Game Theory and Poker

Jason Swanson

April 3, 2005

Abstract

An extremely simplified version of poker is completely solved from a game theoretic standpoint. The actual properties of the optimal solution are then compared to various common notions about game theory and bluffing in poker. No specialized mathematical knowledge is needed to understand this material.

1 Introduction

Let us consider a simple card game which we will call "One Card Poker." The game is played with two players and a 3-card deck. The deck consists of an ace, a deuce, and a trey. To begin the game, one of the players is chosen to be the dealer; the other will be called the opener. After the dealer is selected, each player antes $100, forming a $200 pot. Then, the dealer deals one card to each player.

After the players look at their cards, the opener is the first to act. He may check or bet $100. If he bets, then the dealer may either call or fold. If the dealer folds, the opener takes the pot; if the dealer calls, there is a showdown and the high card takes the pot. The ace is considered the lowest card.

If the opener begins the game with a check, then the dealer may either check or bet $100. If the dealer checks, there is a showdown; otherwise, the opener must either call or fold. Note that there is no raising in this game.

With the rules of One Card Poker in place, let us consider the following situation: the dealer is dealt the deuce and the opener bets. The dealer's "hand" can beat only a bluff. What is the game-theoretic optimal frequency with which the dealer should call in this situation?

In *The Theory of Poker*, David Sklansky discusses using game theory to call a possible bluff. He writes,

> Usually when your hand can beat only a bluff, you use your experience and judgment to determine the chances your opponent is bluffing . . . However, against an opponent whose judgment is as good as yours or better than yours, or one who is capable of using game theory to bluff, you in your turn can use game theory to thwart that player or at least minimize his profits.

He then gives the following example:

> If your opponent bets \$20 to win \$100, he is getting 5-to-1 on a bluff. Therefore, you make the odds 5-to-1 against your folding. That is, you must call five times and fold once.

Motivated by this, let us return to our example. The dealer has the deuce and can beat only a bluff. The opener has bet \$100 into a \$200 pot, giving him 2-to-1 on a bluff. We might therefore think that the dealer should call twice and fold once. Or, in other words, the dealer should call with his deuce with probability $\frac{2}{3}$. As we will see, however, this is wrong. The game-theoretic optimal frequency with which the dealer should call with the deuce is $\frac{1}{3}$.

Another interesting example is the following: the dealer is dealt the ace and the opener checks. Should the dealer bluff? If the opener holds the trey, then a bluff is futile, but if the opener holds the deuce, then a bluff might work. Suppose that the dealer knows that the opener will *always* come out betting when he holds the trey. Then since the opener checked here, the dealer can be certain that the opener holds the deuce. In this situation, what is the game-theoretic optimal frequency with which the dealer should bluff?

Again, let us return to the examples in *The Theory of Poker*:

> when I bet my \$100, creating a \$300 pot, my opponent was getting 3-to-1 odds from the pot. Therefore my optimum strategy was ... [to make] the odds against my bluffing 3-to-1.

Since the dealer will always bet with the trey in this situation, he should bluff with the ace $\frac{1}{3}$ of the time in order to make the odds 3-to-1 against a bluff.

But the analysis in this example began with the supposition that *the opener will always bet when he holds the trey.* What if the opener is a tricky player who will sometimes check with the trey, trying to goad the dealer into bluffing? Or, for that matter, what if the opener has the peculiar habit of *never* betting when he holds the trey? If he also checks whenever he holds the deuce, then what can the dealer conclude? The dealer is holding the ace and the opener has checked. Half the time, the opener will have the trey and a bluff is futile. The other half of the time, the opener will hold the deuce, and $\frac{1}{3}$ of those times, the dealer ought to bluff. So perhaps, against this opponent, the optimal frequency with which the dealer should bluff is $\frac{1}{6}$.

But again, as we will see, this is wrong. The dealer's game-theoretic optimal bluffing frequency is $\frac{1}{3}$, *regardless* of how often the opener checks with the trey. In fact, if the dealer tries bluffing only $\frac{1}{6}$ of the time, there is a strategy that the opener can employ – a strategy which involves *always checking with the trey* – that capitalizes on this mistake.

Finally, let us consider one last example before solving One Card Poker and answering all our questions. Consider the situation in which the opener is dealt the deuce. He checks and the dealer bets. His hand can beat only a bluff. What is the optimal frequency with which he should call this bet? Should he call $\frac{2}{3}$ of the time, as suggested by the excerpt from *The Theory of Poker*? Or should he call $\frac{1}{3}$ of the time, as he would if faced with the same situation as the dealer?

As we will see, there is no unambiguous answer to this question. For example, it may be optimal for the opener to call here $\frac{2}{3}$ of the time, but only if he never checks with the trey and bluffs with the ace $\frac{1}{3}$ of the time. On the other hand, it may be optimal for the opener to call in this situation, say, 19 times out of 30, but only if he checks with the trey 10% of the time and bluffs with the ace 30% of the time. So we see here a somewhat counter-intuitive example where the optimal play for the opener depends not only on the current situation, but also on how he *would have* responded if the situation were different. After all, if this one hand is the only hand these two players will ever play with each other, then why should it matter that the opener intended to check with the trey 10% of the time? He was never and will never be dealt a trey!

2 Obvious Plays and Stupid Mistakes

We now want to analyze One Card Poker using game theory. What this means is that we will consider a group of strategies for the opener and a group of strategies for the dealer. From these groups, we will try to find strategies which are "optimal" in a certain sense. But it would definitely be a waste of time if any part of that analysis were devoted to concluding that one should never fold a trey, or that one should never call a bet with an ace. For this reason, let us devote this section to enumerating the Stupid Mistakes in One Card Poker. We will thereafter assume that the players do not make these Stupid Mistakes. This will allow us to reduce the sizes of the groups of strategies we must consider.

Stupid Mistake No. 1: Folding the trey. We will assume that neither player will ever fold the trey.

Stupid Mistake No. 2: Calling with the ace. We will assume that neither player will ever call a bet while holding the ace.

Stupid Mistake No. 3: Checking with the trey "in position." We will assume that if the dealer holds the trey and the opener checks, then the dealer will automatically bet.

Stupid Mistake No. 4: Betting with the deuce. If a player bets with the deuce, then according to the other Stupid Mistakes, his opponent will fold the ace and call with the trey. If the dealer holds the deuce and is checked to, then checking and betting have the same effect when the opener holds the ace, but betting loses an additional $100 when the opener holds the trey. Similarly, when the opener holds the deuce, checking with the intention of calling a bet has the same effect as betting when the dealer holds the trey, but wins additional money when the dealer holds the ace and decides to bluff. So there is always a better option than betting the deuce. Betting the deuce is a no-win situation for the bettor, and we will assume that neither player will bet the deuce.

3 Strategic Plays and Expected Value

Given that the players will flawlessly avoid the Stupid Mistakes, they now have a limited number of choices. The dealer must decide how frequently he will bluff with the ace when the opener checks, and he must decide how frequently to call with the deuce when the opener bets. Let us denote these probabilities by q_1 and q_2, respectively.

As for the opener, he has three decisions. He must decide how frequently to bet out as a bluff when he holds the ace. Call this p_1. He must decide how frequently to call with the deuce when the dealer bets. Call this p_2. And finally, he must decide how frequently to bet out with the trey. Call this p_3. Note that it is not an obvious play to choose $p_3 = 1$. If, for example, the opener held the trey, and he somehow knew that the dealer held the ace, then the only correct play would be to check.

To summarize, p_1 is the probability the opener bluffs with the ace, p_2 is the probability the opener calls with the deuce, p_3 is the probability the opener bets with the trey, q_1 is the probability the dealer bluffs with the ace, and q_2 is the probability the dealer calls with the deuce.

We now want to use these values to compute the opener's *post-ante* expected value (EV). What this means is that we will regard the \$200 pot, which was created by the antes, as belonging to neither player. So, for example, if the opener checks, the dealer bets, and the opener folds, then we regard this as a \$200 win for the dealer and a \$0 win (or loss) for the opener. After computing the opener's post-ante EV, it will be a simple matter to convert this into a pre-ante EV; that is, the opener's EV for the entire hand, including his original \$100 ante.

There are three possible non-zero post-ante results for the opener. Either he loses \$100, wins \$200, or wins \$300. We will begin by computing the probabilities of each of these outcomes.

Case 1: the opener has the ace, the dealer has the deuce. In this case, the opener loses \$100 if he bluffs and is called. This happens with probability p_1q_2. He wins \$200 if he bluffs and the dealer folds. This has probability $p_1(1-q_2)$. He cannot win \$300.

Case 2: the opener has the ace, the dealer has the trey. Here, the opener can only lose \$100, which happens whenever he tries to bluff. The probability of this is p_1.

Case 3: the opener has the deuce, the dealer has the ace. In this case, the opener checks. With probability $1-q_1$, the hand is checked through and the opener wins \$200. With probability q_1p_2, the dealer bluffs and gets called, and the opener wins \$300. The opener cannot lose \$100.

Case 4: the opener has the deuce, the dealer has the trey. The opener cannot win anything. He loses \$100 when he calls the dealer's bet with probability p_2.

Case 5: the opener has the trey, the dealer has the ace. Here, the opener wins \$300 when he checks and the dealer bluffs. This has probability $(1-p_3)q_1$. Otherwise, with probability $1-(1-p_3)q_1$, the opener wins \$200.

Case 6: the opener has the trey, the dealer has the deuce. Here, the opener wins \$300 when he bets and the dealer calls. This has probability p_3q_2. Otherwise, with probability $1-p_3q_2$, the opener wins \$200.

Since each of these cases has probability $\frac{1}{6}$, we can combine them to see that the opener loses \$100 (cases 1, 2, and 4) with probability

$$Q_1 = \frac{1}{6}(p_1q_2 + p_1 + p_2), \tag{1}$$

he wins \$200 (cases 1, 3, 5, and 6) with probability

$$\begin{aligned} Q_2 &= \frac{1}{6}(p_1(1-q_2) + 1 - q_1 + 1 - (1-p_3)q_1 + 1 - p_3q_2) \\ &= \frac{1}{6}(p_1 - p_1q_2 + 3 - 2q_1 + p_3q_1 - p_3q_2), \end{aligned} \tag{2}$$

and wins \$300 (cases 3, 5, and 6) with probability

$$\begin{aligned} Q_3 &= \frac{1}{6}(q_1p_2 + (1-p_3)q_1 + p_3q_2) \\ &= \frac{1}{6}(p_2q_1 + q_1 - p_3q_1 + p_3q_2). \end{aligned} \tag{3}$$

This gives him a post-ante EV of

$$-Q_1 + 2Q_2 + 3Q_3,$$

where one unit of EV represents a \$100 average win per hand.

If we want to include his ante, then, in each of these cases, we must reduce his final outcome by \$100. Also, we must acknowledge that, with probability $1-(Q_1+Q_2+Q_3)$, he loses his original ante of \$100. This gives him a total EV for the entire hand of

$$-2Q_1 + Q_2 + 2Q_3 - (1 - Q_1 - Q_2 - Q_3),$$

which simplifies to

$$-Q_1 + 2Q_2 + 3Q_3 - 1. \tag{4}$$

4 Game Theory Analysis

In a way, the idea behind game theory is very basic. In this case, we simply want to understand how the players' choices of the p and q values affect their EVs. To accomplish this, we will need to simplify and rewrite the expressions derived in the previous section. Plugging (1), (2), and (3) into (4), we find that the opener's total EV for the entire hand is

$$\frac{1}{6}(-p_1q_2-p_1-p_2)+\frac{1}{6}(2p_1-2p_1q_2+6-4q_1+2p_3q_1-2p_3q_2)+\frac{1}{6}(3p_2q_1+3q_1-3p_3q_1+3p_3q_2)-1,$$

which simplifies to

$$\frac{1}{6}(-3p_1q_2 + p_1 - p_2 - q_1 - p_3q_1 + p_3q_2 + 3p_2q_1).$$

In order to make use of this, it will be convenient to rewrite this as

$$\frac{1}{6}[p_1(1-3q_2) + p_2(3q_1-1) + p_3(q_2-q_1) - q_1] \tag{5}$$

and also as

$$\frac{1}{6}[q_1(3p_2 - p_3 - 1) + q_2(p_3 - 3p_1) + (p_1 - p_2)]. \tag{6}$$

From (5), we see that something special happens when $q_1 = q_2 = \frac{1}{3}$. In this case, the opener's EV is simply $-\frac{1}{18}$, and this *does not depend* on the opener's choices of the numbers p_1, p_2, and p_3. If the dealer chooses these q-values, then the dealer is *indifferent* as to how the opener plays the game. The dealer's EV will be $\frac{1}{18}$, regardless of how the opener plays. For this reason, we will refer to $q_1 = q_2 = \frac{1}{3}$ as the dealer's "indifferent strategy."

If the dealer deviates from the indifferent strategy, then he is making a "mistake" in the sense that the opener (if he can spot the deviation) can increase his EV above $-\frac{1}{18}$ (and possibly even give himself a positive EV) by choosing the appropriate counter-strategy.

Let us consider some examples. Suppose that the dealer does not bluff enough with the ace, does not call enough with the deuce, and bluffs with the ace more often than he calls with the deuce. In other words, the dealer has chosen a strategy in which $q_2 < q_1 < \frac{1}{3}$. In order for the opener to maximize his EV in this case, he should never call with the deuce (since the dealer does not bluff enough), always bluff with the ace (since the dealer does not call enough), and never bet out with the trey (since the dealer will bluff more often than he will call). In other words, the opener should choose the strategy $p_1 = 1$, $p_2 = 0$, and $p_3 = 0$. From (6), we see that the opener's EV will be

$$\frac{1}{6}[-q_1 - 3q_2 + 1].$$

Since $q_2 < q_1 < \frac{1}{3}$, this EV will always be greater than $-\frac{1}{18}$, and may very well be positive. For example, if $q_2 = \frac{1}{6}$ and $q_1 = \frac{1}{4}$, then the opener's EV will be $\frac{1}{24}$.

As another example, suppose that the opener is the kind of "passive" player that will never bet out with the trey. The dealer sees this and reasons, as in the introduction, that he should only bluff $\frac{1}{6}$ of the time. If he does this, then the opener can counter by simply doing nothing! The opener can simply never bet out and never call, unless he has the trey. In other words, the opener chooses the strategy $p_1 = p_2 = p_3 = 0$. Referring to (5), we see that by doing this, the opener has an EV of $-\frac{1}{36}$, which is better for the opener than $-\frac{1}{18}$. In fact, if the players alternate positions and the passive player uses the indifferent strategy when he is the dealer, then his EV over every pair of hands will be $-\frac{1}{36} + \frac{1}{18} = \frac{1}{36}$ and he will be a long term winner.

The point here is that the dealer wants to bluff in a way that puts the opener "on the edge" and makes it difficult for him to decide whether or not he should call. This is why the bluffing frequency should match the pot odds. But no amount of bluffing is going to make it difficult for him to call with a trey. So the bluffing frequency must be targeted at making his decision with the deuce a difficult one. This is accomplished precisely by making $q_1 = \frac{1}{3}$.

For one final example, suppose that the dealer has read *The Theory of Poker* and has decided that he should bluff with the ace $\frac{1}{3}$ of the time. That is, he chooses $q_1 = \frac{1}{3}$. He also knows that, with the deuce, his hand can beat only a bluff. So when the opener bets \$100 into a \$200 pot, the dealer should make the odds 2-to-1 against his folding. That is, he should call twice and fold once. So he decides upon $q_2 = \frac{2}{3}$.

By (5), we see that the opener's EV is now

$$\frac{1}{6}\left[-p_1 + \frac{1}{3}p_3 - \frac{1}{3}\right].$$

The opener can maximize his EV by choosing $p_1 = 0$ (never bluffing), $p_3 = 1$ (always betting with the trey), and doing whatever he likes with the deuce. By doing this, the opener's EV is zero.

In other words, by calling $\frac{2}{3}$ of the time with the deuce, the dealer has forfeited his $\frac{1}{18}$ EV advantage, which is his natural advantage due to acting last. If the players take turns being the dealer, he could end up a long term loser by forfeiting this advantage every other hand.

The *only* way for the dealer to prevent the opener from being able to seize back some of this advantage is to play the indifferent strategy, $q_1 = q_2 = \frac{1}{3}$. It is for this reason that the indifferent strategy is more commonly referred to as the "optimal" strategy.

As we have seen, in this game it is optimal to call with the deuce $\frac{1}{3}$ of the time, not $\frac{2}{3}$ of the time. So is there a flaw in the examples from *The Theory of Poker*? In fact, there is not. When your opponent bets \$100 into a \$200 pot, you should play in a way that makes the probability you will call $\frac{2}{3}$. But it is not the probability that you will call from *your* perspective which should be $\frac{2}{3}$, but the probability *from your opponent's perspective.* When the opener bets with the ace, he does not know whether the dealer has the deuce or the trey. Half the time, the dealer will have the trey and will automatically call. The other half of the time, the dealer will have the deuce. Of those times, the dealer will call with frequency $\frac{1}{3}$. So the probability that the dealer will call, *from the opener's perspective*, is $\frac{1}{2} + \frac{1}{2} \cdot \frac{1}{3} = \frac{2}{3}$.

This phenomenon might be better illustrated if we were using a 4-card deck: one ace, one deuce, one trey, and one four. In that case, what would be the optimal frequency with which you should call a bet with the deuce? Well, that would be related to how frequently you would call a bet with the trey. In order to play optimally, you would have to play the deuce *and* the trey (and the four as well) in a way which makes the *overall* probability that you will call a bluff $\frac{2}{3}$. But this does not mean you would call specifically with the deuce $\frac{2}{3}$ of the time. In fact, you would probably call much less frequently than that.

5 Multiple Optimal Strategies

We just saw that the dealer has a unique optimal strategy, $q_1 = q_2 = \frac{1}{3}$. The situation for the opener, however, is different. From (6), we see that any set of p-values which satisfies

$$\begin{aligned} 3p_2 - p_3 - 1 &= 0 \\ p_3 - 3p_1 &= 0 \end{aligned}$$

will be an indifferent strategy for the opener and will give him an EV of $\frac{1}{6}(p_1 - p_2)$, regardless of how the dealer plays. In other words, the opener can choose p_3 arbitrarily, and then choose

$$\begin{aligned} p_1 &= \frac{1}{3}p_3 \\ p_2 &= \frac{1}{3}p_3 + \frac{1}{3}. \end{aligned}$$

By doing this, he will assure himself an EV of

$$\frac{1}{6}\left(\frac{1}{3}p_3 - \left(\frac{1}{3}p_3 + \frac{1}{3}\right)\right) = -\frac{1}{18},$$

no matter what strategy the dealer selects. In other words, all such strategies are optimal strategies.

For example, choosing $p_3 = 1$ gives $p_1 = \frac{1}{3}$ and $p_2 = \frac{2}{3}$. So always betting out with the trey, bluffing with the ace $\frac{1}{3}$ of the time, and calling with the deuce $\frac{2}{3}$ of the time is an optimal strategy for the opener. On the other hand, choosing $p_3 = 0$ yields $p_1 = 0$ and $p_2 = \frac{1}{3}$. So it is also optimal for the opener to never bet with the trey, never bluff with the ace, and call with the deuce $\frac{1}{3}$ of the time. For each different value of p_3 between 0 and 1, there is a different optimal strategy that the opener may employ. As is the nature of optimal strategies, each of these guarantees the opener an EV of $-\frac{1}{18}$, regardless of the strategy the dealer chooses. However, if the opener chooses a strategy outside of this group, then the dealer can capitalize on this "mistake."

For example, suppose the opener does not call enough with the deuce; that is, $p_2 < \frac{1}{3}p_3 + \frac{1}{3}$. By (6), we see that the dealer can capitalize on this by choosing $q_1 = 1$ (always bluffing with the ace). By (5), if the dealer chooses $q_2 = \frac{1}{3}$ (calling $\frac{1}{3}$ of the time with the deuce), then he nullifies the effect of p_1 on his EV, allowing him to isolate the opener's mistake. With these choices, using (5), the opener's EV is

$$\frac{1}{6}\left[2p_2 - \frac{2}{3}p_3 - 1\right] < \frac{1}{6}\left[2\left(\frac{1}{3}p_3 + \frac{1}{3}\right) - \frac{2}{3}p_3 - 1\right] = -\frac{1}{18}.$$

6 Paradoxes

Let us return now to the examples in the introduction. Consider the situation where the opener is dealt the deuce and checks. The dealer bets. What is the opener's optimal play? Suppose he has a perfect random number generator and can call with any probability he likes. Which probability is optimal?

As we have seen, any probability p_2 of the form $\frac{1}{3}p_3 + \frac{1}{3}$ is optimal, provided he bets out with the trey with probability p_3 and bluffs with the ace with probability $\frac{1}{3}p_3$. In other words, it will be optimal for the opener to call here with any probability between $\frac{1}{3}$ and $\frac{2}{3}$, provided he plays the ace and trey in a manner which is consistent with that choice.

To some degree, this flies in the face of conventional wisdom. We know the opener's card: the deuce. We know the betting sequence so far: check, bet. We can even know everything there is to know about his opponent's tendencies. That is, we can know the precise values of q_1 and q_2. And still we do not have enough information to answer the question, "What is the opener's optimal play?" His optimal play depends not only on his cards, the betting, and his opponent. It also depends on how he intends or intended to play in other situations.

This seems very counter-intuitive. A play has either a positive expectation, a negative expectation, or zero expectation. Other hands, at least in our mathematical model, are independent of this one. How can the opener's actions in other hands affect the value of his chosen action in this hand? How could these other actions play any role in determining the best play here?

But therein lies the catch. At no time until now did we ever ask, "What is the *best* play?" We only asked, "What is the *optimal* play?" So let us now address the *best* play. If we know everything about our opponent, then we know the value of q_1. Looking at (5), if $q_1 > \frac{1}{3}$, then the opener maximizes his EV by taking $p_2 = 1$. In other words, the *best* play is to call. If $q_1 < \frac{1}{3}$, then the opener maximizes his EV by taking $p_2 = 0$. In other words, the *best* play is to fold. If $q_1 = \frac{1}{3}$, then (5) shows us that it *does not matter* what p_2 is. The opener gets the same EV no matter what frequency he calls with. So *all* plays are "best" plays.

We see, then, that there is no contradiction to common sense after all. The best play can in fact be determined from our cards, the betting, and our opponent's tendencies, without having to consider what we will do in other, unrelated situations. But to come to this realization, we must acknowledge the difference between what is "optimal" and what is "best." When your opponent plays optimally, then it doesn't matter what you do. You will always have the same EV, so all plays are "best." But when your opponent does not play optimally, the best play for you – the play which maximizes your expectation – will not be optimal either.

We will return to this idea in a moment, but for now let us move in another direction with this example. For the moment, we will set aside the question of what is best and return to the problem of determining what is optimal. We have seen that, as the dealer, when you hold the deuce and your opponent bets, you should play in a way that gives a $\frac{2}{3}$ probability that you will call, with the probability calculated *from your opponent's perspective*. Does the same principle hold for the opener? It seems reasonable to assume that it does. So suppose the opener holds the deuce. He checks and the dealer bets. Since the opener would always call here with the trey, how could it ever be optimal for him to call with the deuce more often than $\frac{1}{3}$ of the time?

For instance, it would be optimal for the opener to choose

$$p_3 = 0.9 \qquad\qquad p_1 = 0.3 \qquad\qquad p_2 = 0.3 + \frac{1}{3} = \frac{19}{30}.$$

In other words, he checks with the trey 10% of the time, bluffs with the ace 30% of the time, and calls with the deuce 19 times out of 30. So when the dealer bluffs with the ace, the opener will call every time he has the trey and 19 out of every 30 times that he has the deuce. If he has the trey half the time and the deuce half the time, then he is calling the dealer's bluff 49 times out of 60, which is more than the requisite $\frac{2}{3}$. Since the dealer is only getting 2-to-1 odds on his bluff, it seems his bluff has a negative EV. (Out of 60 bluffs, 49 times he loses \$100, and 11 times wins \$200, for a net profit of $-4900 + 2200 = -2700$ dollars.) Hence, the less frequently the dealer bluffs, the higher his EV. But the opener's strategy is supposed to be optimal. So the dealer's bluffing frequency cannot affect his EV. What went wrong?

The key to resolving this apparent contradiction is to fall back on the phrase, "from your opponent's perspective." From the dealer's perspective, the opener does *not* have the trey half the time and the deuce half the time. The opener was first to act and he checked. This check gave the dealer information. If the dealer knows the opener's frequencies, then he knows that the opener will check with the trey 10% of the time. So out of every 20 times that the dealer has the ace, the opener will have the deuce 10 times and check all 10 of those times, and he will have the trey 10 times and check only once. So out of 11 checks, he has the deuce 10 times. Therefore, from the dealer's perspective, there is a $\frac{10}{11}$ probability that the opener has the deuce. If the opener has the deuce, then he will fold to a bluff 11 times out of 30; that is, with probability $\frac{11}{30}$. In other words, from the dealer's perspective, *given that the opener checked*, he will fold to a bluff $\frac{10}{11} \cdot \frac{11}{30} = \frac{1}{3}$ of the time, which is exactly what he should be doing.

So using game theory to call a possible bluff is a subtle thing indeed. If your opponent bets \$100 into a \$200 pot, you should play in such a way that he thinks there is a $\frac{2}{3}$ chance you will call. But this does not mean you simply call $\frac{2}{3}$ of the time with any marginal hand. As we saw before, since your opponent does not know what you have, you need to play *all* your individual hands in such a way that the *overall* probability you will call is $\frac{2}{3}$. And, as this example shows, you must also take into account how the information at your opponent's disposal affects his estimations of the probabilities that you have each of these individual hands.

7 How To Win

Perhaps the most important lesson to take from all of this analysis is the following:

You cannot win with the optimal strategy.

If you always play the optimal strategy, then as the opener you will have an EV of $-\frac{1}{18}$ no matter how your opponent plays, and as the dealer you will have an EV of $\frac{1}{18}$ no matter how your opponent plays. In the long run, you will only be a break-even player. If you use the optimal strategy, your opponent cannot profit through superior play. But he also cannot suffer through inferior play. By playing "optimally" you have created a situation in which your opponent's choices, good or bad, will have no effect on your EV. Clearly, by symmetry, you cannot do this and win. (Unless your opponent lets you always play as the dealer.)

So the object of the game is not to play optimally. It is to spot the times when your opponent is not playing optimally, or even to induce him not to play optimally, to recognize the *way* in which he is deviating from optimality, and then to choose a *non*-optimal strategy for yourself which capitalizes on his mistakes. You must play non-optimally in order to win. To capitalize on your opponent's mistakes, you must play in a way that leaves you vulnerable.

For instance, your opponent may be bluffing too much. To capitalize on this, you begin to call more frequently than is optimal. Once you do this, however, your opponent could stop bluffing altogether and take advantage of you. When you realize he has done this, you would start calling much less frequently than is optimal. In this way, you and your opponent's bluffing and calling frequencies would oscillate, sometimes higher than optimal, sometimes lower.

In game theory, an optimal solution is also called an "equilibrium." The idea is that, through this back-and-forth struggle, the players would eventually settle upon the optimal frequencies and reach an equilibrium with one another. While this might be true in certain "real-world" situations (in politics or economics, for example), it is certainly not true in One Card Poker. An "expert" One Card Poker player would rather quit playing altogether than settle for the monotonous compromise of playing a zero EV optimal strategy. A battle between two One Card Poker experts would *not* be a battle in which both sides played optimally. Such a battle would be a complete waste of time. Rather, it would be a back-and-forth struggle like the one described above; a struggle which never slows down and never reaches equilibrium.

The heart of the game is the struggle. Playing optimally erases this struggle. Playing optimally prevents your opponent from taking advantage of you, but it also prevents him from being punished for his mistakes. As such, using game theory to "optimally" bluff or to "optimally" call a bluff can only be regarded as a *defense*. But since it defends both you *and* your opponent, a better defense is to simply not play at all.

That being said, it should be pointed out that there *are* times when the optimal strategy will be profitable; namely, when your opponent makes Stupid Mistakes. In that case, you can play optimally and be a long term winner. By playing optimally, you ensure that changes in his bluffing and calling frequencies will not affect his EV. Since his EV will be intrinsically negative due to his Stupid Mistakes, you will have a positive expectation. (On the other hand, if he is making Stupid Mistakes, then you can probably outwit him without game theory.)

When your opponent does not make Stupid Mistakes, playing optimally is an exercise in futility. But nonetheless, there is still value in understanding the theoretical aspects of optimal play. In order to profit, you must know, for example, when your opponent is bluffing too much. But what does it mean to bluff "too much" in a situation. It means, of course, bluffing more than is optimal. So you must know what the optimal strategy is (even though you will consciously avoid it) in order to decide on the proper counter-strategy against your opponent.

References

[1] David Sklanksy, *The Theory of Poker*. Two Plus Two Publishing, 1989