner would be clear.

Now suppose there is already \$100 in the pot and Alice can either check or bet before the river card comes out. If Alice bets, Bob has the option to re-raise. There are 9 hearts remaining in the deck, which would give Bob a flush, beating Alice. The remaining 35 cards would allow Bob's aces to hold. Suppose Alice is to act first. Since Alice is favored to win the hand, Alice has reason to bet here. The amount she should bet is derived from calculating expected value (EV).

The expected value is calculated as the probability of Alice winning the pot times the new pot amount, deducted by the amount she bets. Note that this calculation emphasizes that as soon as Alice places a bet, she should no longer consider that money to be her's to lose, but rather part of the pot that she can win (sunk cost).

$$\begin{aligned}\mathbb{E}(A) &= \frac{35}{44}(100+2x) - x \\ &\approx 80 + 0.6x\end{aligned}$$

Note that if the probability of winning here is less than $\frac{1}{2}$, it is not profitable to bet. This however is complicated when we consider a real game where both players do not have complete information and bluffing is a valid strategy.

Also note that Alice's EV is strictly increasing as her bet increases *if* Bob always calls. Bob however, should only call if it is positive EV for him.

$$\begin{aligned}\mathbb{E}(B) &= \frac{9}{44}(100+2x) - x \\ &\approx 20 - 0.6x\end{aligned}$$

Bob should thus only call if Alice's bet is below around \$33 or $\frac{1}{3}$ of the pot pre-betting. This $\frac{1}{3}$ is what we refer to as pot odds. It is important to keep in mind that Bob can call larger bets or even re-raise because of something we refer to as ***implied odds***, which take into consideration further betting on the river due to it being unknown who has the better hand.

3. Implied Odds

Implied odds refer to the potential to make more money when a draw hits. Remember that we previously assumed both players had complete information. This is not true in a real game, which means betting on the river can be profitable. In the case of our previous example, Alice does not know what Bob has, so if Bob hits his flush, he can potentially make more off Alice than was estimated by our EV calculations on the turn.

Example 3.1.
Let us continue with our previous example. If Bob hits a flush on the river, we will assume that he knows correctly that he has the better hand (for now we will ignore the possibility Alice has a higher flush, because the probability is relatively low). Suppose Alice bet $50 on the turn and Bob called. The final card comes 7♥. Now the pot is $200 and Alice acts first. Recall that the board currently shows K♦9♥2♣Q♥7♥. Now Alice doesn't know what Bob has and believes it's likely he has top pair (a king that pairs with the king showing on the board). Alice thinks she can bet again to get value off of Bob. Here Bob can fairly safely call or re-raise Alice's bet.

Let's look at what Alice should do when the river card comes out. Suppose she's fairly certain Bob either hit his flush or just has the top pair on the board. Estimating the probabilities of these two cases is more complicated (has to take into account what kinds of hands Bob generally plays and the history of actions on the current hand), but it's fair to assume Bob has more hands involving kings in his range than two hearts.

This means that if Bob knows Alice will bet on the river even if he hits his flush, he is willing to call larger bets from Alice on the turn or even re-raise or bet if Alice checks.

Definition 3.2. We refer to a player's range as the hands he plays in a given situation. In general, a player's range does not change from hand to hand. That is not to say that the player should be predictable (see Section 4.2 regarding balancing range).

4. Game Theory Optimal Strategies

4.1. Exploiting the Opponent

In actuality, the size of bets should not be proportional to how good your hand is, nor should you only bet when you have a good hand, as that is exploitable by the opponent over time. In the previous sections, we looked at examples constrained to a single hand, in which case we only care about maximizing EV on that hand. However, poker is all about beating the odds over time, so it's important to realize that a strategy optimized for a single hand may not be optimal or even profitable in the long run.

As a simple but realistic example, suppose your opponent only bets and raises hands that they think will win the pot, but still calls some of your bets with weaker hands (this is not an uncommon type of play from risk-adverse beginners). It's easy to exploit a player like this by simply using a strategy which folds to every bet or raise the opponent makes and still betting our good hands. Of course, eventually the opponent will catch on and counter-exploit by bluffing their hands if they know we will fold. On this end of the spectrum, suppose a player bluffs too many hands. To exploit this play style, we can afford to play a larger portion of hands and make large profits when we hit a top hand.

This leads us to the idea of balancing our range, or deciding the hands we play in a given situation such that an opponent cannot exploit our strategy.

4.2. Balance

To play non-exploitable game theory optimal (GTO) poker, ranges should be "balanced", meaning Often this means that we have a variety of possible hands in the eyes of the opponent in any situation. This means adding in a range of hands with which you bluff and not betting only when you have a good hand or betting a larger amount when you have the winning hand.

Definition 4.1. We define ***defensive value*** as the expected value of a strategy against the opponent's most exploitative strategy. Note the difference between this value and EV as we've previously looked at is that this assumes the opponent knows how we play and can exploit any patterns over time in our strategy.

A more rigorous definition of balanced strategy is minimizing the gap between defensive value (Definition 4.1) and expected value. In other words, the expected payoff of the strategy in a given hand should not change over time as your strategy is gradually exposed to your opponent: your opponent plays the same way regardless whether your strategy is known to them.

Definition 4.2. A ***pure strategy*** dictates a player's action in any situation i.e. the player will always make the same decision under given circumstances.

Definition 4.3. A ***mixed strategy*** is one in which the player assigns a probability distribution over all pure strategies (Definition 4.2).

Definition 4.4. ***Nash equilibrium*** is a strategy set in a multi-player game where neither player alone can increase their payoff. Because of this, it is a stable point where neither player wants to deviate from their current strategy.

Definition 4.5. A game in which the sum of all players' scores is equal to 0 is called a ***zero-sum game***.

Definition 4.6. ***Indifference*** refers to a game state where a player gets the same expected payoff regardless what strategy is chosen.

Definition 4.7. An ***indifference threshold*** is a value for a parameter that a player can choose to force *indifference* (Definition 4.6) on the opponent.

It is a known fact of game theory that all multi-player games with finite payout matrices have at least one *Nash equilibrium* (Definition 4.4). Additionally, poker is a *zero-sum game* (Definition 4.5) and it is known that all zero-sum two-player games have an optimal strategy as long as *mixed strategies* (Definition 4.3) are allowed. This leads to the concept of ***indifference*** (Definition 4.6. By setting expected payoff equations equal to each other, we can obtain values for parameters that force a player to be indifferent to choosing among strategies. The value of the parameter found by solving these equations is an ***indifference threshold*** (Definition 4.7). Let us take a look at the following example.

Example 4.8.
Suppose Bob has three of a kind and on a particular board is only scared of Alice having a flush. Let us assume that Alice has a flush here 20% of the time. How often can Alice bluff? For this example suppose there is \$300 in the pot and Alice can choose to bet a fixed amount of \$100. To keep it simple, we will say Bob either calls or folds when Alice bets.

How often should Alice bluff here? If Alice bets \$100, Bob can pay \$100 to potentially win \$400. Suppose Alice only bets when she has the flush. Bob can exploit this strategy by folding every time Bob bets, preventing him from getting any additional value from hitting his flush and taking the pot 80% of the time. Alice has a defensive value of $0.2 \cdot \$300 = \$60$ with this strategy, where she only profits when she has the flush. Now suppose Alice bets all her hands here. 20% of the time she has the flush and the other 80% of the time she has nothing. If Bob calls, his EV is $0.8 \cdot \$400 - \$100 = \$220$ and if he folds, his EV is 0, so Bob will exploit Alice's strategy here by always calling. The defensive value of Alice's strategy is $0.2 \cdot \$400 - \$100 = -\$20$.

The two strategies mentioned so far (always checking a dead hand and always betting a dead hand) are what are known as *pure strategies* (Definition 4.2) and neither is optimal for Alice in this situation. We know this, because both are exploitable – Bob alone can change his strategy and increase his payoff. This indicates we are not at a equilibrium point.

Now, we explore mixed strategies. Let $P\langle A, \text{bluff}\rangle$ be faction of all hands Alice has on the river that she bluffs with. Bob's EV for calling when Alice bets can be computed as

$$\mathbb{E}_B\langle B, \text{call}\rangle = \frac{P\langle A, \text{bluff}\rangle}{0.2 + P\langle A, \text{bluff}\rangle} \cdot \$400 - \$100$$

Alternatively, Bob can fold when Alice bets.

$$\mathbb{E}_B\langle B, \text{fold}\rangle = \$0$$

Alice's EV can be computed as

$$\mathbb{E}_A\langle B, \text{call}\rangle = \frac{0.2}{0.2 + P\langle A, \text{bluff}\rangle} \cdot \$400 - \$100$$
$$\mathbb{E}_A\langle B, \text{fold}\rangle = (0.2 + P\langle A, \text{bluff}\rangle) \cdot \$300$$

Alice's strategy is least exploitable when Bob's EV for calling and folding are equal (i.e. Bob is not able to change his strategy to exploit Alice even if over time he figures out how often Alice bluffs). By setting $\mathbb{E}_B\langle B, \text{call}\rangle = \mathbb{E}_B\langle B, \text{fold}\rangle$, we can solve for Alice's optimal bluff frequency such that Bob is indifferent to calling versus folding. It turns out that it is optimal for Alice to bluff around 6.7% of her hands.

Example 4.9.
Consider the general scenario where we only have one round of betting, the pot has B bets, Alice can make a bet of size 1, and Bob can call or fold if Alice bets. The payout matrix is as follows

	Bob	Check-call	Check-fold
Alice			
Winning hand	Bet	P + 1	P
	Check	P	P
Dead hand	Bet	-1	P
	Check	0	0

As we can see from the payout matrix, it is always in Alice's favor to bet when she has a winning hand. It is less obvious what Alice should do when she has a dead hand. Depending on Bob's calling versus folding frequency, it can be beneficial for Alice to bluff a percentage of her dead hands.

According to the concept of indifference, Alice wants to choose a bluffing frequency such that Bob's EV for calling is equal to his EV for folding. Let b represent $\frac{\text{bluffs}}{\text{bluffs+value bets}}$.

$$\mathbb{E}_B\langle \text{call} \rangle = b(P+1) - 1$$
$$\mathbb{E}_B\langle \text{fold} \rangle = 0$$

We have $\mathbb{E}_B\langle \text{call} = \mathbb{E}_B\rangle$fold when

$$b = \frac{1}{P+1}$$

Likewise, Bob should choose a calling frequency such that Alice is indifferent to checking versus bluffing her dead hands. Let c be the frequency with which Bob calls.

$$\mathbb{E}_A\langle \text{check} \rangle = 0$$
$$\mathbb{E}_A\langle \text{bluff} \rangle = (1-c)(P+1) - 1$$

By setting these two EVs equal to each other, we find the value c with which Bob should call when Alice bets.

$$c = \frac{P}{P+1}$$

It turns out these two quantities are quite useful, so we will give the quantity $\frac{1}{P+1}$ its own letter, α. Alice's optimal bluff to bet ratio is equal to α and Bob's optimal calling frequency is equal to $1-\alpha$.

This can be generalized to different bet sizes. Bets are generally thought about as a fraction of the pot (according to pot odds). Suppose Alice can bet any fraction of the pot xP.

$$b = \frac{xP}{P + xP}$$

$$= \frac{x}{1 + x}$$

5. Multi-street Games

Thus far we have mainly discussed single street (one round of action) scenarios, but in reality, action on a given street depends on everything that has happened before. In Example 4.8, we assumed that Alice has a flush 20% of the time. In reality, this probability depends on everything that happened before the river.

Example 5.1.
Let's set up the following scenario:

- The board shows $K\diamondsuit 9\heartsuit 2\clubsuit Q\heartsuit$.
- Alice has a pair of aces.
- Bob has a hand from a distribution which contains $\frac{1}{10}$ hands with two hearts and $\frac{9}{10}$ dead hands.
- The pot contains \$4, and players can either check or bet \$1.
- Alice is first to act.
- Bob is confident he has the winning hand if he hits his flush and a dead hand otherwise.

The flush comes around 20% of the time (in actuality, it's a little less but for simplicity's sake we will use 20%). From what we studied before, Alice should bet and Bob should call if he has the odds. Note that implied odds should be considered here rather than just pot odds, because Bob can get more value on the river by hitting his flush.

Assume Alice and Bob make it to the river, and now there is \$6 in the pot (Alice bets on the turn and Bob calls). Now Alice has no reason to bet here, because we have assumed Bob knows whether he has the winning hand at this point. Thus Alice checks and Bob can choose to either bet or check. According to Example 4.9, Bob should bluff with $\alpha = \frac{1}{7}$ as many hands as she value bets with and Alice should call with a frequency of $1 - \alpha = \frac{6}{7}$.

However, this is actually incorrect, because our prior calculations relied on a single street game. Consider how this situation is different. Suppose Bob wants to bluff on the river. This means he had to have called Alice's bet on the turn with a dead hand. Also note that Bob can only bluff on the river if a heart comes out. Thus in this multi-street game, Alice should be considering indifference of Bob folding versus playing a dead hand through both streets. Let c be Alice's optimal calling frequency on the river.

$$\begin{aligned}
\mathbb{E}_B\langle \text{dead hand}, fold\rangle &= 0 \\
\mathbb{E}_B\langle \text{dead hand}, play\rangle &= P(\text{flush})P(\text{Alice calls})(-2) \\
&\quad + P(\text{flush})P(\text{Alice folds})(5) \\
&\quad + P(\text{no flush})(-1) \\
&= (0.2)(-2)c + (0.2)(5)(1-c) + (0.8)(-1) \\
c &= \frac{1}{7}
\end{aligned}$$

So we see that it turns out Alice's optimal calling frequency is actually $\frac{1}{7}$ rather than $\frac{6}{7}$ from analysis of a single street game. By analyzing a single street game, we are able to reason about strategies, but the determined frequencies cannot be blindly applied to multi-street games where there are added layers of complexity.

6. Conclusion

GTO strategy explores the concepts of balance and indifference which minimizes exploitability. When you have minimal knowledge of the opponent's play style, it is a good defensive strategy to play close to GTO, which aims to optimize for the worst case by minimizing your own exploitability. GTO strategy assumes an opponent who also plays optimally, or knows how to exploit weaknesses in any strategy.

However, the assumption that the opponent is always perfectly rational and plays according to GTO strategy is rarely true and the discrepancy is what allows players who know how to take advantage profit. Even good players do not play a perfectly balanced game and open themselves up to exploitability in which case it is beneficial to play to their weaknesses whenever you have the information to do so.

To play exploitative poker, it is important to consider past information about the opponent's overall play style, ranges, and strategy as well as what actions took place on earlier streets of a given hand. However, keep in mind that strategies that stray too far from GTO can be counter-exploited. Suppose we have two perfectly rational players who start off with very different exploitable strategies. In theory, over time the two players would learn to exploit and counter-exploit each others' strategies and eventually their strategies would converge to near GTO.

In conclusion, depending on the opponent, playing GTO may not always be the most profitable, but it minimizes exploitability, making it a safe strategy to play against any opponent.

References

[1] Bill Chen and Jerrod Ankenman, *The mathematics of poker*, ConJelCo, 2006.