Advanced Project Second Year - FIT2044
Bayesian Poker Player

David Northfield
Student ID: 20274793

Supervisors: Steven Mascaro and Ann Nicholson

October 31, 2007

1 Introduction

The Bayesian Poker Player, or BPP, is an AI, artificially intelligent, poker player that uses a Bayesian network for its decision making process. It was originally written by Nathalie Jitnah for her honors thesis under the supervision of Ann Nicholson and Kevin Korb in 1993. The first version played 5-card stud but it has since been modified to play Texas Hold'em. Since its creation, BPP has been managed by Ann and Kevin on a part-time basis, and a number of other Monash University staff and students, including Steven Mascaro, have added their own alterations.
BPP is a rather large program consisting of a variety of components in a number of programming languages. There are python modules that encompass the main functions of BPP, a Bayesian network setup with initial probabilities for a Texas Hold'em poker game, and java script and html that facilitate the GUI.
The purpose of this particular project was to modify BPP to allow it to bluff better and/or sandbag. BPP was already programmed with a primitive bluffing algorithm, but was not programmed to sandbag. I will discuss both of these concepts later on.

This was a solo project, however there were three of us working on our own solutions. Robiah Elsaadi, Kym McGain, and I had regular meetings with Steven or Ann to discuss our progress, but the actual work was done individually.

As BPP is a poker player, I will be using a number of poker terms and some shorthand to explain a few things. To clarify this, I'll outline the basics now.
Since poker is a card game, it involves the use of a standard deck of cards. These cards can be divided in two ways that are relevant to poker: suit and rank.
Card suits can be one of four types: clubs, hearts, diamonds, or spades. In shorthand, these are represented by 'c', 'h', 'd', and 's' respectively.
Card ranks are the number/face values of each card. These values are, in rank order: 2, 3, 4, 5, 6, 7, 8, 9, 10, Jack, Queen, King, and Ace. In shorthand, these are: 2, 3, 4, 5, 6, 7, 8, 9, 10 ('*T*'), 'J', 'Q', 'K', and 'A' respectively. The ace is a special case in ranks, as there are times when it is the lowest ranked card, below the two. I will explain these circumstances later.
As an example of both, the ace of spades would be written as 'As', and the five of hearts would be written as '5h'.

1.1 Poker - Texas Hold'em

First of all, poker is a card game where players attempt to form the best 'hand' with their cards that they can in order to beat every other player. Hands are combinations of five

cards that are given a particular name and standing depending on their contents.
The following is a list of poker hands from lowest to highest standing, where each hand beats all those preceding it:

- A **High Card** is where a player has no two cards of equal rank, the five cards are not of the same suit and are not of consecutive ranks. The strength of this hand is marked by its highest ranked card.
  For example, the hand 3c, 5c, 6d, 7h, Kd is King high.
- A **Pair** is where a player has just two cards of equal rank. The strength of the hand is based on the rank of the paired cards.
  For example, the hand 3d, 6h, 6d, 8h, Js is a pair of sixes.
- **Two Pair** is where a player has two different pairs of matching ranked cards. The strength of the hand is first based on the higher ranked pair, then the lower ranked pair.
  For example, the hand 4s, 4d, 8h, 8c, 9s is two pair, eights over fours.
- A **Triple** or **Three of a Kind** is where a player has three cards of the same rank. The strength of the hand is based on the rank of the tripled cards.
- A **Straight** is where a player has five cards all of successive rank. The strength of a straight is based on the highest ranked card. In the case of a straight, an ace can either be the highest or lowest card.
  For example, if the straight is 10, J, Q, K, A, then the ace is considered the highest card, thus making this the highest possible straight. However, if the straight is A, 2, 3, 4, 5, then the ace is considered the lowest card, thus making this the lowest possible straight.
- A **Flush** is where the player has five cards of the same suit. The strength of a flush is based on the highest ranked card.
  For example, the hand 5h, 8h, 10h, Jh, Qh, is a King high hearts flush.
- A **Full House** is where a player has both a pair and a triple. The strength of a full house is first based on the rank of the tripled cards, and then, in certain game types, on the paired cards.
  For example, the hand 7h, 7c, Ks, Kd, Kc is a full house, 'kings full of sevens'.
- **Four of a Kid** is where a player has, as one would expect, four cards of the same rank. The strength of this hand is based on the rank of the matched cards.
  For example, the hand Ad, Ac, Ah, As, Kh is a four of a kind of aces.
- A **Straight Flush** is where, as one might assume, a player has both a straight and a flush. This means that their cards are all of the same suit and are of consecutive rank. The strength of a straight flush is based on its highest ranked card.
  For example, 5h, 6h, 7h, 8h, 9h is a straight flush.

- A **Royal Flush** is a special type of straight flush. It is formed when the straight runs from ten to ace.
  For example, 10s, Js, Qs, Ks, As is a royal flush of spades.

In Texas Hold'em, each player only receives two cards for themselves, called 'hole' or 'pocket' cards. In addition, five 'community' or 'board' cards are dealt to the table for all players to use. Players use a combination of their hole cards and the community cards to form a hand, as described above.
The game is played in rounds as follows:

- **Blinds** are paid. Blinds are pre-set amounts that must be paid by the two players to the left of the dealer. Generally, there is a big blind and a small blind. The big blind is paid by the player to the immediate left of the dealer, and the small blind is paid by the player two seats left of the dealer.
- **Hole cards** are then dealt to each player.
- **Betting** now takes place starting with the player on small blind. When betting, players have four options. They can fold, meaning that they will take no further part in the hand. They can check, meaning that they have already paid the highest amount bid and do not wish to increase the stakes. They can bet, meaning that no amount has been wagered yet and they are placing the original stakes for the round. Finally, they can raise, meaning that an amount has already been bet and the player wishes to increase the wager. At the end of a betting round, all of the money is placed in a 'pot' that is collected by the winning player at the end of the game.
- **The Flop** comes next, where the first three community cards are dealt. At this stage, players can form a full hand with their two hole cards and the three community cards.
- **Betting** occurs again after the flop, starting with the player to the left of the dealer.
- **The Turn** is the fourth community card. The players can now form a hand with their two hole cards and any three community cards, or with the four community cards and one of their hole cards.
- **Betting** occurs again.
- **The River** is the fifth and final community card. Players can form hands like those in the Turn, however if the five community cards form a better hand than any combination including at least one of the player's hole cards, then the player is said to be "Playing the table."
- **Betting** occurs for the last time.

- **The Showdown** is where the remaining players show their hole cards in order to reveal the value of their hand. The player with the best hand wins.

If, at any point during the game, all players fold except for one, then that remaining player wins.
BPP plays by a particular set of additional rules, as explained below:

Fixed-Limit. This means that betting is limited to a fixed value. In this case, big blind is $10, small blind is $5, bets before the turn are $10 and $20 after the turn.

Raise Limit. This means that only a fixed number of raises may be made. In this case, only three per round.

Heads-Up. BPP is capable of playing in multi player tournaments, however the standard interface is set up for a one-on-one game style.

1.2 Bayesian Networks

A Bayesian network is a series of connected nodes that each contain a table of probabilities related to a single event or stimulus. Nodes are connected based on the relations of the things that they represent, and allow their probabilities to affect one another appropriately.
For example, let's say that a person's chances of getting cancer are based solely on whether or not they smoke, and the amount of pollution where they live. Nodes representing smoking and pollution would be considered parent nodes, and a node representing a person's chance of getting cancer would thus be the child node of smoking and pollution. Due to this relation, if the probability of an individual being a smoker increased then the chance of them getting cancer would also increase.
BPP uses a program called Netica[3] to generate its Bayesian network. The cancer example is a sample network that comes with the Netica software available at the Norsys website.

2 Project Aims

The aim of this particular project was to add functionality to BPP in bluffing and/or sandbagging. Both of these relate to a poker concept known as pot odds, which involves the calculation of bets based on your chance of winning and the pot size.

2.1 Bluffing

Bluffing is a poker strategy where a player attempts to convince their opponent that their hand is better than it really is in order to make their opponent fold. Its main purpose is to provide a winning option for a player with a losing hand. In Texas Hold'em, bluffing can be made easier by the visible community cards. When there is, for example, a 7, 8 and 9 on the table after the flop, a player whose hole cards are 5 & 6, 6 & 10 or 10 & J would have a straight. If a player does not have one of these combinations, they might bet heavily in an attempt to convince their opponent that they do. This is the basic strategy of bluffing.
Initially, BPP could bluff by using a fixed set of thresholds and generating a random number[1] to determine when it would bluff. The bluffing method itself was simple. BPP would check the probability of each possible hand based on the community cards displayed and two unknown cards. It then takes the hand that has the highest probability that is better than its actual hand and plays as though that were its hand.

2.2 Sandbagging

Sandbagging, or slow play as it is also called, is a strategy whereby a player attempts to give the impression that their hand is weaker than it really is in order to make their opponent believe that they will win. The use of this strategy increases the chance that an opponent will to play all the way to the showdown, thus increasing the final pot. In many ways, it is the complete opposite to bluffing. Sandbagging is also a risky strategy, as a player may decide early on that their hand is good enough to sandbag, yet after the *turn* they may be outdrawn**??** by their opponent and end up handing them the pot.
BPP had no method for sandbagging before we began the project, and that is why I decided to focus on developing a sandbagging function.

2.3 Pot Odds and Implied Odds

"Pot odds" is the name given to a mathematical principle that helps a player decide how much their hand is worth, and when and how much to bet. Both bluffing and sandbagging are related to pot odds in the sense that both of them augment the sensible nature of play by altering them. In the case of bluffing, a player goes against pot odds by overstating the value of their hand, whereas in sandbagging they understate their hand's value.
Pot odds use a number of formulae to calculate precise values, though players will have

[1] In this context, a random number is a floating point number from 0 to 1 that is generated by the python function "random.random()".

to make the calculations in their head and as such the values will generally be approximations. A simple pot odds formula could be as follows:

$$\frac{a}{b} = c : 1$$

Where 'a' is the current value of the pot, 'b' is the amount that must be called to play on, and 'c:1' is the required win rate in order to break even in this situation.
Let's say that you're at the river and the pot is \$100. Now let's assume that in order to see the showdown you must call a bet of \$20. Following the above formula:

$$\frac{\$100}{\$20} = 5 : 1$$

This means that you would need to win at least once every five times you are in this position in order to break even. Now this formula isn't very tough to remember, but you can use it with a more complicated formula to work out whether or not a *bet* or *call* is wise. In this case, you would assume that your hand can beat any other hand if you get a certain card or one of a certain set of cards and will lose otherwise. That formula looks something like this:

$$(pot \times P(out)) + (-bet \times P(miss)) = Expected\ winnings$$

Where 'pot' is the amount currently in the pot plus all current bets, 'P(out) is the probability of drawing an out[2], 'bet' is the amount a *bet/call* will cost you, and 'P($miss$)' is the probability of not drawing the card you need. P($miss$) and P(out) can be calculated as the number of misses or outs respectively, divided by the number of cards remaining in the deck plus the number of cards dealt to other players.
For example, let's say you're at the turn, you're playing against only one opponent, the pot has \$100 in it, your opponent has bet \$20, you have one heart in your hand, and there are three hearts on the table. This means that you have a flush draw on the river. To substitute all this into the formula, you would get:

$$((\$100 + \$20) \times \frac{9}{46}) + (-\$20 \times \frac{37}{46}) = \$23.48 - \$16.09 = \$7.39$$

So in this example, a call of \$10 would have an average pay out of \$7.39, assuming that if you make the flush on the river that you'll win, but you'll lose otherwise. It should also be noted that:

$$\frac{outs}{misses} = Probability\ of\ winning\ the\ hand$$

Which means that, in this example, the chance of winning the hand is $\frac{9}{37}$, or a little over $\frac{1}{4}$.

Let's now combine the two situations. The probability of winning the hand is around $\frac{9}{37}$, and the first formula says that for this bet you need a win ratio of 5:1, or $\frac{1}{5}$. Since the chance of winning is greater than the required win ratio, it's definitely a good idea to bet, and, thanks to the second formula, you know that you have an expected win of \$7.39 per hand.[2][4]

[2] An 'out' is any card that, when combined with a player's cards, completes a poker hand.

3 Research

3.1 Personal Experience

As I was already a fan of poker, this project was a natural choice for me. I have experience playing against friends in person and against other people online. As such, I have a fair amount of personal experience at bluffing and sandbagging, which are strategies that I implement quite frequently online, though less in person as my friends can often read me. This brings up the point of physical tells. BPP is incapable of reading a player's expression or other physical signs that a human would be able to observe. It has only the opponent's betting behavior to base its judgement on. This aided me to a degree in my decisions for forming a sandbagging strategy.

3.2 Friends

As previously stated, I have played poker quite a bit with my friends. I decided to play a simple scenario with one of my friends in order to determine what my personal actions would be given the sandbagging opportunity. In this scenario, I ran through each round twice, once for each player as the dealer, and determined the best action for each round, taking into account the possible actions of my opponent. Upon finishing, I wrote up the decisions into a table and discovered that I had the same decisions for when my opponent *checked* or *bet* as I did when they *called* or *raised* respectively.

3.3 Professional Strategies

Poker professionals are often asked their opinion on play styles or techniques. My research came across a number of professional tips regarding sandbagging, including that of Howard Lederer. Lederer warns of the dangers of sandbagging, and not just because it gives an opponent the chance to make a better hand. Sandbagging also needs to take into account the way new cards will affect the behavior of an opponent. I will illustrate this with an excerpt from an article written by Lederer:

> "Consider the following example.
>
> You're in late position in a No-Limit Hold 'em ring game. A player raises in early position. You look at your cards, see pocket 8s, and decide to call. The flop is absolutely perfect: Qh 8h 2d. You've hit your set and, with the Queen

> out there, chances are your opponent has something maybe A-Q, maybe pocket Kings or Aces. He bets the flop.
>
> Many players will just call in this spot, hoping to get their opponent to bet on the turn. But a raise is usually the better play. If you just call, you risk seeing a heart on the turn. I don't think you need to be especially worried about the flush beating your set. You might get your set beat by a flush draw even if you raise. However, you do need to be concerned about the effect the third heart will have on your opponent. He very well might suspect that you were on the flush draw and he'd no longer be willing to commit a lot of money to the hand, even if he has Aces.
>
> In fact, any King, Jack, 10, 9 or a card that pairs the board is likely to give your opponent pause. If he bets on the turn and you raise, you're signaling that the turn card helped you. In effect, you're saying that you liked the flop enough to call and the turn improved your hand in some way. You're announcing that you can beat one pair.
>
> So the flop very well may be the only time when your opponent is willing to make a stand with a single pair. If he bets the flop of Qh 8h 2d and you raise, he's likely to think that you're semi-bluffing – raising on a flush draw. At that point, he might feel compelled to protect his hand with large re-raise or perhaps an all in. When this happens, you'll take down a monster pot."[1]

The general rule that Lederer is promoting is that it is better to bet more heavily in earlier rounds and only to sandbag when an opponent is less likely to change their behavior, such as after the turn or flush, though this may vary depending on the specific game.

4 Chosen Solution

4.1 Instruction Vector

A hard-coded array of instructions holds the actions to be taken based upon four variables: who the dealer is, what round it is, what round BPP started sandbagging, and what the opponent's last action was. I created the array by combining all of my research methods. I began with my initial array, discussed a few games with a friend, and then applied what I had learned from professional opinions and strategies.

4.2 Shifting Threshold

I decided that to best reflect the professional opinion that sandbagging later is a better strategy, a simple static threshold would not suffice. As such, I opted for the static base threshold with an increasing additional value based upon the round. This allowed BPP to partly base its decision to sandbag on the current round.

5 Results

5.1 Explanation

BPP has a program called 'assessbot' that plays BPP against itself for a set number of sessions each consisting of a set number of hands. Results from tests using the 'assessbot' program have the following format:

bpp3_1_4 10:100 Session: 1 - ($675.00 | 0.68 sbu | 100.00% | 53.00)
bpp3_1_4 10:100 Session: 2 - ($-710.00 | -0.02 sbu | 50.00% | 50.50)
bpp3_1_4 10:100 Session: 3 - ($-595.00 | -0.21 sbu | 33.33% | 49.00)
bpp3_1_4 10:100 Session: 4 - ($-460.00 | -0.27 sbu | 25.00% | 48.75)
bpp3_1_4 10:100 Session: 5 - ($-290.00 | -0.28 sbu | 20.00% | 48.00)
bpp3_1_4 10:100 Session: 6 - ($1005.00 | -0.06 sbu | 33.33% | 49.00)
bpp3_1_4 10:100 Session: 7 - ($255.00 | -0.02 sbu | 42.86% | 49.29)
bpp3_1_4 10:100 Session: 8 - ($910.00 | 0.10 sbu | 50.00% | 50.25)
bpp3_1_4 10:100 Session: 9 - ($955.00 | 0.19 sbu | 55.56% | 50.56)
bpp3_1_4 10:100 Session: 10 - ($130.00 | 0.19 sbu | 60.00% | 49.90)
10 Sessions, 100 Hands
Sessions won: 60.00% (6)
Sessions won t(9) - value: 0.63
Earnings: 0.19 sbu ($187.50)
Earnings Sd: 0.67 ($674.12)
Earnings t(9) - value: 0.83
Hands won per session: 49.90% (49.90)
Showdowns per session: 55.60% (55.60)

Consider the line:
'bpp3_1_4 10:100 Session: 5 - ($-290.00 | -0.28 sbu | 20.00% | 48.00)'

'**bpp3_1_4**' is the version number, in this case version 3.1.4.
'**10:100**' is the number of sessions followed by the number of hands per session, in this case it is 10 sessions of 100 hands.
'**Session: 1**' is the current session, in this case 5.
'**$-290.00**' is the total gain by BPP over the entire session, in this case a gain of $675. A negative value is a net loss.
'**-0.28 sbu**' is the approximate cumulative gain in sbu (small betting units) per hand, in this case an approximate loss of .28 small betting units per hand thus far.
'**20.00%**' is the cumulative session win percentage, in this case 20%.
'**48.00**' is the percentage of hands won per session.

For the summary results:
'**10 Sessions, 100 Hands**' are the test parameters.
'**Sessions won: 60.00%** (**6**)' is the percentage and number of sessions won.
'**Sessions won t(9) - value: 0.63**' shows whether or not the number of sessions won is significant (ie. more or less than 50% and to what degree).
'**Earnings: 0.19 sbu** (**$187.50**)' is the approximate number of small betting units ($10) won per hand and the .
'**Earnings Sd: 0.67** (**$674.12**)' is the standard deviation of total earnings.
'**Earnings t(9) - value: 0.83**' shows whether or not the winning margin is significant (ie. more or less than breaking even and to what degree)
'**Hands won per session: 49.90%** (**49.90**)' is the cumulative average percentage of hands won over a full session.
'**Showdowns per session: 55.60%** (**55.60**)' is the cumulative average percentage of showdowns reached over a full session.

All simulations were tested with the unmodified bpp3.1.4 as the opponent.

5.2 Original BPP Version

The following results were taken by running the original BPP3.1.4 against itself.

bpp3_1_4 10:100 Session: 1 - ($675.00 | 0.68 sbu | 100.00% | 53.00)
bpp3_1_4 10:100 Session: 2 - ($-710.00 | -0.02 sbu | 50.00% | 50.50)
bpp3_1_4 10:100 Session: 3 - ($-595.00 | -0.21 sbu | 33.33% | 49.00)
bpp3_1_4 10:100 Session: 4 - ($-460.00 | -0.27 sbu | 25.00% | 48.75)
bpp3_1_4 10:100 Session: 5 - ($-290.00 | -0.28 sbu | 20.00% | 48.00)
bpp3_1_4 10:100 Session: 6 - ($1005.00 | -0.06 sbu | 33.33% | 49.00)
bpp3_1_4 10:100 Session: 7 - ($255.00 | -0.02 sbu | 42.86% | 49.29)
bpp3_1_4 10:100 Session: 8 - ($910.00 | 0.10 sbu | 50.00% | 50.25)

bpp3_1_4 10:100 Session: 9 - ($955.00 | 0.19 sbu | 55.56% | 50.56)
bpp3_1_4 10:100 Session: 10 - ($130.00 | 0.19 sbu | 60.00% | 49.90)
10 Sessions, 100 Hands
Sessions won: 60.00% (6)
Sessions won t(9)-value: 0.63
Earnings: 0.19 sbu ($187.50)
Earnings Sd: 0.67 ($674.12)
Earnings t(9)-value: 0.83
Hands won per session: 49.90% (49.90)
Showdowns per session: 55.60% (55.60)

bpp3_1_4 10:100 Session: 1 - ($-1525.00 | -1.52 sbu | 0.00% | 40.00)
bpp3_1_4 10:100 Session: 2 - ($1065.00 | -0.23 sbu | 50.00% | 47.50)
bpp3_1_4 10:100 Session: 3 - ($750.00 | 0.10 sbu | 66.67% | 50.33)
bpp3_1_4 10:100 Session: 4 - ($-1560.00 | -0.32 sbu | 50.00% | 48.75)
bpp3_1_4 10:100 Session: 5 - ($-1425.00 | -0.54 sbu | 40.00% | 47.80)
bpp3_1_4 10:100 Session: 6 - ($270.00 | -0.40 sbu | 50.00% | 47.83)
bpp3_1_4 10:100 Session: 7 - ($-285.00 | -0.39 sbu | 42.86% | 48.57)
bpp3_1_4 10:100 Session: 8 - ($2045.00 | -0.08 sbu | 50.00% | 49.75)
bpp3_1_4 10:100 Session: 9 - ($560.00 | -0.01 sbu | 55.56% | 50.44)
bpp3_1_4 10:100 Session: 10 - ($-725.00 | -0.08 sbu | 50.00% | 49.40)
10 Sessions, 100 Hands
Sessions won: 50.00% (5)
Sessions won t(9)-value: 0.00
Earnings: -0.08 sbu ($-83.00)
Earnings Sd: 1.23 ($1228.96)
Earnings t(9)-value: -0.20
Hands won per session: 49.40% (49.40)
Showdowns per session: 58.00% (58.00)

bpp3_1_4 10:100 Session: 1 - ($10.00 | 0.01 sbu | 100.00% | 47.00)
bpp3_1_4 10:100 Session: 2 - ($-2130.00 | -1.06 sbu | 50.00% | 41.50)
bpp3_1_4 10:100 Session: 3 - ($-945.00 | -1.02 sbu | 33.33% | 42.00)
bpp3_1_4 10:100 Session: 4 - ($855.00 | -0.55 sbu | 50.00% | 46.25)
bpp3_1_4 10:100 Session: 5 - ($-50.00 | -0.45 sbu | 40.00% | 47.40)
bpp3_1_4 10:100 Session: 6 - ($30.00 | -0.37 sbu | 50.00% | 47.17)
bpp3_1_4 10:100 Session: 7 - ($-450.00 | -0.38 sbu | 42.86% | 47.14)
bpp3_1_4 10:100 Session: 8 - ($125.00 | -0.32 sbu | 50.00% | 47.50)
bpp3_1_4 10:100 Session: 9 - ($-550.00 | -0.35 sbu | 44.44% | 47.56)
bpp3_1_4 10:100 Session: 10 - ($-200.00 | -0.33 sbu | 40.00% | 47.90)
10 Sessions, 100 Hands

Sessions won: 40.00% (4)
Sessions won t(9)-value: -0.63
Earnings: -0.33 sbu ($-330.50)
Earnings Sd: 0.79 ($791.12)
Earnings t(9)-value: -1.25
Hands won per session: 47.90% (47.90)
Showdowns per session: 57.40% (57.40)

5.3 Changes v1

The following results were obtained by running the modified bpp3.1.4 with my original changes in it.

Modifications: Fixed sandbagging threshold of 80%.

bpp3_1_4 10:100 Session: 1 - ($835.00 | 0.83 sbu | 100.00% | 62.00)
bpp3_1_4 10:100 Session: 2 - ($-410.00 | 0.21 sbu | 50.00% | 55.00)
bpp3_1_4 10:100 Session: 3 - ($-1065.00 | -0.21 sbu | 33.33% | 51.67)
bpp3_1_4 10:100 Session: 4 - ($205.00 | -0.11 sbu | 50.00% | 53.00)
bpp3_1_4 10:100 Session: 5 - ($-2155.00 | -0.52 sbu | 40.00% | 50.80)
bpp3_1_4 10:100 Session: 6 - ($85.00 | -0.42 sbu | 50.00% | 52.00)
bpp3_1_4 10:100 Session: 7 - ($1155.00 | -0.19 sbu | 57.14% | 52.57)
bpp3_1_4 10:100 Session: 8 - ($-2345.00 | -0.46 sbu | 50.00% | 50.38)
bpp3_1_4 10:100 Session: 9 - ($-850.00 | -0.51 sbu | 44.44% | 50.00)
bpp3_1_4 10:100 Session: 10 - ($5.00 | -0.45 sbu | 50.00% | 50.00)
10 Sessions, 100 Hands
Sessions won: 50.00% (5)
Sessions won t(9)-value: 0.00
Earnings: -0.45 sbu ($-454.00)
Earnings Sd: 1.16 ($1163.67)
Earnings t(9)-value: -1.17
Hands won per session: 50.00% (50.00)
Showdowns per session: 59.00% (59.00)

bpp3_1_4 10:100 Session: 1 - ($-1240.00 | -1.24 sbu | 0.00% | 41.00)
bpp3_1_4 10:100 Session: 2 - ($75.00 | -0.58 sbu | 50.00% | 46.00)
bpp3_1_4 10:100 Session: 3 - ($-965.00 | -0.71 sbu | 33.33% | 45.33)
bpp3_1_4 10:100 Session: 4 - ($-795.00 | -0.73 sbu | 25.00% | 46.75)
bpp3_1_4 10:100 Session: 5 - ($-110.00 | -0.61 sbu | 20.00% | 47.40)

bpp3_1_4 10:100 Session: 6 - ($200.00 | -0.47 sbu | 33.33% | 47.50)
bpp3_1_4 10:100 Session: 7 - ($170.00 | -0.38 sbu | 42.86% | 48.00)
bpp3_1_4 10:100 Session: 8 - ($65.00 | -0.33 sbu | 50.00% | 48.13)
bpp3_1_4 10:100 Session: 9 - ($-955.00 | -0.40 sbu | 44.44% | 47.89)
bpp3_1_4 10:100 Session: 10 - ($460.00 | -0.31 sbu | 50.00% | 48.70)
10 Sessions, 100 Hands
Sessions won: 50.00% (5)
Sessions won t(9)-value: 0.00
Earnings: -0.31 sbu ($-309.50)
Earnings Sd: 0.61 ($610.78)
Earnings t(9)-value: -1.52
Hands won per session: 48.70% (48.70)
Showdowns per session: 57.20% (57.20)

5.4 Changes v2

The following results were obtained by running the modified bpp3.1.4 with a dynamic sandbagging threshold.

Modifications: Sandbagging threshold of 60% with an increase of 10% per round.

bpp3_1_4 10:100 Session: 1 - ($845.00 | 0.84 sbu | 100.00% | 56.00)
bpp3_1_4 10:100 Session: 2 - ($300.00 | 0.57 sbu | 100.00% | 55.50)
bpp3_1_4 10:100 Session: 3 - ($-30.00 | 0.37 sbu | 66.67% | 53.00)
bpp3_1_4 10:100 Session: 4 - ($-225.00 | 0.22 sbu | 50.00% | 51.00)
bpp3_1_4 10:100 Session: 5 - ($220.00 | 0.22 sbu | 60.00% | 51.40)
bpp3_1_4 10:100 Session: 6 - ($-470.00 | 0.11 sbu | 50.00% | 50.50)
bpp3_1_4 10:100 Session: 7 - ($-155.00 | 0.07 sbu | 42.86% | 49.86)
bpp3_1_4 10:100 Session: 8 - ($-535.00 | -0.01 sbu | 37.50% | 49.63)
bpp3_1_4 10:100 Session: 9 - ($-1185.00 | -0.14 sbu | 33.33% | 49.33)
bpp3_1_4 10:100 Session: 10 - ($-710.00 | -0.19 sbu | 30.00% | 49.90)
10 Sessions, 100 Hands
Sessions won: 30.00% (3)
Sessions won t(9)-value: -1.26
Earnings: -0.19 sbu ($-194.50)
Earnings Sd: 0.57 ($573.98)
Earnings t(9)-value: -1.02
Hands won per session: 49.90% (49.90)
Showdowns per session: 57.30% (57.30)

bpp3_1_4 10:100 Session: 1 - ($760.00 | 0.76 sbu | 100.00% | 44.00)
bpp3_1_4 10:100 Session: 2 - ($415.00 | 0.59 sbu | 100.00% | 48.50)
bpp3_1_4 10:100 Session: 3 - ($-275.00 | 0.30 sbu | 66.67% | 46.67)
bpp3_1_4 10:100 Session: 4 - ($-1695.00 | -0.20 sbu | 50.00% | 47.25)
bpp3_1_4 10:100 Session: 5 - ($-310.00 | -0.22 sbu | 40.00% | 47.80)
bpp3_1_4 10:100 Session: 6 - ($-1795.00 | -0.48 sbu | 33.33% | 46.17)
bpp3_1_4 10:100 Session: 7 - ($-275.00 | -0.45 sbu | 28.57% | 47.71)
bpp3_1_4 10:100 Session: 8 - ($120.00 | -0.38 sbu | 37.50% | 48.13)
bpp3_1_4 10:100 Session: 9 - ($535.00 | -0.28 sbu | 44.44% | 48.56)
bpp3_1_4 10:100 Session: 10 - ($100.00 | -0.24 sbu | 50.00% | 48.40)
10 Sessions, 100 Hands
Sessions won: 50.00% (5)
Sessions won t(9)-value: 0.00
Earnings: -0.24 sbu ($-242.00)
Earnings Sd: 0.87 ($870.33)
Earnings t(9)-value: -0.83
Hands won per session: 48.40% (48.40)
Showdowns per session: 59.90% (59.90)

5.5 Changes v3

The following results were obtained by running the modified bpp3.1.4 with an adjusted dynamic sandbagging threshold.

Modifications: Sandbagging threshold of 65% with an increase of 7% per round.

bpp3_1_4 10:100 Session: 1 - ($165.00 | 0.16 sbu | 100.00% | 50.00)
bpp3_1_4 10:100 Session: 2 - ($-525.00 | -0.18 sbu | 50.00% | 47.50)
bpp3_1_4 10:100 Session: 3 - ($70.00 | -0.10 sbu | 66.67% | 48.33)
bpp3_1_4 10:100 Session: 4 - ($880.00 | 0.15 sbu | 75.00% | 49.75)
bpp3_1_4 10:100 Session: 5 - ($-1385.00 | -0.16 sbu | 60.00% | 48.20)
bpp3_1_4 10:100 Session: 6 - ($1185.00 | 0.07 sbu | 66.67% | 49.50)
bpp3_1_4 10:100 Session: 7 - ($-325.00 | 0.01 sbu | 57.14% | 48.86)
bpp3_1_4 10:100 Session: 8 - ($-760.00 | -0.09 sbu | 50.00% | 48.88)
bpp3_1_4 10:100 Session: 9 - ($435.00 | -0.03 sbu | 55.56% | 49.00)
bpp3_1_4 10:100 Session: 10 - ($740.00 | 0.05 sbu | 60.00% | 48.80)
10 Sessions, 100 Hands
Sessions won: 60.00% (6)
Sessions won t(9)-value: 0.63
Earnings: 0.05 sbu ($48.00)

Earnings Sd: 0.80 ($802.88)
Earnings t(9)-value: 0.18
Hands won per session: 48.80% (48.80)
Showdowns per session: 54.20% (54.20)

bpp3_1_4 10:100 Session: 1 - ($-825.00 | -0.82 sbu | 0.00% | 42.00)
bpp3_1_4 10:100 Session: 2 - ($295.00 | -0.27 sbu | 50.00% | 45.00)
bpp3_1_4 10:100 Session: 3 - ($-1080.00 | -0.54 sbu | 33.33% | 45.00)
bpp3_1_4 10:100 Session: 4 - ($1145.00 | -0.12 sbu | 50.00% | 48.25)
bpp3_1_4 10:100 Session: 5 - ($-680.00 | -0.23 sbu | 40.00% | 46.60)
bpp3_1_4 10:100 Session: 6 - ($760.00 | -0.06 sbu | 50.00% | 47.50)
bpp3_1_4 10:100 Session: 7 - ($215.00 | -0.02 sbu | 57.14% | 48.00)
bpp3_1_4 10:100 Session: 8 - ($1145.00 | 0.12 sbu | 62.50% | 49.13)
bpp3_1_4 10:100 Session: 9 - ($-375.00 | 0.07 sbu | 55.56% | 48.56)
bpp3_1_4 10:100 Session: 10 - ($-860.00 | -0.03 sbu | 50.00% | 48.30)
10 Sessions, 100 Hands
Sessions won: 50.00% (5)
Sessions won t(9)-value: 0.00
Earnings: -0.03 sbu ($-26.00)
Earnings Sd: 0.85 ($850.87)
Earnings t(9)-value: -0.09
Hands won per session: 48.30% (48.30)
Showdowns per session: 56.80% (56.80)

6 Future Development

As this was a short project, there were many possible ideas that were considered but not implemented. One that I mentioned in my talk was the check-raise technique. Check-raising is an intimidation technique that can also be used as a means of boosting a potential pot by checking first, giving your opponents, most commonly more than one opponent, the chance to bet before you make your move. Assuming that they do bet, you would then raise them, forcing them to either fold, giving you their bet, or call your raise. This is often a good strategy when you have a good hand and are betting early in a larger game of around six or more players. It can still be effective in smaller games, but in larger games there is a greater chance that at least one opponent will bet and thus force others to call them.
Another possible modification of BPP would be to allow it to play a no-limit game, or even a pot-limit game. Both game types permit variable sized bets, but pot-limit restricts the

maximum bet size to that or the current pot, a more reasonable adjustment to BPP than no-limit. A no-limit game has only one restriction on bet sizes, and that's the amount that a player has left to bet. The problem with implementing this game type is that BPP currently plays with no restriction on the amount of money it can play with. Both BPP and its opponent start with $0 and bet with imaginary money. When one player wins, they shift to a positive balance, and the loser shifts equally as far negatively. A no-limit game where there's no limit to the player's stack would basically allow for infinite bets, which is a bit of an unrealistic situation to simulate.

7 Conclusion

A number of things can be observed in my results. First of all, because each test only involved one thousand hands, they can't be considered a great example of the performance of the program. However, longer tests of ten thousand hands or more take a great deal of time, so generating several tests for each version would take several days. I decided that I would be better off with multiple test cases rather than a single larger test case for each version.

BPP was originally fairly varied in it's results. Average result margins of $187.50, $-83.00, and $-330.50 were quite diverse in range and generally negative. However, as BPP was playing against itself, it could also be said that it had results of $-187.00, $83.00, and $330.50. In general, the results suggest that BPP has an average of approximately breaking even against itself, as should be expected.

My adjustments to BPP all involved an array of options to be implemented based on the conditions of the round. The first version I tested would call references to this array within a fixed threshold of 80%. This gave me the results $-454.00 and $-309.50. In both cases, a fairly large negative result against the original, which had a break even average result, so a definite lack of progress there.

My second version used a reduced fixed threshold of 60%, but with a variable increase of 10% per round starting after the flop. This meant that sandbagging would be more likely in later rounds, and the probability of winning will be higher then too. The combination of these should mean that sandbagging will occur in later rounds much more frequently than earlier rounds, while still being possible in earlier rounds. The initial reason for decreasing the static threshold was to make it possible to sandbag in the earlier rounds when BPP would not be likely to have a winning probability of 80%. After testing, I got results of $-194.50 and $-242.00. This is still a negative result, but an improvement on the original modifications.

My third version used a slightly higher static threshold of 65% and a lower increment of 7% per round. I decided to increase the static threshold and decrease the increment to keep the max value reached in the turn at approximately the same value as the last version, around 80%. This was to ensure that sandbagging was never too highly imple-

mented. The results received were favorable. Average gains of $-26.00 and $48.00 showed significant improvement over the previous versions, and even an apparent minor improvement on the original.

As initially stated, the results used in this report are only one thousand hands per test case. As such, they aren't a full demonstration of the effectiveness of my changes. They do however, to their limited extent, suggest that my modifications provide for a degree of improvement over the original BPP.

Over the course of this project, I have learned a great deal about poker. I've discovered terminology that I never knew, techniques and strategies that I have found insightful and will endeavor to use in the future, and formulae and statistics that have helped me understand more about the theory of poker. As for BPP, it has taught me a great deal about reading and modifying existing code, as well as a great deal about AI. The Bayesian networks involved have also taught me a lot about probability based decision making. All tolled, I've gained a wealth of information from this project.

A Glossary

AI: Artificial Intelligence. A program that takes information about its environment and uses it to logically deduce an appropriate action.

Chip: A token used in poker to represent money.

Draw: When a player 'draws' they rely upon the next card dealt to improve their hand. In the case of Texas Hold'em, that card is the next community card.

Hand: 1. The best five card combination that a player can make with the current cards available to them.
2. A single game that begins with the payment of blinds and finishing when a player wins the pot.

Out: A card that will complete a player's hand. For example, if a player has four hearts in their hand then any remaining heart will complete the player's flush, making them outs.

Outdraw: When a player is outdrawn it means that their hand has been beaten by a draw. Using 'a' and 'b' as two Texas Hold'em players, player 'a' has a straight on the turn, but player 'b' has four hearts. On the river, a heart is drawn giving player 'b' a flush. In this case, player 'b' has outdrawn player 'a'.

Pot: The collection of all bets for a particular hand. When a hand is over, the winning player takes all of the money that has been placed in the pot.

SBU: Small Betting Units. A betting unit is the amount that can be bet in a limit game of poker. In the case of BPP, a small betting unit is $10.

Session: A number of hands of poker played one after another.

Stack: The amount of money, normally represented by poker chips, that a player has to bet with.

References

[1] Howard Lederer. The Other Danger in Slow Playing. http://www.pokertipsfromthepros.com/lesson58.htm.

[2] www.cardschat.com. Pot Odds & Implied Odds. http://www.cardschat.com/poker-odds-pot-odds-implied-odds.php.

[3] www.norsys.com. Norsys Software Corp. www.norsys.com/netica.html.

[4] www.pokerpaydayz.com. Poker Theory. http://www.pokerpaydayz.com/index.php/poker-theory.