MANAGEMENT SCIENCE
Vol. 17, No. 12, August, 1971
Printed in U.S.A.

OPTIMAL BLUFFING STRATEGIES IN POKER*

LAWRENCE FRIEDMAN

Mathematica, Princeton, New Jersey

Bluffing in poker is examined as a problem in game theory. A very common situation occurs where the 'kitty' contains K, player B has the apparent high hand, and player A has an apparent probability, P, of having a better hand than B, and considers a bluff. Under these conditions it can be shown that A should raise 1 unit with probability $(1/(1 + K))(P/(1 - P)) \cdot B$, in turn, should call A's potential bluff with probability $K/(1 + K)$. This says that in pot limit poker, if the potential bluffer has the appearance of having the winning hand with probability .25, he should bluff $\frac{1}{2} \cdot \frac{1}{3} = \frac{1}{6}$ of the time. The optimal strategy is to call a potential bluffing hand half the time.

The game of poker is one of the most popular of all card games. It has many features which make it appealing to an analytically minded person: (1) the calculation or estimation of probabilities, (2) the use of expected gain as a decision criterion, (3) the management of money, and (4) the use of mixed strategies in bluffing situations.

Many aspects of the game are subject to exact analysis which, however, can be exceedingly complicated. The necessity of making rapid decisions at each play and the large number of possible situations precludes the possibility of complete analysis before the game or exact analysis of special situations arising during the course of the game. Also a large part of the interest in the game arises not from mathematical but from psychological problems. For example, one is constantly trying to detect behavior patterns of one's opponents and wherever possible take advantage of poor play.

One of the most interesting aspects of the game of poker is the problem of bluffing. Bluffing can be seen as a game theory problem in its purest sense. The interest arises because it is quite clear to those who have played much poker that some sort of mixed strategy with regard to bluffing must be used.

Many of the bluffing situations arising in poker are subject to exact analysis, and optimal strategies can be determined. However, there are so many types of situations that one cannot analyze all of them. Also the situations involving more than two players are subject to the uncertainties of n-person game theory.

This paper will attempt to describe two of the most common bluffing situations arising in poker and determine the optimal behavior strategies for the players.

The "Classical" Bluffing Situation

A very common bluffing situation arising in poker is described by the following conditions:

1. Two players A and B remain in contention for the pot.
2. Player B apparently has the high hand.
3. Player A has a chance of having a hand which player B could not possibly beat.
4. Previous betting or action of player A does not give away his true situation.

This situation is best described by a few examples:

(1) In 5 card stud A has 4 low hearts showing and a hidden card. Player B has a pair of aces. Player B checks to player A. A must decide whether or not to bluff, and if A raises, B must decide whether or not to call.

* Received November 1969; revised May 1970.

(2) In draw poker, B has opened and drawn 3 cards. A has called and drawn one card, indicating he was probably trying for a straight or flush. It is very unlikely that B has improved his hand enough to beat a straight or flush so he checks to A. If A bets, B must decide whether or not to call.

Many variations of this situation can occur in poker. The essential characteristics of this particular situation are that A may attempt a bluff and there is no reraise by B. (Unless B was quite sure that A was bluffing a reraise would be foolish.)

To clarify the situation, we shall assume that there is 1 unit in the pot at the final betting stage and that the rules of betting enable the maximum raise to be the amount already in the pot (pot limit poker). Poker games with a fixed limit on the size of bet are very discouraging to bluffing. These situations can also be analyzed, but are of less interest since optimal strategies usually involve very little bluffing.

Let α = the fraction of the time that A bluffs when he does not have the winning hand. (If A does have the winning hand he will, of course, raise the pot.)

β = the fraction of the time that B should call A's raise.

P = the probability that A has the winning hand given the exposed cards (e.g. in example 1 this would be the probability that A's hidden card is a heart). The estimation of P is made by player A and in example 1 can be calculated exactly. Certainly, in a fast moving poker game, one cannot take the time to do exact calculations and must be satisfied with an estimate.

The problem is to determine optimum behavioral strategies for players A and B. The first principle to consider in this analysis and indeed in all poker situations is that the player making a decision should not consider the money in the pot as partially belonging to him. His decision on betting should be based on whether or not his expectation of gain exceeds the amount to be bet.

The situation in our example is shown by the alternative decision diagram below. Note that because of our basic principle the money previously put into the pot belongs to no one, the game at this point is nonzero sum.

If A bluffs with probability α when he has the losing hand, and B calls A's raises with probability β we have the following expected gains for hands of this type.

$$\begin{aligned} A\text{'s Expected Gain} &= (1 - P)[-\alpha\beta + \alpha(1 - \beta)] + P[2\beta + (1 - \beta)] \\ &= \alpha(1 - P)(1 - 2\beta) + P(1 + \beta), \\ B\text{'s Expected Gain} &= (1 - P)[2\alpha\beta + (1 - \alpha)] + P(-\beta) \\ &= \beta[-P + 2\alpha(1 - P)] + (1 - P)(1 - \alpha). \end{aligned}$$

A's Situation	A's Decision	B's Decision	Payoffs	
			A's Gain	B's Gain
A has losing hand	A bluffs raises pot	B calls	−1	+2
		B drops	+1	0
	A drops	B wins	0	+1
A has winning hand	A raises pot	B calls	+2	−1
		B drops	+1	0

FIGURE 1. *Decision Tree for 'Classical' Bluffing Situation*

Examining A's expected gain we find that A should always bluff ($\alpha = 1$) if $(1 - 2\beta) > 0$ and A should never bluff ($\alpha = 0$) if $(1 - 2\beta) < 0$.

Examining B's expected gain, we find that B should always call ($\beta = 1$) if $[-P + 2\alpha(1 - P)] > 0$ and B should never call ($\beta = 0$) if $[-P + 2\alpha(1 - P)] < 0$.

The game has an equilibrium point solution which can be viewed as arising from the fact that both players wish to make things hard for each other. The optimum strategies are

$$A \text{ sets } \quad \alpha = P/2(1 - P) \quad \text{for} \quad P \leqq \tfrac{2}{3}$$

$$B \text{ sets } \quad \beta = \tfrac{1}{2}.$$

If A and B use their optimal strategies, we find

$$A\text{'s expected gain} = \tfrac{3}{2}P,$$

$$B\text{'s expected gain} = 1 - \tfrac{3}{2}P.$$

The result is interpreted as follows. A estimates P, the a priori probability of A's having the winning hand. For example, suppose P is estimated at .2, or 1 chance in 5. A should then bluff with probability .125. B should call A's raise with probability .5.

Since P will change with every hand, the bluff probability will also change. The call probability, however, is independent of P.

It can be similarly shown that if there are K units in the 'kitty' and the maximum raise is 1 unit, A should bet with probability α where

$$\alpha = (1/(1 + K))(P/(1 - P)).$$

And B should call with probability β where $\beta = K/(1 + K)$.

The smaller the allowable bet relative to the kitty, K, the less bluffing in the game, and the more likely the potential bluff will be called.

Detecting and Taking Advantage of Nonoptimal play

An interesting problem in game theory is the way to detect nonoptimal strategy and the best method of taking advantage of poor play. For example, in poker if one's opponent is a player who refuses to be bluffed (never drops in a situation where he may be bluffed) then it is clear that one should never bluff against him, since every time he calls, you will have a good hand, and he will lose. Unless he is very stupid he must eventually realize that his strategy of calling all raises is a poor one and that he should change it by dropping some times in certain bluffing situations.

Since his original strategy is a poor one, one does not wish him to change it. Therefore, it might be psychologically good strategy on one's part to deliberately get caught bluffing every once in a while by this player to try and convince him that his strategy of refusing to be bluffed is a good one. Thus one might deliberately sacrifice in order to induce an opponent to continue a bad strategy. Exactly how to do this properly is an aspect of game theory problems which has not been solved by game theorists, probably because it is more of a psychological problem than a mathematical one.

Because the value of P will vary from hand to hand in the situation just discussed, it is probably more difficult to detect whether or not A is using correct strategy than it is to detect B's strategy. One method of detecting A's strategy might be to keep a table showing the estimate of P, the resulting optimal value of α, and what A actually did. For example, 5 similar bluffing situations of the type described might yield the following table:

Estimate of P	Estimate of Optimal α	A's Action (Bluff 1, Nonbluff 0)
.20	.125	0
.25	.167	0
.10	.056	0
.20	.125	1
.33	.250	0
Sums	.723	1

With correct play, the sum of optimal bluff probabilities should converge on the sum of A's actions as given in column 3. In the example just given with only 5 pieces of data it is clear that one cannot say that A bluffs more than he should. If the sum in A's action column equalled 3, one might suspect that A bluffs too much. Tests of significance can be used to detect nonoptimal play. However, the real problem is how to properly take advantage of it.

In our classical bluffing situation, we have the following extreme reactions to nonoptimal play:

1. B detects that A seldom bluffs: B tends to drop on A's raises
   A's average gain approaches P
   B's average gain approaches $1 - P$
2. B detects that A bluffs very much: B tends to call A's raises
   A's average gain approaches $3P - 1$
   B's average gain approaches $2 - 3P$
3. A detects that B usually calls his raise: A seldom bluffs
   A's average gain approaches $2P$
   B's average gain approaches $1 - 2P$
4. A detects that B usually drops on a raise: A bluffs very often
   A's average gain approaches 1
   B's average gain approaches 0.

It is interesting to note that the use of the bluff by A in an optimal manner increases his expected gain by 50% from the expected gain which he would obtain if he never bluffed.

Another way of detecting A's strategy is to examine the fraction of the time in which A's raise is due to a bluff. In following the optimal strategy A will raise in a given bluffing situation $100P$ percent of the time due to the fact that he actually holds the winning hand and $(1 - P)\cdot 100P/2(1 - P)$ percent of the time on a bluff. Therefore, the fraction of A's raises which are bluffs will be

$$\frac{(1 - P)(P/2(1 - P))}{P + (1 - P)P/2(1 - P)} = \frac{1}{3}.$$

Therefore if A is using optimal strategy $\frac{1}{3}$ of his raises will be bluffs. B can detect A's strategy by examining A's hands when B calls A's raise to see whether or not A's raises are bluffs more or less than $\frac{1}{3}$ of the time. When B drops on A's raise, it is good strategy on A's part not to reveal whether or not he bluffed. The less data B has to evaluate A's strategy the better it is for A.

This method of detection of A's strategy is simpler in actual practice than the method previously described which required an estimate of P for every potential bluffing situation. However, since only about half the data can be used (the fraction of time B calls) the first method of detection may lead to a quicker estimate of A's strategy.

B's strategy is easier to detect since if he is using an optimal strategy; he should be calling about half of A's raises.

Simple Bluffing Situation with Possible Reraise

This situation is similar to the previous one except that B has a possibility of reraise. To fix ideas, consider a 5 card stud game in which A and B have the following hands.

A: 10 of clubs, 8 of hearts, 5 of spades, 4 of diamonds, one hidden card;

B: King of diamonds, Queen of clubs, 2 of clubs, 2 of hearts, one hidden card.

Let P_1 = probability that A's hidden card will beat B's hand if B does not have a good hidden card,

P_2 = probability that B's hidden card will give him the high hand even if A has a good hidden card.

This poker situation is found very often and is not limited to 5 card stud. Similar situations may arise in all types of poker usually with greater difficulty in estimating P_1 and P_2.

The following assumptions will be used in our analysis of the optimal strategy in this situation.

1. Pot limit poker is the game and at this stage there is 1 unit in the pot.
2. Previous play has given no indication of the hidden cards.
3. Player B, having the high hand, opens and checks to player A. This is considered optimal strategy for the following reasons:
   a. If A has a good hand and B has not improved, B's betting might cause a higher raise by A if A has a good hidden card.
   b. If B has improved, his betting may discourage A from attempting a bluff if he has no good hidden card. If B checks, A will certainly raise on a good card and B will have the opportunity of a reraise.

4. If B reraises A, A will drop if he was bluffing originally and may call or drop if his original bet was due to a good hidden card. In the face of a raise by B, a second raise by A would be poor strategy.

Let α_1 = A's probability of bluffing by raising the pot when he does not have a good hidden card,

α_2 = A's probability of calling B when B reraises A (A will only consider calling when he has a good hidden card),

β_1 = B's probability of raising after A's raise when B does not have a good hidden card. (With a good hidden card B must raise.)

β_2 = B's probability of calling A's raise when B does not have a good hidden card,

$1 - \beta_1 - \beta_2$ = B's probability of dropping after A's raise.

The alternative courses of play and the payoffs to both A and B are shown in the following chart.

A's expected Gain

$$= \alpha_1[(1 - P_1)(1 - P_2)(1 - 2(\beta_2 + \beta_2)) - (1 - P_1)P_2] + \alpha_2 3P_1[-P_2 + 2\beta_1(1 - P_2)] - P_1P_2 + P_1(1 - P_2)[1 + \beta_2 - 2\beta_1].$$

B's Expected Gain

$$= 2(1 - P_2)\beta_1[P_1 - 3\alpha_2P_1 + \alpha_1(1 - P_1)] + \beta_2(1 - P_2)[-P_1 + 2\alpha_1(1 - P_1)] + P_1P_2[5\alpha_2 + 2(1 - \alpha_2)] + (1 - P_1)(1 - P_2)(1 - \alpha_1) + (1 - P_1)P_2[2\alpha_1 + (1 - \alpha_1)].$$

Players Hidden Cards	A's First Decision	B's Decision	A's Second Decision	Payoffs	
				A's Gain	B's Gain
Both A and B have good hidden cards P_1P_2	A raises	B raises,	A calls, α_2	−4	+5
			A drops, $(1-\alpha_2)$	−1	+2
A has good cards and B doesn't $P_1(1-P_2)$	A raises	B raises, β_1	A calls, α_2	+5	−4
			A drops, $(1-\alpha_2)$	−1	+2
		B calls, β_2		+2	−1
		B drops, $1-\beta_1-\beta_2$		+1	0
A doesn't have good card B doesn't have good card $(1-P_1)(1-P_2)$	A raises, α_1	B raises, β_1	A drops	−1	+2
		B calls, β_2		−1	+2
		B drops, $1-\beta_1-\beta_2$		+1	0
	A drops, $(1-\alpha_1)$			0	+1
A doesn't have good card B does $(1-P_1)P_2$	A raises, α_1	B raises	A drops	−1	+2
	A drops, $(1-\alpha_1)$			0	+1

FIGURE 2. *Decision Free for Bluffing Situation with Possible Reraise*

If A and B wish to maximize their expected gain we find

$$A \text{ sets } \alpha_1 = \begin{Bmatrix}1\\0\end{Bmatrix} \quad \text{if} \quad [(1-P_1)(1-P_2)(1-2(\beta_1+\beta_2)) - (1-P_1)P_2]\begin{Bmatrix}>\\<\end{Bmatrix} 0,$$

$$A \text{ sets } \alpha_2 = \begin{Bmatrix}1\\0\end{Bmatrix} \quad \text{if} \quad [-P_2 + 2\beta_1(1-P_2)]\begin{Bmatrix}>\\<\end{Bmatrix} 0,$$

$$B \text{ sets } \beta_1 = \begin{Bmatrix}1\\0\end{Bmatrix} \quad \text{if} \quad [P_1 - 3\alpha_2P_1 + \alpha_1(1-P_1)]\begin{Bmatrix}>\\<\end{Bmatrix} 0,$$

$$B \text{ sets } \beta_2 = \begin{Bmatrix}1\\0\end{Bmatrix} \quad \text{if} \quad [-P_1 + 2\alpha_1(1-P_1)]\begin{Bmatrix}>\\<\end{Bmatrix} 0.$$

An equilibrium point solution exists and is found by setting the above expressions equal to zero. We find as the optimal strategies

$$\alpha_1 = P_1/2(1-P_1) \qquad P_1 \leqq \tfrac{2}{3}$$

$$\alpha_1 = 1 \qquad \tfrac{2}{3} < P_1 \leqq 1$$

$$\alpha_2 = \tfrac{1}{2}$$

$$\left.\begin{aligned} \beta_1 &= P_2/2(1-P_2); && \text{Raise}\\ \beta_2 &= \tfrac{1}{2} - P_2/(1-P_2); && \text{Call}\\ \beta_3 &= \tfrac{1}{2} + P_2/2(1-P_2); && \text{Drop}\end{aligned}\right\} \quad P_2 \leqq \tfrac{1}{3},$$

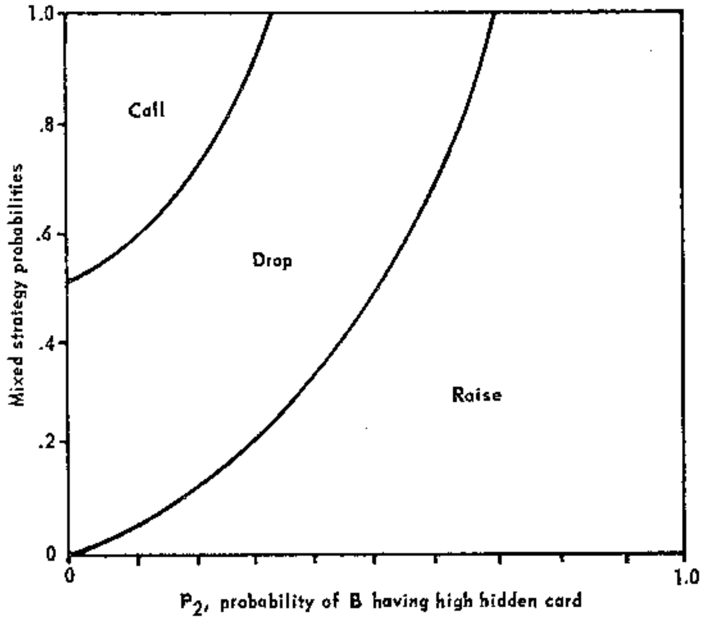

FIGURE 3. Optimal Strategy for Player B

$$\left.\begin{aligned} \beta_1 &= P_2/2(1 - P_2); && \text{Raise} \\ \beta_2 &= 0; && \text{Call} \\ \beta_3 &= 1 - P_2/2(1 - P_2); && \text{Drop} \end{aligned}\right\} \tfrac{1}{3} < P_2 \leqq \tfrac{2}{3},$$

$$\left.\begin{aligned} \beta_1 &= 1; && \text{Raise} \\ \beta_2 &= 0; && \text{Call} \\ \beta_3 &= 0; && \text{Drop} \end{aligned}\right\} \tfrac{2}{3} < P_2 .$$

In almost all real poker situations of this type both P_1 and P_2 will be less than $\frac{1}{3}$. B's mixed strategy as a function of P is shown in Figure 3.

B's strategy has some surprising features. For $P_2 < \frac{1}{3}$ the more threatening B's hand the higher the probability he will drop against A's raise. Of course, B's raise probability will increase as his hand becomes more threatening.

Using the optimal strategies we find the region $P_2 \leqq \frac{1}{3}$.
A's Expected Gain $= \frac{3}{2}P_1[1 - 3P_2]$,
B's Expected Gain $= 1 - \frac{3}{2}P_1[1 - 3P_2]$.

In the previous example where P_2 was zero and B had no reraise possibilities, we found A's expected gain was $\frac{3}{2}P_1$ with optimal strategies. Thus the probability P_2 of B also having an improved hand and his possibility of bluffing has reduced A's expected gain by the fraction $(1 - 3P_2)$.

If A and B are using the optimal strategy $\frac{1}{3}$ of their raises will be on bluffs.

Conjectures on Bluffing

In this paper we have calculated exact bluffing strategies in two common poker situations. In both examples the bluffer bluffs with probability $P/2(1 - P)$ or $\frac{1}{2}$ the inverse of the 'odds' that he has the winning hand. The gambler's 'odds' for any gambling situation is the probability the bettor loses, $(1 - P)$, divided by the probability the bettor wins, P.

In both situations the optimal strategy was to call a potential bluff half the time At this point we would conjecture that these fundamental optimal bluffing probabilities would occur in other poker situations which may be more complicated but in which a pure bluff was possible and pot limit was played.

References

1. GILLIES, MAYBERRY AND VON NEUMANN, "Two Variants of Poker," *Annals of Mathematical Statistics*, No. 28 (1953), Princeton University Press, pp. 13–50.
2. KARLIN, S. AND RESTREPO, RODRIGO, "Multistage Poker Models," *Annals of Mathematical Studies*, No. 39 (1957), Princeton University Press, p. 337.
3. KUHN, H. W., "A Simplified Two-Person Poker," *Annals of Mathematics Studies*, No. 24 (1950), Princeton University Press.
4. NASH, J. AND SHAPLEY, L., "A Simple Three-Person Poker Game, "*Annals of Mathematics Studies*, No. 24 (1950), Princeton University Press.