When to hold 'em

Kaity Parsons, Peter Tingley* and Emma Zajdela

Dept. of Mathematics and Statistics,
Loyola University Chicago
ptingley@luc.edu

Summary   We consider the age-old question: how do I win my fortune at poker? First a disclaimer: this is a poor career choice. But thinking about it involves some great math! Of course you need to know how good your hand is, which leads to some super-fun counting and probability, but you are still left with questions: Should I hold 'em? Should I fold 'em? Should I bet all my money? It can be pretty hard to decide! Here we get some insight into these questions by thinking about simplified games. Along the way we introduce some ideas from game theory, including the idea of Nash equilibrium.

Introduction

So you want to win at poker? Certainly you need to know how good any hand is, but that isn't the whole story: You still need to know what to do. That is, in the words of Don Schlitz [**15**] (made famous by Kenny Rogers [**14**]), you gotta know when to hold 'em, know when to fold 'em, know when to walk away, know when to run. Well, you're on your own for when to walk away and when to run. That leaves when to hold 'em, when to fold 'em, and a crucial question that was left out: when to bet.

We should pay attention to the real world, and the real world tells us the answers are probably interesting. Good poker players do some strange-looking things:

- Bet with very bad hands.
- Fail to bet with very good hands.
- Fold with good hands.

They even have names for these. The first is called bluffing and the second is called slow-playing. The third I guess is just called folding. One thing to think about is, are these actually good strategies? In some sense they must be, since the best players use them, but why? One possibility is that it is psychological: the players are messing with each other, trying to get each other to make mistakes. Let's eliminate that explanation. What if you are playing against a computer, and the computer is programmed to play perfectly. Then does it make sense to bluff? Or to slow-play?

We will answer these questions, but starting with real poker is too complicated. Instead we think about some simplified games to gain insight into the real thing. Simplified poker has of course been studied before, including by Borel [**2**] and von-Neumann and Morgenstern [**17**], and more recently by Chen and Ankenman [**4**] and Ferguson, Ferguson and Gawargy [**9**]. This last we draw on quite heavily.

Dice poker

Here are the rules to the first game we will consider:

- There are only two players, P1 and P2.
- Each player begins by putting $1 in the pot (the "ante").

	4	5	6	Average
4C	0	-2	-2	-1.33
4F	-1	-1	-1	-1
5C	2	0	-2	0
5F	-1	-1	-1	-1

Figure 1   Deciding on a response to P1 betting on 4,5 and 6.

- Each player's hand is determined by rolling a die, so the possible hands are 1,2,3,4,5,6 and all are equally likely. The roll is hidden from the other player.
- After seeing their hand, P1 can either bet another $1 or pass.
- If P1 bet, P2 can either **call** by also placing an extra $1 in the pot, or **fold**, in which case P1 gets the money in the pot.
- If P1 passed or P1 bet and P2 called they compare hands and the higher number gets all the money in the pot. If there is a tie, they split the pot.

What we call passing is often called checking in poker. We use the term pass partly because call and check both start with C, which messes up notation.

We want to understand how P1 should play, where, as in the introduction, P2 is a computer that plays perfectly. But first we need to understand what we mean when we say that P2 plays perfectly. This is confusing because how P2 should play certainly depends on how P1 is playing...it gets circular!

Here is how we get around this: Instead of just letting them play, we make the players each write a computer program to play for them (like in pokerbots [**3**]). Furthermore, we make P1 write their program first, and **let P2 see it**. We assume P2 is an expert poker player/programmer, and writes the perfect program to do as well as possible, using their knowledge of P1's program. The question is, under these conditions, what is P1's best strategy?

Letting P2 see P1's strategy seems unfair, but it is also sort of realistic: Players can usually observe each other and learn about each other's strategies. So another way of wording this is that P2 has been observing P1 for a long time, and knows how P1 plays. Note that P2 does not get to see P1's hand.

Anyway, we are looking for P1's best strategy, but what exactly is a strategy? Well, it is the program P1 writes, which has to take P1's hand and decide whether to bet or pass. So a strategy should be the information of what to do with each possible hand.

Let's try a reasonable looking strategy and see what happens: say P1 passes if they have 1,2 or 3 and bets if they have 4,5 or 6. P2 gets to take this information and figure out what to do. They should only think about what to do if P1 bets, since if P1 passes then hands are revealed without P2 ever having to make a decision. If P1 bets then P2 **knows they have 4, 5 or 6**. If P2 has 1,2 or 3, they will definitely lose, so should fold. If P2 has 6 they definitely don't lose, so should call. The only question is what to do with 4 or 5. To figure that out, P2 can use the table in Figure 1. The rows show what happens after P1 bets if P2 has 4 or 5 and either calls or folds. In each case the expected payout to P2 is the average the three entries in that row. If P2 has 4, folding is better ($-1 > -\frac{4}{3}$), and if P2 has 5 calling is better. We have just figured out that P2's **best response** to P1's strategy is to fold if they have 1234, and call if they have 5 or 6.

So, is this a good strategy for P1? To decide we need to find the expected payout, which we do by making a table of possibilities, showing payouts to P2. See Figure 2. All pairs of hands are equally likely so the expected payout to P2 is the average of the 36 possibilities, which is $\frac{1}{36}$. P1 is losing money! Is this reasonable? Maybe, since P2 did get to look at P1's strategy...but actually P1 has an obvious strategy to break even:

Hand	1P	2P	3P	4B	5B	6B	Total of row
1F	0	-1	-1	-1	-1	-1	-5
2F	1	0	-1	-1	-1	-1	-3
3F	1	1	0	-1	-1	-1	-1
4F	1	1	1	-1	-1	-1	0
5C	1	1	1	2	0	-2	3
6C	1	1	1	2	2	0	7

Figure 2   Possible outcomes to P2 if P1 bets with 456 and P2 calls with 56.

Hand	1B	2P	3P	4P	5B	6B	Row Total
1F	-1	-1	-1	-1	-1	-1	**-6**
1C	0	-1	-1	-1	-2	-2	-7
2F	-1	0	-1	-1	-1	-1	-5
2C	2	0	-1	-1	-2	-2	**-4**
3F	-1	1	0	-1	-1	-1	-3
3C	2	1	0	-1	-2	-2	**-2**
4F	-1	1	1	0	-1	-1	-1
4C	2	1	1	0	-2	-2	**0**
5F	-1	1	1	1	-1	-1	0
5C	2	1	1	1	0	-2	**3**
6F	-1	1	1	1	-1	-1	0
6C	2	1	1	1	2	0	**7**

Figure 3   P2's possible responses if P1 bets on 156, showing payouts to P2. P2's best response is to chose the behaviors (call or fold) that has the highest expected outcome with each possible hand. So they should fold with 1 and call with 23456.

never bet! Then both players are equally likely to have the better hand and win \$1.

Can P1 do better than breaking even? We could try some other straightforward strategies. Similar calculations show that betting on 56 or betting on only 6 both lead to breaking even. Betting on more highish hands, such as on 3456, loses even more money. There is no way for P1 to make money with this type of strategy.

Right, but we forgot one of our main questions: should P1 bluff? We need to try a strategy involving bluffing! That is, P1 should try betting on some bad hands. They should probably bet on more good hands than bad hands, so let's try betting on 156. To figure out P2's best response, we fill out the whole table of possibilities, see Figure 3. This shows that P1 should fold with 1 and call with 23456. The expected outcome to P2 is the average of the resulting 36 possibilities, which is now $-\frac{2}{36}$. This time P1 is making money. And to do so their strategy **must** involve bluffing!

You should pause and think about this: We just showed that, even though P2 knows P1's strategy, and plays perfectly, P1 can still make money, but only if their strategy includes bluffing. So bluffing is a necessary part of a good strategy for P1, it is not just psychological. This is a simplified game, but we hope you agree it is realistic enough to suggest that bluffing is a good idea more generally (which it is).

Notice though that P2 is playing pretty strangely. For instance, does it really make sense to call with 2? This should suggest that we don't have the full answer yet.

Hand	**1P**	1B	**2P**	2B	**3P**	3B	4P	**4B**	5P	**5B**	6P	**6B**
1F	0	1	1	1	1	1	1	1	1	1	1	1
2C	-1	-2	0	0	1	2	1	2	1	2	1	2
3C	-1	-2	-1	-2	0	0	1	2	1	2	1	2
4C	-1	-2	-1	-2	-1	-2	0	0	1	2	1	2
5C	-1	-2	-1	-2	-1	-2	-1	-2	0	0	1	2
6C	-1	-2	-1	-2	-1	-2	-1	-2	-1	-2	0	0
Column total	**-5**	-9	**-3**	-7	**-1**	-3	1	**1**	4	**5**	5	**9**

Figure 4   P1's options if P2 will call with 23456. Payouts are to P1. For each hand P1 should choose the behavior (bet or pass) with the highest expected payout, so they should bet with 56 and pass with 123, and either choice is equally good with 4.

If both players know each other's strategies: Nash equilibrium

One part of this setup tends to make people uncomfortable: letting P2 see P1's strategy. We argued that this was reasonable because over time P2 could observe what P1 does and figure out their strategy. But by that logic P1 can just as easily observe P2's strategy. So, if good players play each other many times, it seem they should end up each playing a best response to the other. This is essentially the idea of a Nash equilibrium.

This idea is usually attributed to Nash for his work [**11, 12**] from the early 1950's, but it was around earlier, at least in special cases. See for instance Cournot's work [**6**] from 1838. Our games are zero-sum (gain to one player equals loss to the other), and Nash equilibria in that setting were studied by von Neumann and Morgenstern in the 1940's [**17**], where they are called saddle points in strategies. Anyway, this is so crucial we will actually define it!

Definition 1. A **Nash equilibrium** for a two player game is a pair of strategies, one for each player, such that each is a best response to the other.

This means that, if one player changes their strategy and the other does not, the player who changes does no better (on average) than they would have done by following their original strategy. Essentially, each is doing as well as possible assuming their opponent is playing well, so in that sense the strategies are optimal.

Let's think about dice poker some more. We found that P1 betting on 156 is a pretty good strategy, and that P2's best response is to call with 23456. But, if P1 knows that P2 will call with 23456, do they still want to bet on 156? The answer is no: As in Figure 4, P1 would now like to bet with 56 and pass with 123, and it doesn't matter with 4. So P1 betting on 156 is not their best response, and hence these strategies (P1 bet on 156, P2 call on 23456) are not a Nash equilibrium.

Is there a Nash equilibrium? Well, if you let two good players play repeatedly, they should settle on strategies somehow, and it seems that these should be an equilibrium. Or they could just keep changing strategies, so maybe not...

In fact, with the type of strategy we've been using, there is no Nash equilibrium. One way to see this is to check that, for any P1 strategy, if you (i) find P2's best response, and then (ii) find P1's best response to that, you don't get back the same P1 strategy. So, that strategy cannot be part of an equilibrium pair. Doing this in every case is annoying, but you can probably convince yourself it is true by checking a few plausible P1 strategies (maybe bet on 1456 or on 1256).

It might seem that we are stuck, but we missed another important thing you can learn from the real world: poker players like to be unpredictable, and don't always do the same thing in the same situation. We need more randomness!

	Hand	1: $\frac{2}{3}$B	$+\frac{1}{3}P$	2P	3P	4P	5B	6B	Total of row
1	F	$\frac{2}{3}(-1)$	$+\frac{1}{3}0$	-1	-1	-1	-1	-1	-5.66
2	F	$\frac{2}{3}(-1)$	$+\frac{1}{3}1$	0	-1	-1	-1	-1	-4.33
3	$\frac{1}{3}F$ $+\frac{2}{3}C$	$\frac{1}{3}\frac{2}{3}(-1)$ $\frac{2}{3}\frac{2}{3}(2)$	$+\frac{1}{3}\frac{1}{3}1$ $+\frac{2}{3}\frac{1}{3}1$	$\frac{1}{3}1$ $\frac{2}{3}1$	$\frac{1}{3}0$ $\frac{2}{3}0$	$\frac{1}{3}$(-1) $\frac{2}{3}$(-1)	$\frac{1}{3}(-1)$ $\frac{2}{3}$(-2)	$\frac{1}{3}(-1)$ $\frac{2}{3}$(-2)	-2.33
4	C	$\frac{2}{3}(2)$	$+\frac{1}{3}1$	1	1	0	-2	-2	-0.33
5	C	$\frac{2}{3}(2)$	$+\frac{1}{3}1$	1	1	1	0	-2	2.66
6	C	$\frac{2}{3}(2)$	$+\frac{1}{3}1$	1	1	1	2	0	6.66

Figure 5   Table of payouts to P2 for the Nash equilibrium in dice poker. Some outcomes now have probabilities associated with them. For instance, if P1 has 1 and P2 has 3, there is a probability of $\frac{1}{3}\frac{2}{3}$ that the outcome is -1 to P2, because P1 bets with probability $\frac{2}{3}$, and then P2 folds with probability $\frac{1}{3}$.

Mixed strategies

A **mixed strategy** is a strategy where a player doesn't always do the same thing in the same situation. For instance, P1 can decide that, if they have 4, they will bet half the time and pass half the time. Then, whenever they are dealt 4, they randomize, maybe by flipping a coin, to decide which to do.

The famous (but kind of difficult) von-Neumann mini-max theorem from [**16**] (see also e.g. [**1**, Chapter 1.2]) implies that, if mixed strategies are allowed, Nash equilibria always exist (provided the number of unmixed strategies is finite, or some other more technical conditions hold). In any case, such an equilibrium does exist for our game. It is: P1 bets with 56, and bets $\frac{2}{3}$ of the time with 1. P2 calls with 456, and also calls $\frac{2}{3}$ of the time with 3. The table of payouts is given in Figure 5, and works out to an expected payout to P2 of $\frac{-3.33}{36}$. So P1 is making significantly more money than with the strategy "bet with 156" we studied before (which is actually the best possible unmixed strategy).

To check that this is a Nash equilibrium just check that, in every situation, each player's expected value does not go up if they change strategies. For instance, if P1 decided to pass with 5, their expected payout would be 3, which is worse than the 3.66 they are currently getting. If P2 has a 4 and decides to fold they get exactly the same as if they call (-0.33), which is still alright.

Something interesting: If P1 has 1, sometimes they bet and sometimes they pass. If this is a Nash equilibrium, in either case they should not want to change. That means the two possible strategies must lead to the same expected payout: if one was worse P1 would never want to play it! This observation is called the **principle of indifference**. By the same argument P2 should get the same payout for passing and calling with 3.

If we somehow knew P1 and P2 should randomize with 1 and 3 respectively, this gives a way to find the Nash equilibrium: Let x be the probability $P1$ bets with 1. If P2 gets a 3, then their expected payouts are:

$$\text{Fold: } \frac{x(-1) + (1-x)1 + 1 + 0 - 1 - 1 - 1}{6} = \frac{-1 - 2x}{6}.$$

$$\text{Call: } \frac{x(2) + (1-x)1 + 1 + 0 - 1 - 2 - 2}{6} = \frac{-3 + x}{6}.$$

These must be equal, and solving gives $x = \frac{2}{3}$. A similar calculation using P1's indif-
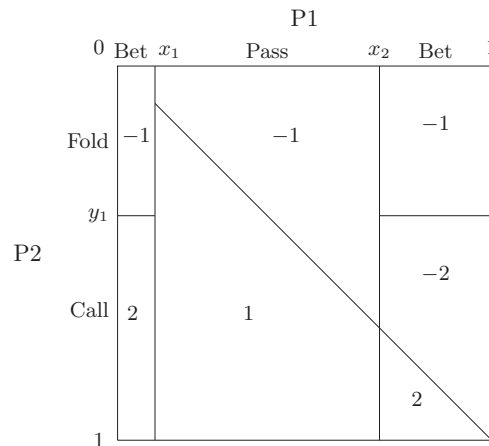
Figure 6   Strategies and payouts for the continuous game. This is interpreted as follows: If P1 gets hand x and P2 gets hand y, plot x and y on the top and left respectively, and then read off what they do. The number in the region containing (x, y) is the resulting payout to P2.

ference with 1 gives that P2 should fold with probability $\frac{1}{3}$ with a 3.

We had to guess which hands to randomize with. If we had guessed wrong we would have figured that out. For instance, if we guessed that P2 should fold with 3 and randomize with 4, then P1's indifference with 1 would imply that P2 should call with probability $\frac{5}{3}$ with 4, which is impossible. We could keep trying until we got it right, but that is annoying. We now move to a more continuous situation, which eliminates some of the guessing.

Allowing infinitely many possible hands

Now let's think about the same game but with one change: each player's hand is now a random number in the interval [0,1]. This is actually closer to real poker, where there are lots of possible hands and each has a different probability of winning, which can be expressed as a number in $[0, 1]$. This has been studied many times before, dating to von-Neumann and Morgenstern [**17**, Chapter 19.14-19.16] (see also e.g. [**7**], [**9**], [**13**, Chapter 5], [**4**, Example 11.3]). A variant where P1 must fold if they don't bet was studied even earlier by Borel [**2**, Chapter 5] (see also [**10**, §9.2]).

We want to find optimal strategies, which we now know means we are looking for a Nash equilibrium. We start by guessing that things are qualitatively similar to the dice game, and P1 should use a bluffing strategy: For some $x_1 < x_2$ P1 should bet with hands $h < x_1$ and $h > x_2$, and pass for $x_1 \leq h \leq x_2$. P2 should call with hands better than some cutoff value y_1. It also seems reasonable to guess that $x_1 \leq y_1 \leq x_2$. The strategies and payouts to P2 are then described by a "table," as in Figure 6.

We want to find the cutoffs. The key idea is that, if e.g. P1 has hand exactly x_1, it shouldn't matter if they bet or pass. This is because with a slightly better hand passing is better, and with a slightly worse hand betting is better. But if the behavior (bet or pass) is kept fixed and the hand varies the payout changes continuously. So the payout for a bet and a pass with exactly x_1 must be identical. This is really an indifference principle as above: P1 can play a mixed strategy with x_1 if they like, although getting exactly x_1 has probability zero, so never happens.

To find the payout if P1 has x_1 and bets, we average over possible P_2 hands:

- If P2 has hand $h < y_1$ they fold and P1 wins \$1. This has probability y_1
- If P2 has hand $h > y_1$ they call and P1 loses \$2. This has probability $1 - y_1$.

If P1 has hand x_1 and passes,

- If P2 has hand $h < x_1$ then P1 wins \$1. This has probability x_1
- If P2 has hand $h > x_1$ then P1 loses \$1. This has probability $1 - x_1$.

So we get the **indifference equation**:

$$y_1 - 2(1 - y_1) = x_1 - (1 - x_1).$$

We also get equations from P1's indifference with x_2 and P2's indifference with y_1:

$$\text{P1 with } x_2: \quad -x_2 + (1 - x_2) = -y_1 - 2(x_2 - y_1) + 2(1 - x_2),$$
$$\text{P2 with } y_1: \quad -x_1 - (1 - x_2) = 2x_1 - 2(1 - x_2). \tag{1}$$

This is just a system of three equations and three unknowns! Solving,

$$x_1 = \frac{1}{10}, \quad x_2 = \frac{7}{10}, \quad y_1 = \frac{4}{10}.$$

We made a guess: that the solution has the rough form shown in Figure 6. So we need to check that our answer really is a Nash equilibrium. There are three cases:

1. If P2 has hand $h = \frac{4}{10}$ they get the same payout for calling and folding. After a bet the payout for folding is constant (-1) and the payout for calling is weakly increasing with h, so folding is best if $h < \frac{4}{10}$ and calling is best if $h > \frac{4}{10}$.
2. P1 gets the same payout for passing as bluffing with hand $h = \frac{1}{10}$. The payout for bluffing is independent of h for $h \leq \frac{4}{10}$ and the payout for passing is increasing, so this implies that bluffing is better than passing for $h < \frac{1}{10}$, and that passing is better than bluffing for $\frac{1}{10} < h < \frac{4}{10}$.
3. If P1 has a hand $h > \frac{4}{10}$, their expected payout if they bet is $\frac{4}{10} + 2(h - \frac{4}{10}) - 2(1 - h)$, and their expected payout if they pass is $2h - 1$. Subtracting, they expect to win $2h - \frac{7}{5}$ more by betting, which is positive for $h > \frac{7}{10}$ and negative for $h < \frac{7}{10}$, so passing is better if $\frac{4}{10} < h < \frac{7}{10}$ and betting is better if $h > \frac{7}{10}$.

Neither player wants to change in any situation, so this is in fact a Nash equilibrium.

Things often work this way: you make some guesses, then solve indifference equations to get an answer, then you have to check that it really is an equilibrium, which proves your guesses were correct...or they weren't, and you try again.

Anyway, we have an answer! P1 bluffs with the worst 10% of hands, **value bets** on the best 30% of hands, and otherwise passes. P2 calls a bet with any hand better than 0.4. To find the expected payout just add up the payout in each region times its area (the probability of landing in that region). It works out to P2 losing \$0.1 a hand.

This game makes sense for any bet size a. The 'table' is the same, except all 2s become $1 + a$. The cutoffs are the solutions x_1, x_2, y_1 to the indifference equations

$$2(y_1 - x_1) = a(1 - y_1), \ \ 1 - x_2 = x_2 - y_1, \ \ (2 + a)x_1 = a(1 - x_2). \tag{2}$$

Different bet sizes lead to different payouts, and it turns out that $a = 2$ is best for P1, giving a payout of $\frac{1}{9}$ (see [**9**, §1.1]). However, if P1 can bet different amounts with different hands, then they can do even better, up to a payout of $\frac{1}{7}$ (see [**13**, §5.2]).

Figure 7   Betting tree if P2 can bet after a P1 pass. Here payouts are to P2 and $\pm$ means $+$ if P2 has the higher hand and $-$ if P1 has the higher hand.

More betting

We now allow P2 to bet, but only if P1 passes. That is, we consider the game as above, with P1's bet size being a, and the following new options:

- If P1 passes then P2 can choose to bet a fixed amount b.
- If P2 bets then P1 can call or fold.

We do not assume P1 and P2 use the same bet size, but we do assume each uses a single fixed bet size whenever they bet. This can be described by the tree in Figure 7. This game has been studied before, for instance in [**9**], where all the results here can be found. See also [**4**, Example 17.1].

As usual we have to guess the rough layout, and this time we make two separate guesses. See Figure 8. For any value of a, b we solve indifference equations to find the cutoffs. The diagram also implies an assumed ordering on the cutoffs. For example, in (L), $x_1 \leq y_1 \leq x_2, y_2 \leq x_3 \leq y_3 \leq x_4$. These constraints must hold in addition to the indifference equations. We will see that (L) works when $a \leq b$ and (R) works when $a \geq b$. If $a = b$ this means there are two natural Nash equilibria!

Each cutoff determines one indifference equation: If a player has exactly that hand, then the strategies on the two sides of the cutoff must give the same expected payout. For instance, for (L), there are 7 equations:

$$-y_2 + (1+a)(1-y_2) = 1 \tag{3}$$

$$y_1 + (1-y_3) = -(1+b)y_1 + (1+b)(1-y_3) \tag{4}$$

$$-(1+b)y_1 - (x_3 - y_1) + (y_3 - x_3) + (1+b)(1-y_3) =$$
$$-y_2 - (1+a)(x_3 - y_2) + (1+a)(1-x_3) \tag{5}$$

$$-y_2 - (1+a)(x_4 - y_2) + (1+a)(1-x_4) =$$
$$-(1+b)y_1 - (y_3 - y_1) - (1+b)(x_4 - y_3) + (1+b)(1-x_4) \tag{6}$$

$$(x_2 - y_1) - (1+b)(x_3 - x_2) - (1+b)(1-x_4) = -(x_3 - y_1) - (1-x_4) \tag{7}$$

$$-x_1 - (x_4 - x_3) = (1+a)x_1 - (1+a)(x_4 - x_3) \tag{8}$$

$$(x_3 - x_2) - (1-x_4) = (1+b)(x_3 - x_2) - (1+b)(1-x_4). \tag{9}$$

Figure 8 Two candidate Nash equilibria. Here P1's hand is increasing left to right and P2's hand is increasing top to bottom. The strategies are given on the top and left respectively, where B stands for bet, P for pass, C for call, and F for fold. There are two rows for P1's strategy, since there are two situations where they may have to make a choice: at the start of the game (S) and after a P2 bet (B). Similarly there are two columns for P2's strategy, showing what P2 does if P1 passed (P) or bet (B). The quantities in the interior show payouts to P2.

We solve using Mathematica, see [**5**]. For $a = b = 1$ we get

$$(L): \quad x_1 = \frac{1}{9}, \; x_2 = \frac{1}{3}, \; x_3 = \frac{1}{2}, \; x_4 = \frac{5}{6}, \; y_1 = \frac{1}{6}, \; y_2 = \frac{1}{3}, \; y_3 = \frac{1}{2}.$$
$$(R): \quad x_1' = \frac{1}{9}, \; x_2' = \frac{1}{3} \; x_3' = \frac{2}{3}, \; y_1' = \frac{1}{6}, \; y_2' = \frac{1}{3}, \; y_3' = \frac{1}{2}. \tag{10}$$

Note that for these values $P2$ uses the same strategy in both (L) and (R).

We made some guesses, so we need to check if these really are Nash equilibria. For now we restrict to $a = b = 1$ when we show that both are. They then give the same payout (as they must, see e.g. [**1**, §1.1]) which is $\frac{1}{18}$ to P2.

We start by showing that P1 is responding optimally to P2's strategy. Consider P1's expected payout as a function of their hand h if P2 uses the cutoffs $y_1 = \frac{1}{6}$, $y_2 = \frac{1}{3}$, $y_3 = \frac{1}{2}$ from (10) and P1 plays

PF: pass then fold if P2 bets,   PC: pass then call if P2 bets,   B: bet.

A little work shows these functions are

$$\text{PF}(h) = \begin{cases} -1 & h \leq \frac{1}{6} \\ -\frac{4}{3} + 2h & \frac{1}{6} \leq h \leq \frac{1}{2} \\ -\frac{1}{3} & h \geq \frac{1}{2} \end{cases} \qquad \text{PC}(h) = \begin{cases} -\frac{5}{3} + 4h & h \leq \frac{1}{6} \\ -\frac{4}{3} + 2h & \frac{1}{6} \leq h \leq \frac{1}{2} \\ -\frac{7}{3} + 4h & h \geq \frac{1}{2} \end{cases}$$

$$\text{B}(h) = \begin{cases} -1 & h \leq \frac{1}{3} \\ -\frac{7}{3} + 4h & h \geq \frac{1}{3}. \end{cases} \tag{11}$$

These are shown in Figure 9. Both of P1's strategies in (10) always choose an optimal

Figure 9   P1's expected payout as a function of hand for their three possible behaviors if P2 plays with cutoffs from (10). Here the dotted line is for PF, the dashed line for PC, and the gray line for B.

Figure 10   P2's payout as a function of hand if P1 plays as in (L), shown on the left, and (R), on the right. The lines are for P2 playing: (i) Solid: FP, fold after a bet and pass after a pass. (ii) Dashed: FB, fold after a bet and bet after a pass. (iii) Dotted: CP, call after a bet and pass after a pass. (iv) Gray: CB, call after a bet and bet after a pass.

response. For instance for $h > \frac{5}{6}$, (L) plays PC and (R) plays B, but both are optimal because the dashed and gray lines coincide and are at the top in that range.

We also need to show that P2 is responding optimally in every situation. This time we need two graphs to show what happens against P1's two different strategies in (10). See Figure 10. In both cases the dashed line is maximal for $h < \frac{1}{6}$, the solid line is maximal for $\frac{1}{6} < h < \frac{1}{3}$, the dotted line is maximal for $\frac{1}{3} < h < \frac{1}{2}$, and the gray line is maximal for $h > \frac{1}{2}$. So this P2 strategy is a best response to both P1 strategies.

We have now shown that, at $a = b = 1$, both (L) and (R) really are Nash equilibria! In (L) P1 is slow-playing their very best hands, by passing and hoping P2 bets. So slow-playing is a reasonable strategy. Although, as discussed in [**9**], (R) may be better for P1 in practice since it penalizes more common P2 errors.

This argument works to show that (L) is an equilibrium in most cases with $a \leq b$, (R) is an equilibrium in most cases with $a \geq b$, and both are equilibria whenever $a = b$. If a and b are far apart some other cases appear, basically because the bluff cutoffs x_1 and y_1 reverse, changing the indifference equations. These are handled in [**5**].

If $b > a$ then (R) is never an equilibrium, even though solving the equations often gives cutoffs in the correct order (this happens at e.g. $a = 1, b = 2$). This is because, if $b > a$, then $\text{PC}(h)$ is increasing faster than $\text{B}(h)$ for $h > x_3'$. But $\text{B}(x_3') = \text{PC}(x_3')$ so, for $h > x_3'$, $\text{PC}(h) > \text{B}(h)$. Since (R) plays B for $h > x_3'$ this violates the Nash condition. Similarly (L) is never a Nash equilibrium if $a > b$.

Remark. A situation like this one shows up in real poker: if there are only two players, and one of them has so few chips that only one bet can be made. But even then there are additional complications, for instance due to a flop or draw. Some of our conclusions, such as that optimal play involves bluffing, certainly still hold. But we will not analyze even the simplest situation in a real game here.

REFERENCES

1. E. N. Barron. Game theory. An introduction. 2^{nd} edition. Wiley Series in Operations Research and Management Science. John Wiley & Sons, Inc., Hoboken, NJ, 2013.

2. É. Borel (1938) Traité du Calcul des Probabilités et ses Applications Volume IV, Fascicule 2, Applications aux jeux des hazard, Gautier-Villars, Paris.

3. Pokerbots: `http://pokerbots.org/`

4. B. Chen and J. Ankenman. *The Mathematics of Poker*. ConJelCo LLC. 2006.

5. Code is available at `http://webpages.math.luc.edu/~ptingley/`

6. A. Cournot. Researches on the Mathematical Principles of the Theory of Wealth.

7. C. Ferguson and T. Ferguson. On the Borel and von Neumann poker models. *Game theory and applications* **9** (2003) 17–32.

8. C. Ferguson and T. Ferguson. The endgame in poker. `https://www.math.ucla.edu/~tom/papers/poker3.pdf`

9. C. Ferguson. T. Ferguson and C. Gawargy. Uniform (0, 1) two-person poker models. *Game Theory and Applications* **12** (2004) 17–37.

10. S. Karlin. Mathematical methods and theory in games, programming and economics. Vol. I: Matrix games, programming, and mathematical economics. Vol. II: The theory of infinite games. Addison-Wesley Publishing Co. 1959.

11. J. Nash. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci. U. S. A.* **36(1)** (1950) 48–49.

12. J. Nash. Non-cooperative games. *Ann. of Math.* (2) **54** (1951) 286–295.

13. V. Mazalov. *Mathematical game theory and applications*. John Wiley & Sons, Ltd., Chichester, 2014.

14. Kenny Rogers. *The Gambler*. From "The Gambler," United Artists Group 1978.

15. Don Schlitz. The Gambler (song). Written 1976.

16. J. von Neumann. Zur Theorie der Gesellschaftsspiele *Math. Annalen.* **100** (1928) 295–320.

17. J. von Neumann and O. Morgenstern. Theory of Games and Economic Behavior. Princeton University Press, Princeton, New Jersey, 1944.

KAITY PARSONS (MR Author ID: 1345241) received a B.S. in mathematics from Loyola University Chicago in 2017. Since graduating, Kaity has been working as a center director at the Town and Country Mathnasium in St. Louis Missouri. She and her partner are currently planning a one-way cross-country road trip with their dog, Angel and their cat, Salem, to seek what they will find.

PETER TINGLEY (MR Author ID: 679482) received a Ph.D. in Mathematics from the University of California, Berkeley in 2008. He spent short periods at the University of Melbourne (Australia) and MIT, and since 2012 has been at Loyola University Chicago. He also helps run the Chicago Math Teachers' circle, which he co-founded in 2015, and is the current blue division racquetball champion at the Evanston YMCA.

EMMA ZAJDELA (MR Author ID: 1345242) received a B.S. in math and physics from Loyola University Chicago in 2016 and an M.S. in math from the University of Illinois Chicago in 2018. She is currently an NSF fellow and Ph.D. student in applied math at Northwestern University. Since 2016 she has served as Assistant to the President of the Malta Conferences Foundation, a nonprofit that uses science as a bridge to peace in the Middle East. She also recently received her yellow belt in judo.