

Best Response to Tight and Loose Opponents in the Borel and von Neumann Poker Models

Casey Warmbrand

May 23, 2006

Abstract

This paper will present two famous poker models, developed by Borel and von Neumann. An analysis of each model will be presented and an equilibrium solution found. Throughout the paper we will show facts about optimal play, for these models, that mirror very general "rules of thumb" that most skilled poker players would argue are true in any poker variant. Analysis will include what to do against opponents who are behaving irrationally, as well as how one should adjust their play with varying bet sizes.

1 Introduction and Preliminaries

The study of poker models dates all the way back to 1938, when Borel introduced his model, 'la relance' in "Applications aux Jeux des Hazard" [4]. In 1944, a similar yet slightly more complex model was introduced by von Neumann in "Theory of Games and Economic Behavior." [7]

Both models greatly simplify the poker that you or I tend to play in a casino or home game. As any game theorist knows, it is quite difficult to analyze a game as complex as poker, thus the need for simplification. Both poker games are two-player zero-sum games, where we assume the hands dealt to each player are uniformly distributed and independent. The assumption of independence is not necessarily a fair assumption. In any ordinary poker variant with a finite deck of cards, let's say a standard 52 card deck, we do not have independence of hands. For example, if I am holding 4 Aces, then it is impossible for anyone else to have an Ace. This fact decreases the expected strength of any opponent's hand. However, in a game such as Texas Hold-em, where each player's starting hand is only 2 cards, the dependence between hands is rather small. Another important

thing to note is that neither of these models are symmetric. Player I's actions will differ from Player II's, this will lead to equilibria in which one player will have a strictly positive expected payoff.

Both the Borel and von Neumann models will be presented completely, with their Nash equilibrium solutions. We will see that the Borel model favors Player II, and the von Neumann model favors Player I. Out of the solution to each of the games, we will be able to construct best-response rules for the player that the game favors. In order to solve the games for their equilibria, we will need to assume that both players are rational. However, as is often the case either amongst friends in a home game or against fellow gamblers in a casino, irrational play is frequently observed. Thus, for the sake of a fruitful analysis, it will be beneficial to later relax the assumption of rationality for one of the two players by allowing them to deviate from equilibrium. Then using the best-response rules, we will be able to analyze how a rational player should react to such an "off-equilibrium" opponent. For both models, we will look at two different types of off-equilibrium playing opponents; one who plays too few hands or a *Tight* player, and one who plays too many hands, a *Loose* player, and discuss strategies against such opponents.

2 The Borel Model

2.1 'La Relance'

We will call the two players I and II. Player I is dealt a random hand X uniformly distributed on $[0, 1]$. The uniform distribution implies that $\text{Prob}(X < k) = k$ and $\text{Prob}(X = k) = 0$ for all $k \in [0, 1]$. Similarly, player II receives a random hand, Y , independent of X , also uniformly distributed on $[0, 1]$. Each player posts an ante of 1 unit. Player I, after receiving his hand, has the option either to fold his hand or bet some fixed amount $\beta \geq 0$. If player I folds, player II wins the antes; if player I has bet, then player II has the option of folding or calling the bet. If the bet is called, a showdown occurs, and the higher hand wins the pot. We may sometimes refer to I as the *bettor* and to II as the *caller*.

The use of a continuous range of hands is something that is not encountered in any conventional poker variant. However, the number of possible hands that can be made in most poker games is quite large. For example in 5-card stud poker the number of hands is $\binom{52}{5} = 2,598,960$, each equally likely. When the number of potential hands is very large, the probability of 2 players hands being the same becomes very small, and it is "as if" the

potential hands come from a nearly uniform distribution on $[0, 1]$.

The betting tree is presented below, where the value at the end of each branch indicates the payoff to player I. Positive payoffs denote a profit to I, and negative payoffs denote a loss to I; equivalently a loss to II and a profit to II, respectively. The $\pm(\beta+1)$ at the end of the branch $\{\text{bet,call}\}$, indicates that I gets a profit of $(\beta+1)$ if $X > Y$ but loses $(\beta+1)$ if $Y > X$.

[Insert Neil's Borel tree here]

Fig. 1. The betting tree for La Relance

2.2 Rationalizable Strategies

Next we consider what sort of strategies each player might want to use. A typical strategy for I should be: Bet if $X > m$ and fold otherwise, for some fixed $m \in [0, 1]$. At the decision node for I in the tree above, the only information available to him is his hand strength. Thus, if it is desirable to bet a hand $X = m_0$, then it must also be desirable to bet any hand stronger than this, $X > m_0$. So player I should only want to use a strategy of this form, referred to as a threshold strategy. A strategy for II should be: Call if $Y > n$ and fold otherwise, for some fixed $n \in [0, 1]$. This is also a threshold strategy. Again, this is the only rationalizable choice, because at II's decision node he only knows his hand strength and the fact that I bet. Thus, if calling with a hand, $Y = n_0$ is desirable, calling with any stronger hand, $Y > n_0$ is too.

Note that neither player gains anything from using a mixed strategy. This is because the randomization of the deal is enough to keep the opponent guessing. From now on, we will refer to I's strategy simply as m and II's strategy as n .

If we assume II knows that I is using the strategy: bet if $X > m$, then clearly if his hand Y is less than m , he should not call a bet made by I. Thus II will always want to choose $n \geq m$. This assumption is not necessarily a good one. Typically, an opponent's strategy will not be known when the game is first played (at least I hopes that II doesn't know his strategy). However in repeated play of the game, assuming I plays the same fixed strategy in each repetition, statistical analysis of the betting patterns of I could be used to make an approximate guess at his strategy, m . The longer the game is played the better this approximation would become.

Assuming each player plays one of these threshold strategies, and that this relation, $n > m$ holds, we may draw a general picture in order to

determine the probability of each of the 4 outcomes at the terminal nodes of the tree in Figure 1.

2.3 The Payoff Function

[insert unit square here]

Fig. 2. Outcome probabilities for La Relance

The above diagram's 4 regions, correspond to the 4 possible outcomes of the game. The area of each region is exactly the probability of that outcome occurring, because X and Y are independent and uniformly distributed on $[0, 1]$. Each outcome, along with its probability and payoff to I, are given in the following table.

Outcome	Probability	Payoff
I folds	$p_1 = m$	$\pi_1 = -1$
I bets, II calls, I wins	$p_2 = \frac{(1-n)^2}{2}$	$\pi_2 = (\beta + 1)$
I bets, II calls, II wins	$p_3 = \frac{(1-n)^2}{2} + (n - m)(1 - n)$	$\pi_3 = -(\beta + 1)$
I bets, II folds	$p_4 = (1 - m)n$	$\pi_4 = 1$

We have that $\sum_{i=1}^4 p_i = 1$, as one of these four outcomes must occur. From the table we can easily find the expected payoff to I.

$$\pi_\beta(m, n) = \sum_{i=1}^4 p_i \pi_i = -\beta n + \beta n^2 + n^2 + \beta m - \beta m n - 2m n. \quad (1)$$

2.4 A Best-Response Rule for the Caller

Player I wishes to maximize this function, while II would like to minimize it. Taking player I's strategy, m to be fixed, and allowing II to choose any strategy, n , we look for an n that minimizes π . Setting $\frac{\partial \pi}{\partial n} = 0$ and solving for n , leads us to a best-response rule for player II, which depends on I's strategy, m , and β , the fixed bet size.

$$n_\beta(m) = \frac{2m + (1 + m)\beta}{2(\beta + 1)} = \frac{\beta + (\beta + 2)m}{2(\beta + 1)} \quad (2)$$

Viewing (2) as a function of β , with m fixed, we see that n is an increasing rational function of β , with horizontal asymptote $\frac{1+m}{2}$. This implies the next two lemmas.

Lemma 2.1. *Regardless of the bet size, to optimize his profit, II should always call, when Y is bigger than the average of 1 and m .*

Lemma 2.2. *For the Borel model of poker, given a fixed strategy for player I, a higher bet size leads player II to call less often, where as a smaller bet size leads him to call more often.*

In some sense, lemma 2.2 allows us to view β as a risk factor. When β is large, there is a risk of a large loss involved in calling a bet. Fearing this, II is less likely to call when β is large.

Looking at II's best-response rule in equation (2) as a function of m , with β fixed, we see that n is a linear function of m with slope $\frac{\beta+2}{2(\beta+1)} > 0$. Thus n will increase with m . Let μ and ν denote the equilibrium strategies for I and II respectively (given later in Theorem 2.4). For an arbitrary fixed strategy, m for I, we denote the best-response by II as n . First consider player I types who don't play enough hands. These players, denoted I_\bullet , use a strategy: Bet if $X > m$, for some $m > \mu$. Since $n(m)$ increases with m , we know that a rational player II, who is playing against an I_\bullet type, should use a strategy: Call if $Y > n$, for $n = \frac{\beta}{2(\beta+1)} + \frac{\beta+2}{2(\beta+1)}m > \frac{\beta}{2(\beta+1)} + \frac{\beta+2}{2(\beta+1)}\mu = \nu$. We can think of a player who plays "too few hands", as a *tight* player, and we now see that the correct response is to call less hands. Similarly, for I_\bullet types who use an $m < \mu$, i.e. play "too many hands" or are *loose*, the best-response rule tells player II to use $n < \nu$. We interpret these two facts in the following lemma.

Lemma 2.3. *For the Borel model of poker, the best-response to a tight bettor is to call less, and the best-response to a loose bettor is call more.*

The extremes here are $m = 1$ and $m = 0$ which give us $n = 1$ and $n = \frac{\beta}{2(\beta+1)}$, respectively. It is interesting to note that even when facing the loosest possible I_\bullet type, who *always* bets, the optimal strategy for II is to call if and only if $Y > \frac{\beta}{2(\beta+1)}$. For example, in the case $\beta = 1$, if I always bets, II's best-response is to only call $\frac{3}{4}$ of the time.

2.5 A Best-Response Rule for the Bettor

Unfortunately, differentiating the profit function in (1) with respect to m , does not lead us to a best-response rule for player I. It does however, give something of use. Taking $\frac{\partial \pi}{\partial m} = 0$, we see that $n = \frac{\beta}{\beta+2}$. This condition on II's strategy, as we will see shortly, is what will make I indifferent to betting on any hand $X < n$.

In order to attain a best-response rule for I we will have to use another technique. Now we take II's strategy, n to be fixed, and allow I to choose any strategy, m . For the purposes of this subsection, it will be convenient to view the antes as a sunk cost to both players. This changes to game from zero-sum to constant-sum=2. We next consider the options available to I at his decision node. He can either bet or fold. A fold will yield a payoff of 0, since the ante is now a sunk cost. We construct a new kind of payoff function that computes the expected payoff to I from betting a hand $X = x$.

$$\pi_{\beta}^I(x) = \begin{cases} -\beta(1-n) + 2n & \text{for } x \leq n \\ -\beta(1-x) + (\beta+2)(x-n) + 2n & \text{for } x > n \end{cases} \quad (3)$$

First note that this is a non-decreasing, continuous, piece-wise linear function of x . Because this function is non-decreasing and continuous, we see that I will in deed wish to use a threshold strategy against any threshold strategy that II uses.

We now only need to find the values of x , such that $\pi_{\beta}^I(x) > 0$, where a bet will yield a higher expected payoff than a fold. Looking at the first piece of the function above, we see that it is 0 precisely when $n = \beta\beta + 2$; this is the same indifference condition we obtained before from $\frac{\partial \pi}{\partial m} = 0$. When $n > \frac{\beta}{\beta+2}$, this piece of the function is positive, and when $n < \frac{\beta}{\beta+2}$, it is negative. So immediately we see that if n is sufficiently large, then player I should always bet. If $n = \frac{\beta}{\beta+2}$, then the second piece of the function above will be positive when $x \geq \frac{\beta+\beta n}{2(\beta+1)}$. We state I's best-response rule as follows: For any $k \in [0, n]$,

$$m_{\beta}(n) = \begin{cases} 0 & \text{if } n > \frac{\beta}{\beta+2} \\ k & \text{if } n = \frac{\beta}{\beta+2} \\ \frac{\beta(1+n)}{2(\beta+1)} & \text{if } n < \frac{\beta}{\beta+2} \end{cases} \quad (4)$$

We interpret a player II using a "large" n , corresponding to calling very few hands, as a "tight" player, and similarly ones using a "small" n as a "loose" player. We see from the best-response rule for I, that against any sort of tight opponent, no matter how tight, it is optimal to always bet. Against loose opponents, the frequency of a bet depends on how small n is. In particular for $n = 0$ (always calling), it is best for I to only bet if $X > \frac{\beta}{2(\beta+1)}$.

Lemma 2.4. *For the Borel model of poker, the best-response to a tight caller is to bet as much as possible and the best-response to a loose caller is to bet less.*

Taking one more look at the indifference condition, $n = \frac{\beta}{\beta+2}$, we find another interpretation. The probability that player II folds, $\text{Prob}(\text{fold})=n$. With the notion of sunk costs, β is the amount II will lose if he calls and loses, where as $\beta+2$ is the amount he will win. Thus, if the probability that II folds is equal to the ratio of expected losses to expected winnings from a showdown, then I is indifferent to betting or folding with any hand $X < n$.

2.6 The Borel Equilibrium

Rather than simultaneously solving the two best-response rules, we can use substitution and calculus to find the unique equilibrium. Taking both players to be rational, we assume II will always use the best-response rule (2). We may now rewrite the p_i 's, probabilities for each of the outcomes and π from (1), in terms of just m and β .

$$p_2 = \frac{((1-m)(\beta+2))^2}{8(\beta+1)^2} \quad p_3 = \frac{(1-m)^2(\beta+2)(3\beta+2)}{8(\beta+1)^2} \quad p_4 = \frac{(1-m)(\beta+\beta m+2m)}{2(\beta+1)}$$

$$\pi_\beta(m) = -\frac{(\beta(1-m))^2 + 4m^2(\beta+1)}{4(\beta+1)} \quad (5)$$

Since this is II's **best**-response, this is the worst that player I could possibly do with any strategy m and fixed bet size β . Notice that the profit to player I is negative. So with rational players, the Borel model favors player II. We can also find I's equilibrium strategy by finding m to maximize $\pi_\beta(m)$ in (5). Setting $\frac{\partial \pi}{\partial m} = 0$ and solving for m , gives I's equilibrium strategy in terms of β . Substituting this m into (2) and (5), gives II's equilibrium strategy and the payoff of the game in equilibrium. The results are given in the following theorem.

Theorem 2.5. *The unique optimal strategy for player II is to call if $Y > \nu$ and fold otherwise. Player I's optimal strategy is to bet if $X > \mu = \nu^2$. The expected outcome is that II wins ν^2 from I. Where $\nu = \frac{\beta}{\beta+2}$.*

We see that in equilibrium II will make I indifferent to betting when $X < \nu$, as seen from I's best-response rule, (4). In order for this strategy by II to be a best-response, it is necessary that I only bets when $X > \mu = \nu^2$. This gives us the uniqueness of the equilibrium.

Looking back at (5), we notice that the payoff is quadratic in m , with critical point at I's equilibrium strategy, μ . This means that the payoff II gets from his best-response doesn't depend on whether I is considered "loose" or "tight," only on how far from equilibrium I is playing.

Since this game favors player II, if player I is allowed to select the bet size β before the hands are dealt, he should choose it as small as possible.

3 The von Neumann Model

3.1 Same Game, With a Twist

The von Neumann model differs from the Borel model in a very small, yet extremely significant way. The set up for the game will be the same, with X and Y being I and II's hands respectively, independent and uniformly distributed as before. The only difference in this game is in I's options after the deal. Instead of folding if he does not wish to bet the fixed amount β , he now checks his hand, and the showdown occurs immediately with the winner receiving the antes. Before, in Borel's model, the options of both players were quite restricted. There was no opportunity for II to bet or for I to check. While II is still left without any opportunities to bet, I is now given the option of checking. Even with this constraint for II, this model seems to represent true poker much better than the Borel model. In most poker variants, players often have the option to check, a strategy that (when available) weakly dominates folding. This is exactly the case here in the von Neumann model, which is why we say that if I doesn't bet he automatically checks. Also, and probably more importantly, the equilibrium solution to this game entails bluffing by I. Bluffing is a key element in many poker games, so a model like Borel's that doesn't include such a feature leaves a lot to be desired.

This added option of check for I, makes the von Neumann model a better game for I. The advantage to I in this game over La Relance will manifest itself in the payoffs to I here, compared with the payoffs to I from the Borel model. In fact, as we will soon see, this model favors I. As a result, much of the analysis for this model will appear to be from the opposite player's perspective than as before.

The betting tree for the von Neumann model is presented below. Notice that the only difference is that the right branch at I's decision node is now labeled "check" rather than "fold" and that the payoff at this terminal node is now ± 1 rather than just -1 . This means that I wins 1 when he checks and $X > Y$, and loses 1 when he checks and $Y > X$.

[insert Neil's vN tree here]

Fig 3. The betting tree for von Neumann's model

It should be noted that von Neumann’s original model in [7] was presented a little differently, with a small and large bet. However, taking his small bet to be the ante, which here we set fixed at 1 (corresponding to a rescaling or choice of currency), and his large bet to be β , the model here is equivalent.

3.2 Rationalizable Strategies

The new option to player I of “check” rather than fold, creates the possibility that a threshold strategy is no longer the only rationalizable strategy for I. However, when player II is at his decision node, after I has bet, the situation is exactly as before. This means that again II’s only rationalizable strategies are threshold strategies of the form: Call if $Y > c$ and fold otherwise, for some fixed $c \in [0, 1]$. Let’s consider how player I might improve upon a threshold strategy of the form: Bet if $X > b$ and check otherwise, for some fixed $b \in [0, 1]$. For certain strategies, c for II and bet sizes β , player I may be able to gain additional value from betting especially weak hands. For this reason the rationalizable strategies for I are of the form: Check if $a < X < b$ and bet otherwise, for fixed $a \leq 1, b \geq 0$, with $a < b$.

Using his strategy above, I makes 2 different kinds of bets, one when $X < a$, betting with a “weak” hand which we call a *bluff*, and one when $X > b$, betting with a “strong” hand, an *honest bet*. The interval (a, b) will be referred to as the check-range. For appropriate values of c and β , it will be possible for a to realize values less than 0. This will correspond to player I never bluffing, and playing a threshold strategy, similar one used in the Borel model. We will denote II’s strategy simply by c , and I’s strategy by the pair $\{a, b\}$.

If we assume that I knows II’s strategy is c , it would not be wise to bet if X is close to c , because if II calls, I is likely to lose. This means that the check-range should always contain c . Thus I will always want $a \leq c \leq b$. Note that if this is the case, whenever I’s bluff is called by II, player I will surely lose the pot. Assuming each player only uses a rationalizable strategy, then using this relation, $a \leq c \leq b$ we can now construct a diagram similar to Figure 2.

3.3 Another Payoff Function

[insert unit square for vN here]

Fig. 4. Outcome probabilities for von Neumann’s model

The above diagram now has 5 regions, corresponding to the 5 possible outcomes of this game. Again, the area of each region is its probability of occurring. The outcomes, probabilities, and associated payoffs are given in the following table.

Outcome	Probability	Payoff
I checks, I wins	$q_1 = \frac{(b-a)^2}{2} + a(b-a)$	$\pi_1 = 1$
I checks, II wins	$q_2 = \frac{(b-a)^2}{2} + (1-b)(b-a)$	$\pi_2 = -1$
I bets, II calls, I wins	$q_3 = \frac{(1-b)^2}{2} + (1-b)(b-c)$	$\pi_3 = (\beta + 1)$
I bets, II calls, II wins	$q_4 = \frac{(1-b)^2}{2} + a(1-c)$	$\pi_4 = -(\beta + 1)$
I bets, II folds	$q_5 = c(1 - (b-a))$	$\pi_5 = 1$

We again have that $\sum_{i=1}^5 q_i = 1$. We now compute the expected payoff to player I.

$$\pi_\beta(a, b, c) = \sum_{i=1}^5 q_i \pi_i = \beta[c(a+b-1) + (b-b^2-a)] + a(2c-a) \quad (6)$$

3.4 The Bettor's Best-Response

As before, I wishes to maximize this function, while II would like to minimize it. This time we find I's best-response strategy, $\{a, b\}$ first. Setting $\frac{\partial \pi}{\partial a} = 0$ and solving for a , and setting $\frac{\partial \pi}{\partial b} = 0$ and solving for b , leads us to two best-response rules, both for I. One of them determines a as a function of c and β , while the other, determines b solely as a function of c .

$$a_\beta(c) = \frac{2c - (1-c)\beta}{2} = \frac{-\beta + (\beta+2)c}{2} \quad (7)$$

$$b(c) = \frac{1+c}{2} \quad (8)$$

The fact that the best-response rule (8) doesn't depend on β , leads to another lemma, similar to Lemma 2.1.

Lemma 3.1. *Regardless of the bet size, to optimize his profit, I should bet honestly if and only if X is bigger than the average of 1 and c .*

Examining (7) closer, we see that a depends linearly on β , with negative slope, and so we have a lemma, analogous to Lemma 2.2.

Lemma 3.2. *In the von Neumann model, given a fixed strategy for player II, a higher bet size leads Player I to bluff less often, and a smaller bet size leads player I to bluff more often..*

As before, we may again interpret β as a measurement of risk. When β is large, being caught in a bluff by II's call is very costly to I. And so he is less likely to bluff when β is large.

Looking at the best-response rules in (7) and (8), we see that a is a linear function of c with slope $\frac{\beta+2}{2}$ and that b is also a linear function of c with slope $\frac{1}{2}$. So both a and b will increase with c . Let a_e, b_e , and c_e denote the equilibrium strategies (given later in Theorem 3.6). Consider II types who don't play enough hands. These again are the ones we classify as "tight", and refer to them as II^\bullet with strategies of the form call if $Y > c$, for some $c > c_e$. Since both a and b increase with c , a rational player I, playing against an irrational II^\bullet type, should use a strategy: Check if $a < X < b$ and fold otherwise, for some $a > a_e$ and $b > b_e$. Similarly, for the "loose" players, denoted II_\bullet , who use a $c < c_e$ and play too many hands, the best-response rules tell player I to use $a < a_e$ and $b < b_e$. We can immediately see the effect that an off-equilibrium value of c has on the hand values for which it is optimal for player I to bluff, which is stated in the next lemma.

Lemma 3.3. *For the von Neumann model of poker, the best-response to a tight caller is to bluff more, where as against a loose caller it is best to bluff less.*

Let us now consider the effects of off-equilibrium c values on the length of the check-range, $b - a$. For an arbitrary c , the length of the check-range is $\frac{(\beta+1)(1-c)}{2}$. This is a linear function of c , with slope $-\frac{\beta+1}{2}$. Thus, the length of the check-range decreases with c . Since whenever player I doesn't check, he bets, the length of the bet-range is $1 - (b - a)$. This means that the length of the betting range is actually increasing with c . The probability that I bets on any given hand is exactly the length of the bet-range, so we may interpret this as a measure of the total betting that player I does. The next lemma follows immediately.

Lemma 3.4. *For the von Neumann model of poker, the best-response to a tight caller is more total betting. Similarly, a best-response to a loose caller is less total betting.*

Finally we consider what happens with the extreme values of c . For $c = 1$, we obtain $a = 1, b = 1$ from (7) and (8). When facing the tightest possible opponent, who never calls a hand, the best-response is to always

bluff. For $c = 0$, we see that $a = -\frac{\beta}{2}, b = \frac{1}{2}$. Here the negative value of a , simply means that it is never wise to bluff against someone who will always call. The best-response against the loosest possible opponent who always calls, is to bet if and only if $X > \frac{1}{2}$.

3.5 The Caller's Best-Response

Unlike the Borel model, differentiating equation (6) with respect to c , does give us something interesting. Setting $\frac{\partial \pi}{\partial c} = 0$ and solving leads us to:

$$(\beta + 2)a = \beta(1 - b) \quad (9)$$

While this is not a best-response rule for II, since there is no c , it is the condition for II's indifference, as described in [1]. What this means is that if this equation is satisfied by whatever strategy player I is using, then player II will be indifferent between calling and folding for all hands $Y \in [a, b]$. In order to attain the desired best-response rule we will use the same technique that we used to find the best-response rule for I in the Borel model. Again in this subsection it is convenient to view the antes as a sunk cost to both players. We assume that I is using a fixed strategy $\{a, b\}$, and allow II to choose any strategy c . Player II has two options when it is his turn to act at his decision node. He may call or fold. A fold yields a payoff of 0, since the antes are already a sunk cost. So we construct a payoff function that computes the expected payoff **to II** from calling a hand $Y = y$.

$$\pi_{\beta}^{II}(y) = \begin{cases} (\beta + 2)\frac{y}{1-b+a} - (\beta)\frac{1-b+a-y}{1-b+a} & \text{for } y \leq a \\ (\beta + 2)\frac{a}{1-b+a} - (\beta)\frac{1-b}{1-b+a} & \text{for } y \in (a, b) \\ (\beta + 2)\frac{a+y-b}{1-b+a} - (\beta)\frac{1-y}{1-b+a} & \text{for } y \geq b \end{cases} \quad (10)$$

Note that this function is a non-decreasing, continuous, piece-wise linear function of y . The fact that it is non-decreasing and continuous reinforces II's desire to play a threshold strategy in response to any threshold strategy that player I uses.

We next look for values of y , such that $\pi_{\beta}^{II}(y) > 0$, where a call will yield a higher expected payoff than a fold. If we consider the middle piece first, we see that this is 0 precisely when equation (9) is satisfied. This shows that the equation $\frac{a}{1-b} = \frac{\beta}{\beta+2}$ is both necessary and sufficient for player II's indifference between calling and folding on hands $Y \in [a, b]$. If the ratio $\frac{a}{1-b}$ is less than $\frac{\beta}{\beta+2}$, then the middle piece of the function is negative, and

so we must look for stronger hands, or larger y 's, to call with. So we find the y 's to make the third piece of the function positive. If the ratio $\frac{a}{1-b}$ is more than $\frac{\beta}{\beta+2}$, then the middle piece of the function is positive, and we might also gain from calling with weaker hands, or smaller y 's. So we find the y 's which make the first piece of the function positive. We state II's best-response rule as follows: For any $k \in [a, b]$,

$$c_\beta(a, b) = \begin{cases} \frac{\beta(1-b+a)}{2(\beta+1)} & \text{if } \frac{a}{1-b} > \frac{\beta}{\beta+2} \\ k & \text{if } \frac{a}{1-b} = \frac{\beta}{\beta+2} \\ \frac{\beta+\beta(b-a)}{2(\beta+1)} & \text{if } \frac{a}{1-b} < \frac{\beta}{\beta+2} \end{cases} \quad (11)$$

We already know that when $\frac{a}{1-b}$ is large, $c_\beta(a, b) < a$, and when it is small $c_\beta(a, b) > b$ from the way we found (11) from (10). We notice that the when I uses the strategy $\{a, b\}$, the probability that he bluffs is a , and the probability he bets honestly is $1 - b$. So when $\frac{a}{1-b}$ is large, I is bluffing in a high proportion to his honest bets. This could be considered "loose" play. II's best-response rule tells him to use a $c < a$, and so II calls frequently. Similarly, when $\frac{a}{1-b}$ is small, the proportion of bluffs to honest bets is small, potentially "tight" play. Against this kind of $\{a, b\}$, (11) tells II to use $c > b$, in other words to fold frequently.

Lemma 3.5. *For the von Neumann model of poker, the best-response to a tight bettor is to fold more, and the best-response to a loose bettor is to call more.*

We can interpret the indifference condition from (9) in another way. Given a fixed strategy $\{a, b\}$ for I, the probability he bluffs, $\text{Prob}(\text{bluff})=a$ and the probability he bets honestly, $\text{Prob}(\text{honest})=1 - b$. As before with sunk costs, $\frac{\beta}{\beta+2}$ is the ratio of expected losses to expected winnings from a showdown. Thus we see that if the ratio of $\text{Prob}(\text{bluff})$ to $\text{Prob}(\text{honest})$ is equal to the ratio of expected losses to expected winnings, then player II is indifferent to calling or folding for all hands $Y \in [a, b]$, and he will definitely want to fold all hands $Y < a$ and call all hands $Y > b$.

3.6 The von Neumann Equilibrium

Now taking both players to be rational, we may assume that I always uses the best-response rules (7) and (8). As before, we rewrite the probabilities of each outcome in terms of just c and β . Then recalculating the payoff in (6), leads to the following two forms of the same equation:

$$\pi_c(\beta) = \frac{4c^2 + \beta(1 - 6c + 5c^2) + \beta^2(2c + c^2)}{4} \quad (12)$$

$$\pi_\beta(c) = \frac{\beta - c(6\beta + 2\beta^2) + c^2(4 + 5\beta + \beta^2)}{4} \quad (13)$$

The reason for writing the same payoff equation in two ways is that (12) allows us to view the payoff as a function of β for a fixed c , while (13) allows us to view it as a function of c for β fixed.

This is the payoff when I uses his **best**-response, so this is the worst that II can ever expect to do with startegy c . Unlike before, it is not immediately clear whether the profit to I is positive or negative. But, assuming player II responds rationally, he will choose c to minimize $\pi_\beta(c)$ in (13). Setting $\frac{\partial \pi}{\partial c} = 0$ and solving for c , gives II's equilibrium strategy in terms of β . Substituting this into (7) and (8), gives the equilibrium strategy for I. And finally, we obtain the equilibrium payoff in terms of β . Only now do we notice the sign of the payoff. It is positive, indicating that with rational players, the von Neumann model favors player I. The next theorem summarizes the results.

Theorem 3.6. *An equilibrium strategy for I is to check if $a_e < X < b_e$ and bet otherwise, where $a_e = \frac{\beta}{(\beta+1)(\beta+4)}$ and $b_e = \frac{(\beta+2)^2-2}{(\beta+1)(\beta+4)}$. An equilibrium strategy for II is to call if $Y > c_e$ and fold otherwise, where $c_e = \frac{\beta(\beta+3)}{(\beta+1)(\beta+4)}$. The expected out come is that I wins $\frac{\beta}{(\beta+1)(\beta+4)}$ from II.*

We now see that in equilibrium, I will always choose a_e and b_e so that they satisfy indifference equation (9), thus making II indifferent between calling and folding when $Y \in [a, b]$. In order for this strategy by I to be a best-response, we need that II only calls when $Y > c_e$. This gives us uniqueness of this equilibrium.

The length of the equilibrium check-range, given by $(b_e - a_e) = \frac{\beta+2}{\beta+4}$. This is an increasing rational function of β , with horizontal asymptote 1, leading to the following corollary.

Corollary 3.7. *The larger the bet size, β , the more checking player I does in equilibrium. In particular for $\beta = 0$, player I checks $\frac{1}{2}$ the time in equilibrium. As $\beta \rightarrow \infty$, player I's check-range, $(b_e - a_e) \rightarrow 1$, he checks all the time in equilibrium.*

This corollary also suggest we might want to view the bet size as a risk measurement. When β is large, I should choose to bet less (check more)

often, than he would if β were smaller. We could even say that when the bet size is unfathomable (maybe you are betting your life), it is best to check, in other words *don't even bet at all*.

Similar to the payoff function in the Borel model, the payoff equation in (13) is quadratic in c , with critical point at II's equilibrium strategy, c_e . This again implies, that for a fixed β , the payoff that I gets from best-responding doesn't depend on whether the opponent he is facing is "loose" or "tight"; it only depends on how far off-equilibrium II is playing.

Suppose now we allow I the option of selecting the bet size, β before the hands are dealt. If we look at the equilibrium value of the game given in Theorem 3.6, and maximize it with respect to β , this will be the optimal bet size that player I should choose. Taking a derivative with respect to β , setting it equal to 0, and solving for β , yields that $\beta = 2$. Hence the optimal bet size is exactly the size of the pot.

4 Other Poker Models

Work has been done to generalize these two poker models into one single model. This is done by allowing player I three options after the deal, fold, bet β_1 , or bet β_2 , with $0 \leq \beta_1 \leq \beta_2$. Setting $\beta_1 = \beta_2$, reduces this model to the Borel model, where the only options are fold or bet a fixed amount. Setting $\beta_1 = 0$, reduces this model to the von Neumann model, since folding is now weakly dominated by bet $\beta_1 = 0$, which will always be called by II. This type of bet and call is equivalent to a check. Work on this generalized model has been done by Bellman and Blackwell in [3] and by Bellman in [2]. A betting tree similar to the ones presented here, for this more general model can be found in [5]

Another model, dubbed 'real' poker was developed by Newman in [6]. In this model, I is allowed to choose the bet size β , **after** he has looked at his hand. The idea here is that this in some way models no-limit poker. The solution for this game requires I to use a strategy similar to the von Neumann model: Check (bet 0) if $a < X < b$, bet $\beta(X) > 0$ otherwise. The value of the bet here depends on the hand strength of I. The bet size increases in both directions, meaning with hands close to the check-range (a, b) , you should bet small, and with hands far from the check range (close to $a = 0$ or $b = 1$), you should bet alot. In fact, $\beta \rightarrow \infty$ as $X \rightarrow 0$ and $X \rightarrow 1$. It truly is a "no-limit" poker game.

Some work toward removing the restriction to independent uniformly distributed hands has been done in [5]. The consider a number of differ-

ent possible distributions for the hands; independent not identically distributed, negative dependence, and the von Neumann model with general dependence.

Developing a poker model may not seem like such a daunting task, but if you wish to analyze the model and find its equilibrium, it will become exponentially difficult with the complexity of the model. This causes major problems for someone wishing to develop a model that simulates any real poker variant closely. We are optimistic that, perhaps slowly, progress will be made in the development and analysis of more advanced poker models, so that other interesting results can be shown in the future.

5 Recap of Our Findings

We have shown that for the Borel model, tight play is best-responded to with tight play, just as loose play by an opponent should be responded to with appropriately loose play. We also saw that as the fixed bet size increased, the strength of hand necessary to call a bet should also increase. For the von Neumann model, we have shown that tight play is best-responded to with increased bluffing, and that loose play by an opponent should correspond to a decrease in your propensity to bluff. We also saw, how the overall probability that you bet should increase versus a tight opponent and decrease against a loose opponent. We also saw that as the bet size increased, the probability of checking should increase and the probability of betting should decrease.

It is our hope that these results, given in the lemmas throughout the paper, agree with your intuition and experience in all poker variants. We feel that these notions of how to react to tight or loose play by an opponent, and to increased bet sizes, could be used as general guidelines for optimal play in nearly any poker variant one might be playing, not just the highly simplified models presented here.

References

- [1] N. J. Bearden, W. Schultz-Mahlendorf, and S. Huettel. An experimental study of von neumann's two person $[0,1]$ poker. *Unpublished*.
- [2] R. Bellman. On games involving bluffing. *Rendiconti del Circolo Math. di Palermo*, 1:139–156, 1952.
- [3] R. Bellman and D. Blackwell. Some two-person games involving bluffing. *Proc. Nat. Acad. Sci.*, (35):600–605, 8/4/49.

- [4] Emile Borel. *Applications Aux Jeux des Hazard*. Gautier-Villars, 1938.
- [5] C. Ferguson and T. S. Ferguson. On the borel and von nuemann poker models. *Game Theory and Applications*, 9:17–32, 2003.
- [6] D. J. Newman. A model for "real" poker. *Operations Research*, 7:557–560, 1959.
- [7] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1953.