Fictitious Play applied on a simplified poker game*

Ioannis Papadopoulos

June 26, 2015

Abstract

This paper investigates the application of fictitious play on a simplified 2-player poker game with the goal of approximating optimal strategies. It is shown that fictitious play indeed produces solutions very close to optimal. The calculated strategies indicate that the player who acts second typically has an advantage.

1 Introduction

Poker is a zero-sum incomplete information game. The game starts with each player being dealt a number of cards which are hidden from his opponents, followed by 1 or more rounds of betting. Additional common cards visible to both players are dealt between betting rounds. 3 common cards (the flop) are dealt after the first betting round and 1 more card is dealt after each of the 2nd and 3rd betting rounds, the turn and river card respectively. After the final betting round each player makes the best possible hand from a combination of his cards and the common cards and the player with the strongest hand wins the total amount bet by all players(the pot). If multiple players have a hand of equal strength they split the pot. The variant being considered is a 2 player game (heads up) with fixed-size(limit) betting and the maximum number of bets is also limited.

The game tree for the complete game of limit holdem is very large. With 4 bets allowed each round there are a total of $3.19*10^{17}$. [2]. We consider a simplified version of the game, where the player's hands are abstracted into a small number of discrete cases and only consider betting on the river. The hand categories for the river are Premium, Strong, Average, Weak which may be subsequently referred to as A, B, C, D respectively (in order of decreasing strength) for conciseness. The fictitious play algorithm will be used to approximate optimal strategies.

1.1 Overview

Section 2 provides a formulation of the game in normal form, followed by applicable mathematical results. Section 3 outlines the fictitious play algorithm. Experiments are presented in Section 4 and a discussion of the results in Section 5.

2 Model

2.1 Formulation

The problem can be formulated as a normal form game as follows:

1. $N = \{1, 2\}$ the set of players.
2. T_i, the type of player i defined by the hand the player holds, for $T = \{A, B, C, D\}$.
3. A_i, for $i = 1, 2$ the set of strategies for player i
4. u: A_1 x T_1 x A_2 x $T_2 \rightarrow \mathbb{R}$, the payoff function

A strategy in the formulation specifies what a player does at all the nodes which are his decision point. As an example with one bet allowed player 1's

*This thesis was prepared in partial fulfillment of the requirements for the Degree of Bachelor of Science in Knowledge Engineering, University of Maastricht, supervisor: Dr. Gijs Schoenmakers

strategy set is {CF, CC, B} for checking and then folding to a bet, checking and calling and betting. The abbreviations used are F for folding, B for betting and C stands for both checking and calling but these two possibilities are never both possible. Player 2 has 2 decision points, one after player 1 checks and one after he bets. His strategy set is {CF, CC, BF, BC} where the first letter is his decision in the case where player 1 checks and the second the decision when player 1 bets.

A player's payoff is his expected winnings, including his share (50%) of the amount already in the pot. For example if the pot is 10, player 1 makes a bet of 10 and player 2 calls and wins by showing a stronger hands their payoffs are -15 and 15 respectively.

2.2 Extensive form

The problem can also be formulated as an extensive form game. An advantage of this formulation is that it allows explicit representation of a number of additional aspects, like the sequencing of players' possible moves, their choices at every decision point, the information each player has about the other player's moves when he makes a decision and his payoffs for all possible game outcomes. A strategy for a player in an extensive form game prescribes the probability of taking each of the available actions at each decision node where it is the player's turn to act. A strategy at a given node is pure if a single action is taken with probability 1 and mixed if multiple actions have a non-zero probability. A best response is a strategy that maximizes a player's winnings against a specified strategy for the opponent.

2.3 Theoretical results

John von Neumann proved in 1928 [7] that for every two-player, zero-sum game with finitely many strategies, there exists a value V and a mixed strategy for each player, such that: (i) Given player 2's strategy, the best payoff possible for player 1 is V (ii) Given player 1's strategy, the best payoff possible for player 2 is $-V$. These strategies are called optimal.

2.4 Nash equilibrium

A pair of strategies such that neither player can benefit by unilaterally changing his own strategy is referred to as a Nash equilibrium. John Nash showed in 1951 [5] that if mixed strategies are allowed an equilibrium exists for every game with a finite number of players in which each player can choose from finitely many pure strategies.

3 Fictitious Play

Fictitious play, introduced by G.W. Brown in 1951 [3], is a mathematical technique for the finding optimal strategies in games. Starting from arbitrary player strategies, iterate the following steps: 1. Each player calculates a best response against the opponent's strategy. 2. The player's updated strategy is to play the best response with probability $\frac{1}{N}$ and his previous strategy with probability $\frac{N-1}{N}$. These steps are repeated until the strategies are stable. Berger [1] points out that what modern mathematicians refer to as fictitious play differs from Brown's version in a subtle detail: Since Robinson's paper in 1951 [6] both players update their strategies simultaneously, while Brown states the players update their strategies alternatingly. The results in the following sections were obtained using the simultaneous update of strategies. Fictitious play has been proven to converge if both players have only a finite number of strategies and the game is zero sum [6]. In the context of heads up poker it can be viewed as a set of learning rules designed to produce agents capable of approaching optimality [4].

4 Experiments

A number of experiments were performed with the variables being:

- The number of bets allowed
- The size of the bet
- The probabilities of the players' hand types

The size of the pot was kept constant at 1 throughout all the experiments as only the ratio of the pot size to the size of the bet affects the player strategies. The term game value will be used to refer to player 1's payoff. Since the pot size is 1 and each player has contributed an equal part of the pot 0.5 and -0.5 are upper and lower bounds for the value, which each player could achieve with the simple strategy of folding all of his hands when faced with a decision. Due to practical constraints the fictitious play algorithm could only be run for a finite number of iterations. Convergence was verified by finding best response strategies for each player against the opponent's strategy and validating that the game value for each player when using the best response strategy is within ϵ of the value with both players using the strategies after the final iteration of the algorithm. In the simulations ϵ was set equal to 10^{-3}.

4.1 Half Bet Game

In the half bet game only the first player is allowed to bet. For the size experiment a bet size of 0.5 chips was chosen. The starting hand distributions are 0.1, 0.2, 0.3 and 0.4 for each player to hold a hand of type A, B, C, D respectively. The strategy set is {C, B} (checking or betting) for player 1 and {F, C} for player 2 (folding or calling after player 1 bets). The optimal strategies are shown in Figure 1. CIP for each player is the total he has bet up to that node. The first value displayed for each hand at every node is the probability of reaching the node with that hand. The number in parentheses is the payoff of holding the hand at that node. For player 1 the probability of having of having a hand of type A or B at node 2 is 1, hence he is playing a pure strategy of always betting with those hands. With a hand of type D he is playing a mixed strategy of checking with a probability of 0.75 and betting with probability 0.25 . When he bets, the 2nd player indifferent between calling and folding a type C hand. If player 2 calls and shows a weaker hand he loses 0.5 (player 1's bet) while if he shows a stronger hand he wins 1.5 (player 1's bet plus 1 that is already in the pot). So the value of calling with a type C hand for player 2 is $0.1*(-0.5)+0.2*-0.5+0*0.5+0.4*0.25*3 = 0$. Player 2 in turn is calling with an average hand with a probability that makes player 1 indifferent to bluffing with a type D hand. This pattern of one player betting a combination of his strongest and weakest with a ratio that makes the other player indifferent between folding or calling with his medium strength hands and that player responding by calling with a probability that makes the first player indifferent between checking or bluffing with his weakest hands was observed in many of the games that were analyzed. Having the strategic option of betting gives player 1 the advantage in this game, with a value of +0.035.

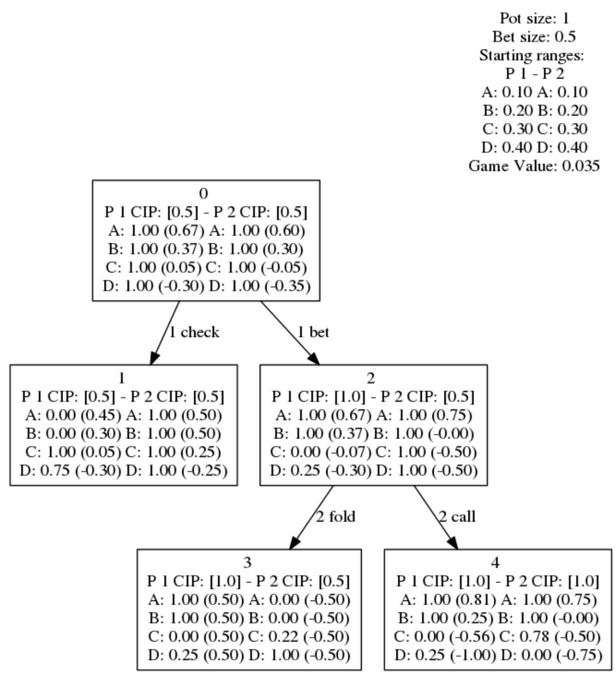

Figure 1: Half Bet game

4.2 One Bet Game

In this game players have the same starting hand distributions as the half-street game but player 2 has the option of betting if player 1 checks. The strategy sets are {CF, CC, B} and {CF, CC, BF, BC} for players 1 and 2 respectively. A number of one bet games were solved for different starting hand distri-

butions and bet sizes, while the size of the pot was held constant at 1. The following hand distributions were used, where the values are the probabilities of holding a hand of type A, B, C, D respectively:
$R1 = [0.1, 0.2, 0.3, 0.4]$
$R2 = [0.25, 0.25, 0.25, 0.25]$
$R3 = [0.5, 0.0, 0.0, 0.5]$
$R4 = [0.05, 0.3, 0.45, 0.20]$
Distribution R1 was chosen because it resembles some situations in the full game of poker. Distribution R3 includes only very strong and very weak hands while distribution 4 mostly consists of medium strength hands. The results for bet sizes of 0.5, 1, 2, 3, 5 are shown in tables 1, 2, 3, 4, 5 respectively. Table 13 in the appendix shows the optimal strategies for each range combination with a bet size of 0.5. The game was also solved for bet sizes of 0.1 and 10 but the results are not included because the player strategies were very simple for both cases. If the bet size is 10 the players are never bluffing and only bet hands of type A. With a bet size of 0.1 the players bet with exactly the hands which are better than 50% of the opponent's range and for most combinations of hand distributions call a bet 100% of the time because they are risking 0.1 to win 1.1 (the pot plus the bet of 0.1) and only need a 1.0/11 chance of winning. For both cases the value of the game is very close to 0 when the players have the same hand distribution.

The game tree for the case when both players have the R1 hand distribution with a bet size of 0.5 is shown in Figure 2. Player 1's betting strategy is the same as the half street game. The value of the game is -0.02, so player 2 has an advantage in this game.

Table 1: 1 Bet, Bet size 0.5 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.010	-0.196	-0.183	-0.080
R2	0.162	-0.023	-0.083	0.137
R3	0.150	0.042	-0.000	0.142
R4	0.047	-0.169	-0.175	-0.008

The optimal strategies for this case can be found in Table 13 in the Appendix

Table 2: 1 Bet, Bet size 1 Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.030	-0.200	-0.225	-0.090
R2	0.166	-0.026	-0.094	0.152
R3	0.185	0.063	0.000	0.212
R4	0.025	-0.187	-0.237	-0.035

Table 3: 1 Bet, Bet size 2 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.014	-0.208	-0.242	-0.054
R2	0.168	-0.021	-0.042	0.177
R3	0.225	0.037	0.000	0.271
R4	0.035	-0.200	-0.287	-0.004

Table 4: 1 Bet, Bet size 3 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.020	-0.200	-0.225	-0.050
R2	0.158	-0.000	-0.000	0.178
R3	0.218	0.000	0.000	0.288
R4	0.026	-0.206	-0.300	-0.011

Table 5: 1 Bet, Bet size 5 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.017	-0.167	-0.158	-0.055
R2	0.148	-0.000	-0.000	0.173
R3	0.157	0.000	0.000	0.283
R4	0.025	-0.200	-0.288	-0.015

4.3 Two Bets Game

The same experiments from 4.2 were repeated for the game were up to 2 bets are allowed. Tables 6 through 10 display the results for bet sizes of 0.5, 1, 2, 3, 5. The optimal strategies with a bet size of 1 and both players having each hand with a probability of 0.25 are shown in Figure 3. Interestingly the first player is checking his strongest hands with a non-zero probability. In that way he protects his average (type C) hands against a bluff by player 2 by forcing him to reduce his bluffing frequency. When player 1 bets, player 2 is indifferent between calling and folding with both type B and C hands. There is no distinction between the two due to player 1 betting only type A and D hands.

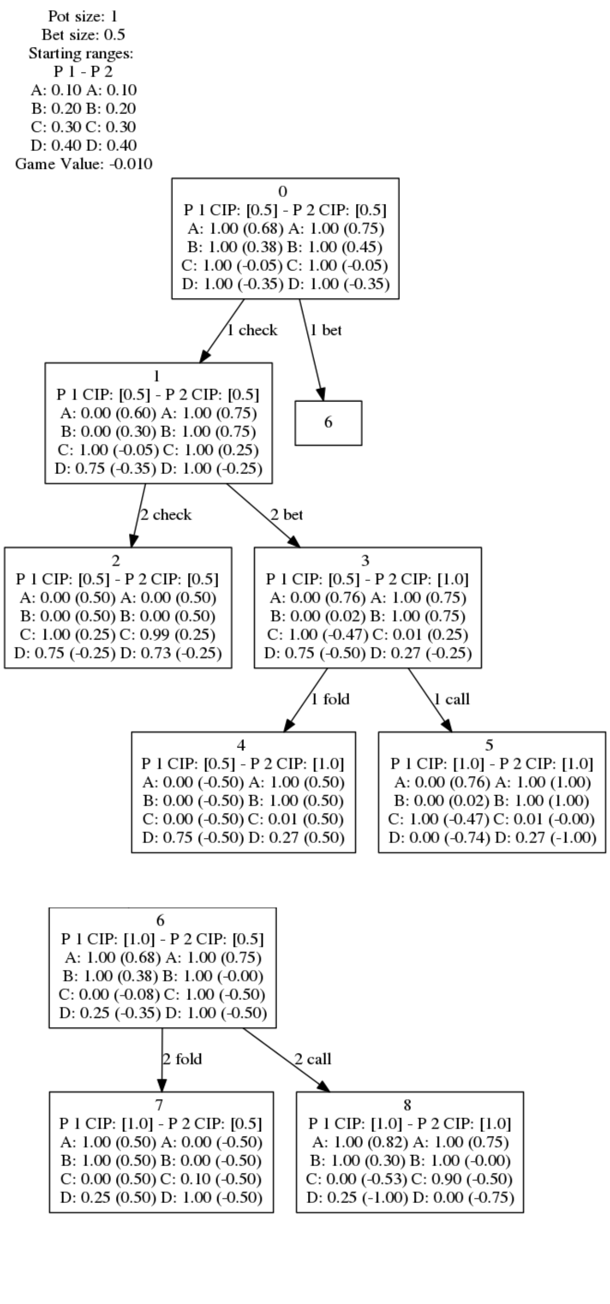

Figure 2: 1 Bet game

Table 6: 2 Bets, Bet size 0.5 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.018	-0.188	-0.183	-0.080
R2	0.158	-0.018	-0.083	0.134
R3	0.150	0.042	-0.000	0.141
R4	0.038	-0.165	-0.175	-0.016

Table 7: 2 Bets, Bet size 1 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.022	-0.200	-0.225	-0.076
R2	0.169	-0.026	-0.094	0.157
R3	0.199	0.063	-0.000	0.212
R4	0.032	-0.181	-0.238	-0.022

Table 8: 2 Bets, Bet size 2 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.014	-0.208	-0.242	-0.052
R2	0.169	-0.021	-0.042	0.177
R3	0.228	0.038	-0.000	0.274
R4	0.035	-0.200	-0.288	-0.004

Table 9: 2 Bets, Bet size 3 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.020	-0.200	-0.225	-0.049
R2	0.158	-0.000	-0.000	0.181
R3	0.218	0.000	-0.000	0.292
R4	0.026	-0.206	-0.300	-0.009

Table 10: 2 Bets, Bet size 5 - Game Value

P1 \ P2	R1	R2	R3	R4
R1	-0.017	-0.167	-0.158	-0.055
R2	0.148	-0.000	-0.000	0.174
R3	0.157	0.000	-0.000	0.283
R4	0.025	-0.200	-0.288	-0.015

5 Results

Tables 11 and 12 show the game value (which is player 1's payoff) averaged over all bet sizes for the one and two bets games. This can give an indication about which player has the advantage in each case. The average values for the two games are similar and in both cases the second player seems to have the advantage. When the players have the same hand distribution the game value is always non-positive. In the cases where the game value is positive it is a result of player 1's distribution being stronger and the corresponding game value for player 2 with the distributions swapped is at least as high. A possible interpretation of this advantage is that player 2 knows player 1's decision before he acts, which gives him an information advantage.

Table 11: Game value averaged over bet sizes -1 bet

P1 \ P2	R1	R2	R3	R4
R1	-0.018	-0.194	-0.207	-0.066
R2	0.160	-0.014	-0.044	0.164
R3	0.187	0.028	0.000	0.239
R4	0.032	-0.193	-0.258	-0.015

Table 12: Game value averaged over bet sizes -2 bets

P1 \ P2	R1	R2	R3	R4
R1	-0.018	-0.193	-0.207	-0.062
R2	0.160	-0.013	-0.044	0.165
R3	0.190	0.028	-0.000	0.240
R4	0.031	-0.190	-0.258	-0.013

5.1 Effect of the bet size

This section examines the effect of the bet size on the value of the game with 1 or 2 bets allowed. Figures 4 and 5 show the game value as a function of the bet size when both player have distribution R1 and R4 respectively with one or two bets allowed. In both cases the advantage of the second player is largest when the size of the bet is equal to the pot.

5.2 Number of betting rounds

In limit poker games offered by casinos the players are typically allowed to make more than 1 bet per betting round, usually 3 or 4 bets. As shown in figures 4 and 5 with a bet size of 2 or higher the difference in the value of the game with 1 or 2 bets is very small. Some of the games were also solved with 3 or 4 bets allowed but that had virtually no effect on the value or the player strategies. In fact when playing the equilibrium strategies the deeper parts of the tree with multiple bets are never reached except when both players hold the strongest possible hand. That holds for different starting hand distributions or even abstracting the starting hands into more than 4 distinct categories. The betting pattern described in section 4.1 can provide some insight into this result. Player 1 usually bets some combination of his strongest and weakest hands. After he bets it is not advantageous for Player 2 to raise with medium strength hands because player 1 could play optimally by folding his weakest hands and calling (or raising) with the strongest. Since player 1 bets his strongest and weakest hands he is left with medium strength hands when he checks and raising is not his best option when facing a bet. For this reason the results for the games with more than 2 bets will not be presented.

5.3 Different starting hand distributions

This section examines the relative value of different hand distributions. Distribution R1 is the weakest, both players are at disadvantage against any of the other 3 distributions. That is primarily due to the

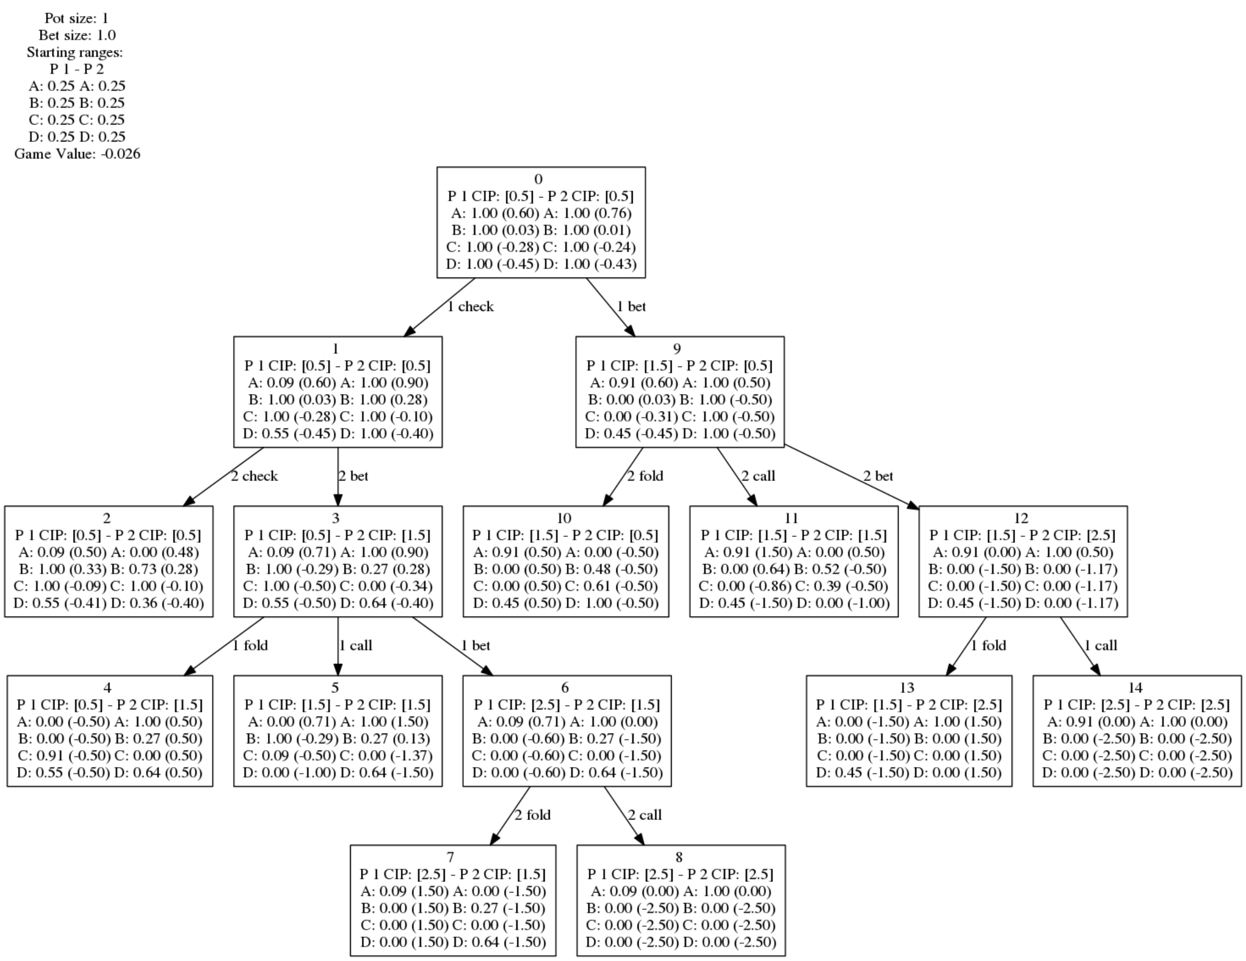

Figure 3: 2 Bets game

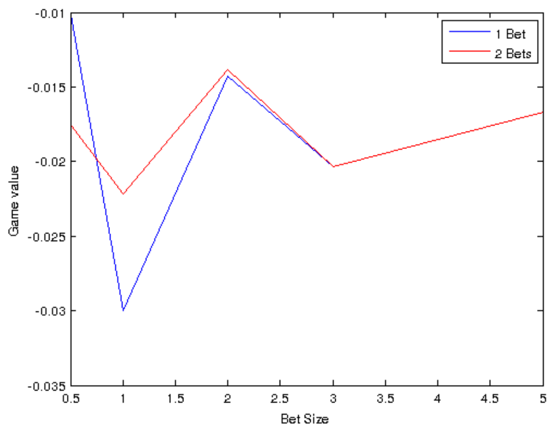

Figure 4: Game value with both players having distribution R1

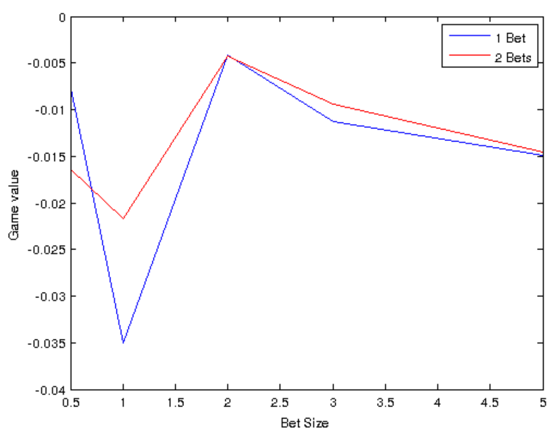

Figure 5: Game value with both players having distribution R4

higher fraction of weak hands. Distribution R3 is clearly the strongest with both players having an advantage in all cases. R4 is weaker than both R2 and R3. These results indicate that having a combination of strong and weak hands, which is referred to as a polarized hand distribution, is preferable to having mostly medium strength hands. Poker is an incomplete information game and the player with a polarized hand distribution has an information advantage by knowing when he is likely to have the best hand. On the other hand holding mostly medium strength hands leaves a player vulnerable to an opponent betting a combination of stronger and weaker hands.

6 Conclusion

It was shown that the fictitious play algorithm can produce strategies which are very close to optimal for this model of a poker game. The results strongly indicate that the player who acts second has an advantage.

An interesting addition to the model would be to allow cards to be dealt between betting rounds with certain probabilities, changing the strength of the players' hands. In extensive form this can be modeled as decision nodes controlled by a fictitious player called nature, with the children corresponding to the chance outcomes.

References

[1] U. Berger. Brown's original fictitious play. In *Journal of Economic Theory*, volume 135, page 572–578, 2007.

[2] M. Bowling, N. Burch, M. Johanson, and O. Tammelin. Heads-up limit holdem poker is solved. volume 347, page 145–149, 2015.

[3] G. W. Brown. Iterative solutions of games by fictitious play. In *Activity Analysis of Production and Allocation*, page 374–377, 1951.

[4] W. Dudziak. Using fictitious play to find pseudo-optimal solutions for full-

scale poker, 2006. Retrieved from http://ww1.ucmss.com/books/LFS/CSREA2006/ICA4506.pdf.

[5] J. Nash. Non-cooperative games. volume 54, page 286–295, 1951.

[6] J. Robinson. An iterative method of solving a game. volume 54, pages 296–301, 1951.

[7] J. v. Neumann. Zur theorie der gesellschaflspiele (on the theory of parlor games). volume 100, pages 295–320, 1928.

Appendix

Table 13: One bet, bet size 0.5 -Strategies

	R1	R2	R3	R4
R1	P1: A:(0, 0, 1) B:(0, 0, 1) C:(0, 1, 0) D:(0.75, 0, 0.25) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.09, 0.74, 0.02, 0.15) D:(0.75, 0, 0.25, 0)	P1: A:(0, 1, 0) B:(0, 1, 0) C:(0.22, 0.78, 0) D:(1, 0, 0) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.42, 0.58, 0, 0) D:(0.33, 0, 0.67, 0)	P1: A:(0, 0.55, 0.45) B:(0.12, 0.51, 0.36) C:(0.12, 0.51, 0.36) D:(0.45, 0, 0.55) P2: A:(0, 0, 0, 1) D:(0.44, 0.22, 0.22, 0.11)	P1: A:(0, 0, 1) B:(0, 0, 1) C:(0, 1, 0) D:(0.75, 0, 0.25) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.33, 0.67, 0, 0) D:(0.24, 0, 0.76, 0)
R2	P1: A:(0, 0, 1) B:(0, 0, 1) C:(0, 0.5, 0.5) D:(0.5, 0, 0.5) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.06, 0.32, 0.1, 0.52) D:(0.87, 0, 0.13, 0)	P1: A:(0, 0.45, 0.55) B:(0, 1, 0) C:(0.54, 0.46, 0) D:(0.82, 0, 0.18) P2: A:(0, 0, 0, 1) B:(0, 0.09, 0.03, 0.88) C:(0.42, 0.58, 0, 0) D:(0.36, 0, 0.64, 0)	P1: A:(0, 0.72, 0.28) B:(0.33, 0.45, 0.22) C:(0.33, 0.45, 0.22) D:(0.28, 0, 0.72) P2: A:(0, 0, 0, 1) D:(0.45, 0.22, 0.22, 0.11)	P1: A:(0, 0, 1) B:(0, 0, 1) C:(0, 0.5, 0.5) D:(0.5, 0, 0.5) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.33, 0.67, 0, 0) D:(0.25, 0, 0.75, 0)
R3	P1: A:(0, 0, 1) D:(0, 0, 1) P2: A:(0, 0, 0, 1) B:(0, 0.67, 0, 0.33) C:(0, 0.67, 0, 0.33) D:(0.34, 0, 0.66, 0)	P1: A:(0, 0, 1) D:(0.67, 0, 0.33) P2: A:(0, 0, 0, 1) B:(0.17, 0.83, 0, 0) C:(0.17, 0.83, 0, 0) D:(0, 0, 1, 0)	P1: A:(0, 0.5, 0.5) D:(0.5, 0, 0.5) P2: A:(0, 0, 0, 1) D:(0.25, 0.25, 0.25, 0.25)	P1: A:(0, 0, 1) D:(0.67, 0, 0.33) P2: A:(0, 0, 0, 1) B:(0.09, 0.42, 0.09, 0.4) C:(0.09, 0.42, 0.09, 0.4) D:(0, 0, 1, 0)
R4	P1: A:(0, 0, 1) B:(0, 0, 1) C:(0, 0.39, 0.61) D:(0.87, 0, 0.13) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.07, 0.35, 0.1, 0.48) D:(0.88, 0, 0.12, 0)	P1: A:(0, 1, 0) B:(0, 1, 0) C:(0.44, 0.56, 0) D:(1, 0, 0) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.35, 0.65, 0, 0) D:(0.33, 0, 0.67, 0)	P1: A:(0, 0.76, 0.24) B:(0.27, 0.55, 0.19) C:(0.27, 0.55, 0.19) D:(0.24, 0, 0.76) P2: A:(0, 0, 0, 1) D:(0.45, 0.22, 0.22, 0.11)	P1: A:(0, 0, 1) B:(0, 0, 1) C:(0, 0.39, 0.61) D:(0.87, 0, 0.13) P2: A:(0, 0, 0, 1) B:(0, 0, 0, 1) C:(0.33, 0.67, 0, 0) D:(0.25, 0, 0.75, 0)

The strategy sets are {CF, CC, B} for player 1 and {CF, CC, BF, BC} for player 2.