To Bluff or Not to Bluff: Principles and Practice From Laboratory Experiments

By

David Yaffe*

Presented to the Department of Economics, University of Oregon,
in partial fulfillment of requirements for honors in Economics.

under the supervision of
Prof. Bill Harbaugh

June 2006

* I would like to specially thank Bill Harbaugh who without his helpful comments, suggestions, and support this paper would not have been possible. I'd also like to thank Sean Wallace who without his excellent programming skills, this paper would not have been possible. All errors are my own responsibility.

To Bluff or Not to Bluff: Principles and Practice From Laboratory Experiments

Abstract: Truth-telling behavior in university students (n=38) was examined using a simple stylized poker experiment. Four treatments were applied that changed either card probabilities or payoff structure. Results indicated that in some cases the students' behavior conformed closely to theoretical predictions. However, other cases indicated robust inconsistencies between observed behavior and mixed-strategy Nash predictions. These findings suggest that there is a strong relationship between the propensity to tell the truth and the costs and rewards thereof. A link between calling behavior and costs and rewards was also found. Implications for the justice system are discussed.

Approved: ____

Prof. Bill Harbaugh   Date

1. Introduction

There are many good reasons to study lying and deception among children. Everyday, parents are put into situations that involve trusting their children. When a child wants to go play outside with the neighborhood kids, a frequent pre-requisite is to insure that they have finished with their household responsibilities such as brushing their teeth, making their beds, cleaning their rooms and so forth.

While it could be argued that the above reasons are somewhat trivial, there are more pressing issues that involve a child's trust such as court testimony. With a substantial and increasing number of young children testifying in North American courts, researchers are inquiring about the accuracy of their testimonies. Children's truthfulness and their understanding of its benefits and consequences of lying are critical for ensuring the viability and value of their testimony. Within the last two decades, extensive theoretical research has been done on children's conceptual understanding of lie-telling and truth-telling, however, very little empirical evidence exists which addresses a child's willingness and ability to make risky decisions based on bluffing and truth-telling in a legal and economic framework. Additionally, little is known about how factors such as payoffs and detection probabilities influence children's propensity to tell the truth. It is that deficit that has inspired my research endeavors.

In this paper I report on the laboratory results of a study that uses a simplified version of poker to examine the strategic behavior of college students, their propensity to bluff and tell the truth and to what extent their behavior is consistent with game theoretical predictions. The ultimate goal of this project is to extend it to younger children between the ages of ten and thirteen. However, the subjects that I report on are only a few years past the age of legal majority (the age at which one acquires full legal rights of an adult).[1]

The game of poker is a simple, stylized representation of some important elements of truth telling and I adapt a version of poker created by Reiley, Urbancic, and Walker (2005) that strips the conventional game of poker to its fundamentals. First, there are information asymmetries where one player has complete information about his card and the other player's card, but the other player only knows his own card. Like a child who knows whether or not they have made their bed but their parent hasn't a clue. Second, the game allows for opportunities to bluff and either get away with it, or get called by the other player. Similarly, a child who hasn't made their bed may report otherwise and can get a way with it if their parent trusts them, or can get called if their parent checks their room. For these reasons stated above, I feel that poker is a plausible representation of some significant elements of truth telling. However, it is also true that not all aspects of truth-telling are captured by this experiment, such as the moral implications of lie-telling.

For this game, two students are paired anonymously and assigned the role of player 1 or player 2. There are three possible cards- A, B, or C- in which A > B > C. Player 1 knows not only his own card, but also the card of player 2. Player 2 only knows his own card and hence has less information than player 1. After observing his card, player 1 then moves first by either betting or folding. If he folds, then the game is over. If he bets, then player 2 can either call or fold. A successful bluff requires two things. Player 1 must bet with card C and player 2 must fold.

[1] The age of legal majority varies state by state in the U.S. Most states give full rights by the age of twenty-one. The age of students who participated ranged from 19-55, though only a few outlier students were above the age of 23.

The experiment was conducted on a total of 38 university students ranging from sophomore to senior status. Four different treatments were applied. Almost all of the students acted strategically by mixing between two strategies: That of betting all the time regardless of their card (Bluffing), and folding when given the weak card (Truth-telling). I found the students were not behaving consistent with game theoretical predictions during some treatments, but their behavior conformed remarkably closely to theoretical predictions under other treatments. There were instances where the students were playing close to their optimal strategies given what the other player was doing. Interestingly, the students significantly, albeit not perfectly, adjusted their behavior in a direction consistent with theory. For example, the third treatment gave a high probability to receiving the losing card, making the optimal strategy to bluff very little of the time. That said, the students bluffed significantly less under treatment 3 than the treatment 1, where the odds of getting both the winning and losing card were equal.

I conclude that college students are not the best poker players. In most cases, they were bluffing either too much when in principle they should have been bluffing less, and bluffing too little when in principle they should have been bluffing more. However, their choices showed a clear understanding that some level of bluffing is optimal and they knew when to change their propensities to tell the truth given changes in the game. For example, when the costs of bluffing were increased, students responded by truth-telling significantly more and almost exactly equal to the optimal amount. These results are quite provocative for they suggest to some degree that incentives and consequences change the propensity to tell the truth.

This paper begins with a brief literature review, largely from psychology, that reports on the surrounding issues of adolescent testimony, surveys previous studies that examined the age at children develop the capacity to engage in deceitful behavior and their understanding of the morality of truth and lie-telling. This is followed by a description of the poker game created by Reiley et al and a walkthrough of the process of determining the mixed-strategy Nash equilibria of the four different treatments used in this study. Results are then presented comparing observed behavior with theoretical predictions followed by a conclusion with suggested policy considerations.

2. Literature Review

By the 1980's, society's sensitivity to and reaction from problems such as abuse and violence suffered by children had changed, and as a result, states revised their criminal procedures to deal more effectively with adolescent victims and defendants (Bruck, Ceci, & Hembrooke, 1998). This led to important changes in the legal system not only in the United States but also in other countries in the Western world (Davies, Lloyd-Bostock, McMurran, & Wilson, 1995). Relaxations of standards that had prevented many children from testifying were perhaps the most important among these changes (Bruck, Ceci, & Hembrooke, 1998).

Children are now increasingly being called to the witness stand to testify for various reasons including alleged child abuse and other serious transgressions committed by their parents or other adults (Bala, Lee, Lindsay, Talwar, 2000). Because of this increase, researchers have been looking for ways to determine the age a child develops the capacity to rationally deceive others. Both justice system professionals and forensic psychologists are interested in whether or not children understand the importance of telling the truth and what measures if any can help facilitate more honesty among children (Bala, Lee, Lindsay, Talwar, 2004).

Within the last couple of decades there has been a plethora of research on the conceptual understanding and moral judgments of lie-telling and truth-telling (e.g., Bussey, 1992, 1999; Lee, Cameron, Xu, Fu, & Board, 1997; Peterson, 1995; Peterson, Peterson, & Seeto, 1983; Siegal & Peterson, 1998;Wimmer, Gruber, & Perner, 1984; for review, see Lee, 2000). Many of the studies on lie-telling are focused on resolving the theoretical debate in regards to the development of morality. Research to date suggest that children's conceptual knowledge of lie-telling and truth-telling develops as early as pre-school (e.g., Bussey, 1992, 1999; Peterson, 1995).

Bussey (1999) investigated the ability of children as young as four to categorize three different types of intentionally false and true statements as truths and lies. Results revealed that older children were better able to categorize false statements as lies and true statements as truths than were the four-year-olds.

Lewis et al (1989) looked at deception in children not quite 3 years old. The experimenter placed a toy behind a child and instructed the child not to peek at the toy when he left the room but informed the child that she could play with it at a later time. The experimenter left and the child remained in the room for 5 minutes if she did not look at the toy, or until she turned around and looked. As soon as the child looked, the experimenter returned, stared at the child, and then asked the child if they peeked. They found that the majority of the subjects peeked when left alone and of the subjects that peeked, 38 per cent told the truth and said they had looked. Polak and Harris (1999) adopted the Lewis paradigm and found similar results. It has also been established beyond a reasonable doubt that children as early as 2 ½ already practice a variety of deceptive strategies that necessarily presuppose a working knowledge of false beliefs (Chandler, Fritz, and Lee 1989).

Bala et al. (2004) took this research further by examining children's lie-telling behavior to conceal a transgression not of their own, but of a parent. Their findings suggest that children were sensitive to the different conditions in which they were asked to conceal their parent's transgression, and responded by adjusting their lie- and truth-telling behavior.

3. Methodology

As mentioned before, the game of poker captures some key elements of lie- and truth-telling and is simple enough for kids to understand. The bluff and fold serve as indicators of deception and honesty respectively and the call serves as a measure of detection. Modeling these interactions is nothing new.

Both John von Neumann (1944), the father of game theory, and French mathematician Émile Borel formulated models that served to illustrate the rationality of bluffing in poker.[2][3] Both Borel's and von Neumann's models are comprised of risk-neutral players and a continuum of possible hands for simplicity.

Reiley, Urbancic, and Walker (2005) provide an even simpler model of poker, coined Stripped-down Poker, which I adopt with a few adjustments.[4] My model

[2] von Neumann and Morgenstern (1944)
[3] Borel (1938)
[4] The majority of the content in this section and the Theory section comes from Reiley et al (2005) with the exception of the various treatments. I am very grateful for the foundation these authors have laid before me.

functions identically to the "Stripped-down" version although I change the names of the game, the actions, the payoff structure, and the type cards given to the players to make it rated PG.

The game consists of two players with two actions. Player 1 can either bet or fold and player 2 can either call or fold. Before the game can be played, each player must ante one dollar into the pot. Once that is completed, player 1 and only player 1 will receive a randomly drawn card with either a king or queen on it. Player 1 will privately observe the card, and decide whether to bet or fold. If player 1 folds, the game ends and player 2 gains his ante in addition to player 1's ante. If player 1 bets, he must place an additional dollar into the pot after which player 2 decides whether to fold or call. If player 2 folds, the game ends and player 1 receives his two dollars in addition to player 1's ante. If player 2 calls, then he is required to add an additional dollar into the pot, and player 1 must show his card. Player 1 wins the pot with the king and loses the pot with the queen. At this point, the pot consists of four dollars, so the winner earns two dollars from the loser.

My adapted version changes a few subtleties of the "Stripped-down" game so that the children are distanced from the idea that they are engaging in gambling activities. This is done for two reasons. First, I would hate to think that a child's parents would perceive this research to be corrupting. Secondly, presenting the game as a poker match may lead to unexpected and uncontrolled variations in behavior. For example, some children (boys) may be familiar with poker and be experienced bluffers in the context of poker.

My modified version uses the words quit, continue, and challenge in replace of fold, bet, and call respectively. Instead of drawing a card with either a king or queen, I use cards with either an A or B. Player 2 is told that he always has card C. Card A beats card B, and card B beats card C. Thus card A becomes equivalent to a king and card C becomes equivalent to a queen. In Stripped-down Poker, the maximum a player can end up losing in one game is two dollars. We change the payoff structure so that instead of losing two dollars, the player would win 1. Essentially we just add 3 points to every possible outcome. In theory, this has no effect on the mixed-strategy Nash equilibrium.

Initially, paper forms were used to conduct the experiment on university students. This method works fine, but programming the game for use with hand held computers is an extraordinarily more efficient way of administering the experiment for reasons that are nine fold. Screen shots of the game can be found in the appendix.

4. Theory

The subsequent sections will guide you through the process of finding the optimal strategy or mixed-strategy Nash equilibrium of the game. Ultimately, the mixed-strategy Nash equilibrium of the baseline treatment is the action set where player 1 bluffs one-third of the time and truth tells two-thirds of the time while player 2 calls two-thirds of the time and folds one-third of the time. Reiley et al (2005) provided much of the theory following section is based on. I've added three additional treatments and provided an explanation of their respective mixed-strategy Nash equilibria.

4.1 Extensive and Strategic Form

Each player moves sequentially, so the game is illustrated first in extensive form. Notice in figure 1 below that there are two decision nodes in player 2's information set. This makes it impossible to find a solution to the game using backwards induction.

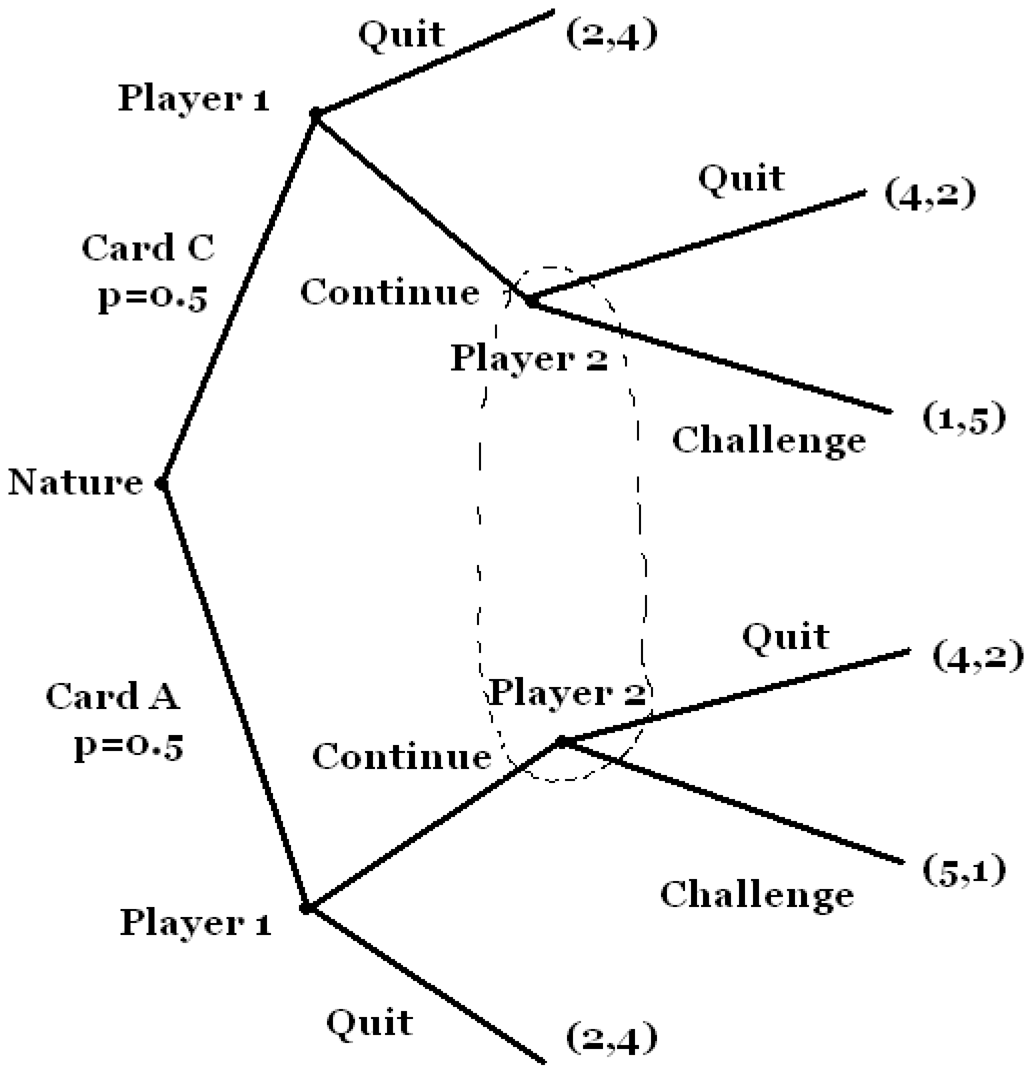

Figure 1. Stripped-down poker for kids in extensive form.

Thus, we naturally want to convert this game into its strategic form and look for a Bayesian Nash equilibrium using a game matrix. Before we can do this, we have to find each player's strategies.

Player 1 has four possible strategies:

1. Continue with A and Continue with C
2. Continue with A and Quit with C
3. Quit with A and Continue with C
4. Quit with A and Quit with C

Player 2 has only two strategies:

1. Challenge
2. Quit

To convert the game into strategic form using a 4x2 matrix, we must use expected payoffs because determining each players' payoff from a strategy requires calculating the weighted average of payoffs in each state of the world. Below is the resulting 4x2 matrix:

Table 1. Stripped-down poker for kids in strategic form.

		Player 2	
		Challenge	Quit
Player 1	CC: Continue if Card A, Continue if Card C	3,3	4,2
	CQ: Continue if Card A, Quit if Card C	3.5,2.5	3,3
	QC: Quit if Card A, Continue if Card C	1.5,4.5	3,3
	QQ: Quit if Card A, Quit if Card C	2,4	2,4

Notice that for player 1, CC strictly dominates QC and QQ. This suggests that regardless of what player 2 does, player 1 is always better off playing CC instead of QC or QQ. That being said, we can purge both those strategies from the matrix leaving a 2x2 matrix.

Table 2. Dominated strategies eliminated and best responses underlined.

		Player 2	
		Challenge	Quit
Player 1	CC (Bluff)	3,3	4,2
	CQ: (Truth-tell)	3.5,2.5	3,3

Notice that there is no pure strategy Nash equilibrium. Naturally, this leads us to find the mixed-strategy Nash equilibrium.

4.2 Mixed-Strategy Nash Equilibrium

We let p be the probability player 1 plays CC (Bluffing)
We let q be the probability player 2 plays Challenge.

In order for player 1 to mix up his strategies, he must be indifferent between them. Otherwise player 1 would always prefer one strategy, but we know that this is not possible because there is no pure-strategy Nash equilibrium. To keep both players indifferent, we have the following two equations:

(1.1) $$3\times p+2.5\times(1-p)=2\times p+3\times(1-p)\Rightarrow p=1/3$$

(1.2) $$3\times q+4\times(1-q)=3.5\times q+3\times(1-q)\Rightarrow q=2/3$$

Hence the mixed-strategy Nash equilibrium is:

(1/3 **CC** + 2/3 **CQ**, 2/3 **C** + 1/3 **Q**)

More intuitively we have:

(1/3 **Bluff +** 2/3 **Truth-tell**, 2/3 **Call** + 1/3 **Fold**)

4.3 Treatments

The preceding game was used as a base treatment for the experiment. After playing the game with the base parameters, three slightly different treatments were administered that varied the card probability and payoff structure. Incentives and costs were thus changed to determine whether or not behavioral changes were induced. In other words, we wanted to see if changes in the subjects' behavior were consistent with theoretical predictions. Each of the three treatments have different mixed-strategy Nash equilibria. The first treatment gives high odds to the good card, the second treatment gives high odds to the bad card, and the third treatment gives a high payoff for player 1 to fold, and a high payoff for player 2 to call. The subsequent paragraphs guide you through their respective mixed-strategy Nash equilibria.

4.4 Treatment 1

The first treatment as mentioned above is designed to entice player 1 to bluff almost ninety-five percent of the time. To make this happen without giving player 2 a dominated strategy, and thus the game a pure-strategy Nash Equilibrium, we must make the probability of getting the card A no more than seventy-four per cent. A probability of seventy-five percent would in principle, make player 2 always fold, thereby giving the game a pure-strategy Nash Equilibrium where player 1 always bluffs and player 2 always folds.

The only difference between the base treatment and treatment 1 are the probabilities. That said, below are the two equations making both players indifferent between their strategies:

(2.1) $2.04\times p+1.78\times(1-p)=2\times p+2.52\times(1-p)\Rightarrow p=37/39\approx 0.95$

(2.2) $3.96\times q+4\times(1-q)=4.22\times q+3.48\times(1-q)\Rightarrow q=2/3$

Hence the mixed-strategy Nash equilibrium is:

(1/3 **Bluff** + 2/3 **Truth-tell**, 2/3 **Call** + 1/3 **Fold**)

4.5 Treatment 2

Treatment 2 is designed to entice player 1 to truth-tell more often. By giving a probability of seventy-five percent to the bad card, player 1 should in theory truth-tell exactly eight-ninths of the time and bluff one-ninth of the time. Like treatment 1, treatment 2 only differs from the base treatment in the probabilities. This leaves us with the following two equations making each player indifferent:

(3.1) $4\times p+3.25\times(1-p)=2\times p+3.5\times(1-p)\Rightarrow p=1/9$

(3.2) $2\times q+4\times(1-q)=2.75\times q+2.5\times(1-q)\Rightarrow q=2/3$

Hence the mixed-strategy Nash equilibrium is:

(1/9 **Bluff** + 8/9 **Truth-tell**, 2/3 **Call** + 1/3 **Fold**)

4.6 Treatment 3

The third and final treatment is designed to change the behavior of both players. The payoff structure is changed in two ways. First, the payoff given to player 1 from folding changes from 2 to 3.5. This can be interpreted as an increase in the opportunity cost of bluffing. Remember that player 1 bluffs with the hope that player 2 will fold, thus resulting in a payoff of 4 for player 1. Since the difference in payoffs between bluffing successfully and truth-telling is now smaller, it seems intuitive that player 2 should expect player 1 to bluff less or truth-tell more.   Player 2 should in principle respond to this by calling less. Second, if player 2 successfully calls a bluff, then he is given a payoff of 10 rather than 5. Similarly, player 1 should respond to this by bluffing less. This leaves us with the following two equations making each player indifferent:

$$5.5\times p + 2.5\times(1-p) = 2\times p + 3\times(1-p) \Rightarrow p = 1/8 \qquad (4.1)$$

$$3\times q + 4\times(1-q) = 4.25\times q + 3.75\times(1-q) \Rightarrow q = 1/6 \qquad (4.2)$$

Hence the mixed-strategy Nash equilibrium is:

(1/8 **Bluff** + 7/8 **Truth-tell**, 1/6 **Call** + 5/6 **Fold**)

6. Results

Base Treatment

Thirty-eight subjects (nineteen pairs) played ten rounds of the adapted version of Stripped-down poker. Players were randomly matched and their respective roles (i.e. player one or player two) were assigned randomly. The same role and pair assignments remained consistent throughout all ten rounds. Table 3 below shows the results aggregated over the ten rounds played.

Table 3. Aggregate experiment results from the base treatment.

	Player 1 given…					
	Card A 85/190 (45%)			**Card C 105/190 (55%)**		
Player 1	**Bet 84/85 (99%)**		**Fold 1/85 (1%)**	**Bet 67/105 (64%)**		**Fold 38/105 (36%)**
Player 2	**Call 65/84 (77%)**	**Fold 19/84 (23%)**		**Call 55/67 (82%)**	**Fold 12/67 (18%)**	
Payoffs	**5, 1**	**4, 2**	**2, 4**	**1, 5**	**4, 2**	**2, 4**

Before the game was started, the subjects were informed that the probability of getting card A and card C was exactly 50 per cent. As you can see from the table, the actual proportions were close, but not perfect. Card A was dealt 85 times out of a possible 190 (45 per cent) and card C was dealt 105 times (55 per cent).

By looking at the right side of the table where player one received card A, we see that player one folded once. In principle, one should never fold with card A. Because this action was chosen in the fifth round of a total of ten rounds, the most reasonable explanation is that player 1 either inadvertently chose to fold, or that they incorrectly thought they were dealt card C. In a real poker match, it might seem reasonable to fold when you have a good hand in order to confuse other players, but it is doubtful that this accurately describes this player's strategy. Now take look at the right side of the table where card C was dealt. Player 1 bet an overwhelming 64 per cent of the time while folding 36 per cent of the time. What does this suggest about the aggregate player's strategy over time?

Looking at player 1's decisions when they received card A reveals no information regarding their strategy unless they folded. This is true because both mixed-strategies involve betting with card A. However, let's assume the one fold with card A was inadvertent. That being said, we can estimate player 1's probability of mixing his/or her strategy by looking at their decisions when given card C.[5]

As mentioned above, player 1 bet 64 per cent of the time given card C. More intuitively, player 1 played the bluffing strategy 64 per cent of the time and played the truth-telling strategy 36 per cent of the time, on average.

Let's look at player 2's strategy when under both states of the world. Because player 2 doesn't know which card player one has been given, in terms of strategy, nothing is gained by looking at player 2's decision individually in each state. Therefore we must add up the number of times player 2 folded (120) regardless of the state, and divide that number by the total amount of times player 1 bet (151). On average, player two called 79.5 per cent and folded 20.5 per cent. Table 4 below compares the mixed strategy Nash equilibrium with the laboratory data for the base treatment.

Table 4. Theory vs. Results for the base treatment

	Player 1	Player 2
Theory	Bluff 33%, Truth-tell 67%	Call 67%, Fold 33%
Results	Bluff 64%, Truth-tell 36%	Call 79.5%, Fold 20.5%

Player 1 bluffed 31 per cent more than the mixed strategy Nash equilibrium and player 2 called 12.5 per cent more than the mixed strategy Nash equilibrium. Moreover, neither player was unilaterally playing their best response given what the other player was doing. Figure 4 below shows both players' best response functions:

5 The probability estimates follow a binomial distribution. Further statistical analyses, including confidence intervals and hypothesis testing can be found in the appendix.

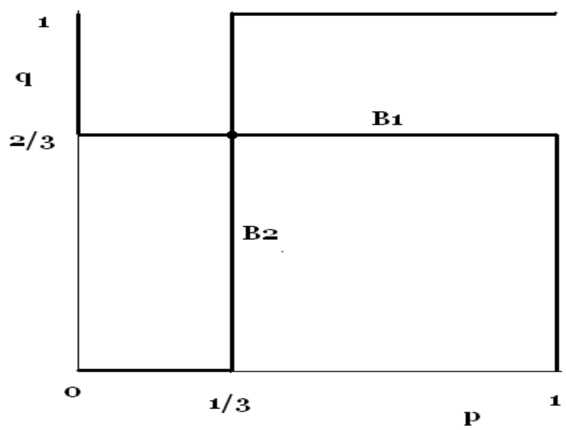

Figure 4. Best response functions for the base treatment.

If q (the probability that player 2 plays call) > 2/3, then player one's expected payoff to truth-tell exceeds his expected payoff to bluffing, and hence exceeds also his expected payoff to every mixed strategy that assigns a positive probability to bluffing. Similarly, if q < 2/3, then player 1's expected payoff to bluffing exceeds his expected payoff to truth-tell, and hence exceeds also his expected payoff to every mixed strategy that assigns a positive probability to truth-tell. If q = 2/3, then both bluffing and truth-tell, and hence all his mixed strategies, yield the same expected payoff.

Likewise, if p (the probability that player 1 plays bluff) > 1/3, then player two's expected payoff to call exceeds his expected payoff to fold, and hence exceeds also his expected payoff to every mixed strategy that assigns a positive probability to fold. If p < 1/3, then player two's expected payoff to fold exceeds his expected payoff to call, and hence exceeds also his expected payoff to every mixed strategy that assigns a positive probability to call. If p = 1/3, then both call and fold, and hence all his mixed strategies, yield the same expected payoff.

Given that player two called 79.5 per cent of the time, player one's best response would have been to always truth-tell. Likewise, given that player one bluffed 64 per cent of the time, player two's best response would have been to always call.

Treatment 1

Table 5. Aggregate experiment results from treatment 1.

	Player 1 given…					
	Card A 135/187 (72%)			**Card C 52/187 (28%)**		
Player 1	**Bet 134/135 (99%)**		**Fold 1/135 (1%)**	**Bet 35/52 (67%)**		**Fold 17/52 (33%)**
Player 2	**Call 90/134 (67%)**	**Fold 44/134 (33%)**		**Call 19/35 (54%)**	**Fold 16/35 (46%)**	
Payoffs	**5, 1**	**4, 2**	**2, 4**	**1, 5**	**4, 2**	**2, 4**

Once again, before the game started, it was explained to the subjects that the probability of getting card A was now 74 per cent, and the probability of getting card C was 26 per cent[6]. Player 1 received card A 72 per cent of the time and received card C 28 per cent of the time. Again, a single player folded once when given card A, but all others bet 134 times. When given card C, player one bet 67 per cent of the time and folded 33 per cent of the time. Player two called 64 per cent of the time and folded 36 per cent of the time. Table 7 compares the mixed-strategy Nash equilibrium to the results from treatment 1.

[6] It should be noted that after making this announcement, the player 2s were very unhappy. Given that player 1 earned 15 more points than player 2 on average over the ten rounds, they had every right to be upset!

Table 6. Theory vs. results for treatment 1.

	Player 1	Player 2
Theory	Bluff 95%, Truth-tell 5%	Call 67%, Fold 33%
Results	Bluff 67%, Truth-tell 33%	Call 64%, Fold 36%

Player 1 bluffed 28 per cent less than the mixed-strategy Nash equilibrium and player 2 called just 3 per cent less than the mixed-strategy Nash equilibrium, though not the best response to player one's actual strategy. Given that player 1 bluffed less than the optimal amount, in this case 95 per cent, player 2's best response was to always fold. Similarly, given that player 2 called less than the optimal amount, player 1's best response was to always bluff. This can be observed by looking at figure 5 below:

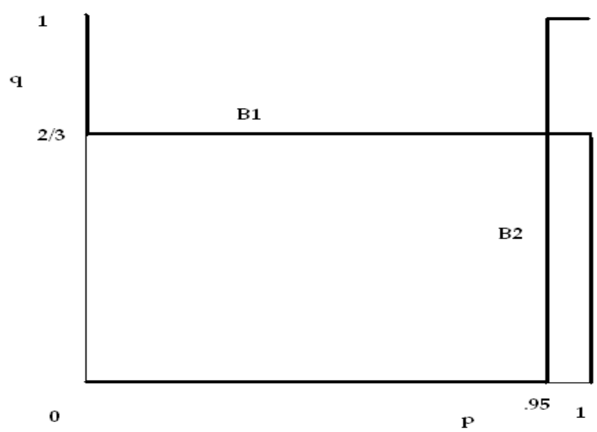

Figure 5. Best response functions for treatment 1.

Treatment 2

Player 1 bluffed only 25 per cent of the time. This is what was intended. In theory, they should have bluffed less, but there responsiveness was significant nonetheless. Player 2 called 93 per cent of the time, markedly above the mixed-strategy Nash equilibrium. Table 7 shows the decisions made in treatment 2 and Table 8 compares theory with the data from treatment 2:

Table 7. Aggregate experiment results from treatment 2.

	Player 1 given…					
	Card A 24/95 (25%)			**Card C 71/95 (75%)**		
Player 1	**Bet 24/24 (100%)**		**Fold 0/24 (0%)**	**Bet 18/71 (25%)**		**Fold 53/71 (75%)**
Player 2	**Call 22/24 (92%)**	**Fold 2/24 (8%)**		**Call 17/18 (94%)**	**Fold 1/18 (6%)**	
Payoffs	**5, 1**	**4, 2**	**2, 4**	**1, 5**	**4, 2**	**2, 4**

Table 8. Theory vs. results for treatment 2.

	Player 1	Player 2
Theory	Bluff 11%, Truth-tell 89%	Call 67%, Fold 33%
Results	Bluff 25%, Truth-tell 75%	Call 93%, Fold 7%

Player one bluffed 14 per cent more than the mixed-strategy Nash equilibrium and player two called 26 per cent more than the mixed-strategy Nash equilibrium.

Given that player 1 was playing his strategy nearly 14 per cent over the mixed-strategy Nash equilibrium, player 2's best response was to call 100 per cent of the time. Remarkably, player two called 93 per cent of the time. Figure 6 below shows each player's best response functions for treatment 2:

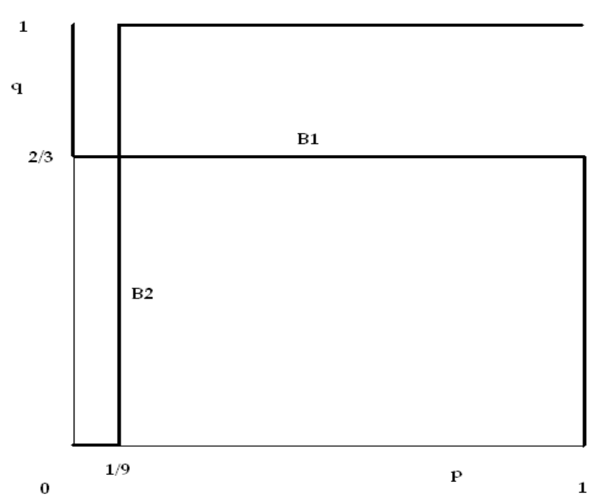

Figure 6: Best response functions for treatment 2.

Treatment 3

Table 9. Aggregate experiment results from treatment 3.

	Player 1 given…					
	Card A 34/86 (40%)			**Card C 52/86 (60%)**		
Player 1	**Bet 31/34 (91%)**		**Fold 3/34 (9%)**	**Bet 9/52 (17%)**		**Fold 43/52 (83%)**
Player 2	**Call 25/31 (81%)**	**Fold 6/31 (19%)**		**Call 6/9 (67%)**	**Fold 3/9 (33%)**	
Payoffs	**5, 1**	**4, 3.5**	**2, 4**	**1, 10**	**4, 2**	**2, 4**

Table 10. Theory vs. results for treatment 3.

	Player 1	Player 2
Theory	Bluff 12.5%, Truth-tell 87.5%	Call 17%, Fold 83%
Results	Bluff 17%, Truth-tell 83%	Call 77%, Fold 23%

Player 1 told the truth 83 per cent of the time. This is only 4.5 per cent less than the Nash prediction. Player 2 called 77 per cent of the time, 60 per cent more than the Nash prediction.

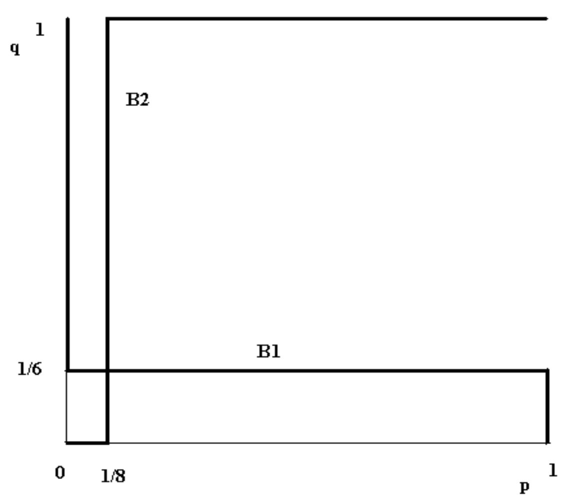

Table 11. Best response functions for treatment 3.

Given that player 1 was bluffing more than one-eighth of the time, player 2's best response was to call 100 per cent of the time. Likewise, given that player 2 was calling more than one-sixth of the time, player 1's best response was to truth-tell 100 per cent of the time.

8. Discussion

The results from the base treatment show a major disconnect between actual behavior and Nash predictions. Player 1 was not truth-telling enough and player 2 was calling too much.

After announcing that card A would be given 74 per cent of the time for treatment 1, the player 1's failed to respond by bluffing more often but the player 2's responded correctly by calling less than they were in the base treatment. Furthermore, player 2's actions followed closely to the Nash predictions, though they were not the best response to player 1's actual strategy.

Both truth-telling and calling increased significantly during treatment 2. The increase in truth-telling was predicted, however the increase in calling was unexpected.

Truth-telling increased significantly when the costs of bluffing were increased. Furthermore, the rate of truth-telling was extremely close to the Nash predictions. However, the rate of calling was much higher than the Nash predictions. These findings suggest that a strong relationship exists between truth-telling and the costs of bluffing.

9. Conclusion

In this paper I have shown that college age students behave strategically when confronted with decisions in which they could rationally deceive one another. The laboratory results suggest that there are indeed robust inconsistencies between theoretical predictions and revealed human behavior. These conclusions support prior reports given by others within the field of economics. Goeree and Holt (2001) suggest that observed behavior from laboratory experiments in which conform to the Nash mixed-strategy prediction seem "to only work by coincidence, when the payoffs are symmetric."

On the other hand, there were cases where changes in the treatment produced behavior that conformed rather nicely to theoretical predictions. When the opportunity cost of bluffing increased, the students' propensity to tell the truth markedly increased.

To the extent that children behave in accordance with adults, these findings are non-trivial. Unlike previous results from psychology studies, these results come from an experiment with both real payoffs and real consequences. The results imply that the propensity to tell the truth is a function of the costs and rewards of doing so and in particular of the costs of being detected. That being said, possible methods designed to help facilitate lowering these costs or sufficiently increasing the rewards could bring the truth out of children and increase the legitimacy of adolescent testimony. These findings ought to be very useful to justice officials, prosecutors, and researchers.

In the future I plan on extending this experiment to kids old enough to understand the game. My current findings coupled with possible data on kids helps strengthen the pre-existing literature on lie- and truth-telling and adds compelling evidence that helps explain some of the factors that change the propensity of truth-telling.

There are of course shortcomings of this research. One could argue that a poker match isn't the most appropriate way of modeling the elements of truth-telling and to an extent those critics are correct. Maybe there are better ways to model and observe this kind of behavior such as developing some sort of mock court where the rate of detection and penalties vary. However, this proposal seems a bit too contrived to be considered as a viable method of studying deception.

The games played were just that, games. Though there are truths in such arguments, the results reported in this paper do not lie. If you find yourself skeptical of these results, go ahead and *call* me.

References:

Bala, N., Lee, K., Lindsay, R. C. L., & Talwar, V. (2000). A legal and psychological critique of the present approach to the assessment of the competence of child witnesses. *Osgoode Hall Law Journal*, *38*(3), 409–452.

Bala, N., Lee, K., Lindsay, R. C. L., & Talwar, V. (2004). Children's Lie-Telling to Conceal a Parent's Transgression: Legal Implications. *Law and Human Behavior, 28(4), 411-435.*

Borel, Émile. *Traité de Calcul des Probabilités et ses Applications*, vol. 4, Gauthier-Villars, 1938.

Bruck, M., Ceci, S. J., & Hembrooke, H. (1998). Reliability and credibility of young children's reports. *American Psychologist*, *53,* 136–151.

Bussey, K. (1992). Lying and truthfulness: Children's definitions, standards and evaluative reactions. *Child Development*, *63,* 129–137.

Bussey, K. (1999). Children's categorization and evaluation of different types of lies and truths. *Child Development*, *70,* 1338–1347.

Davies, G., Lloyd-Bostock, S., McMurran, M., & Wilson, C. (Eds.). (1995). Psychology, law, and criminal justice: International developments in research and practice. Berlin, Germany: Walter de Gruyter.

Goeree, J. K., Holt, C. A., Ten little treasures of game theory and ten intuitive contradtictions. *The American Economic Review*, Vol. 91, No. 5. (Dec., 2001), pp. 1402-1422.

Lee, K., Cameron, C. A., Xu, F., Fu, G., & Board, J. (1997). Chinese and Canadian children's evaluations of lying and truth-telling: Similarities and differences in the context of pro- and anti-social behaviors. *Child Development*, *68,* 921–931.

Lewis, M., Stranger, C., & Sullivan, M. W. (1989). Deception in 3 year olds. *Developmental Psychology*, *25,* 439–443.

Peterson, C. C. (1995). The role of perceived intention to deceive in children's and adults' concepts of lying. *British Journal of Developmental Psychology*, *13,* 237–260.

Peterson, C. C., Peterson, J. L., & Seeto, D. (1983). Developmental changes in ideas about lying. *Child Development*, *54,* 1529–1535.

Polak, A., & Harris, P. L. (1999). Deception by young children following noncompliance. *Developmental Psychology*, *35,* 561–568.

Reiley, D. H., Urbancic, M. B., Walker, M. (2005). Stripped-down poker: A classroom game with signaling and bluffing.

Siegal, M., & Peterson, C. C. (1998). Pre-schoolers' understanding of lies and innocent and negligent mistakes. *Developmental Psychology*, *34,* 332–341.

von Neumann, John and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton, New Jersey: Princeton University Press, 1944.

Wimmer, H., Gruber, S.,&Perner,H. (1984). Young children's conception of lying: Lexical realism-moral subjectivism. *Journal of Experimental Child Psychology*, *37,* 1–30.

Appendix:

Table A.1. Binomial Distribution Properties

$$n \equiv \#\ of\ \ times\ given\ card\ C$$

$$x \equiv \#\ of\ \ times\ bet\ with\ card\ C$$

$$x \sim B(n, p)$$

$$\hat{p} = \frac{x}{n}$$

$$\sigma^2(\hat{p}) = \frac{n \cdot \hat{p} \cdot (1 - \hat{p})}{n^2} = \frac{\hat{p} \cdot (1 - \hat{p})}{n}$$

$$\sigma(\hat{p}) = \sqrt{\frac{\hat{p} \cdot (1 - \hat{p})}{n}}$$

Table A2. Confidence Intervals for estimated parameters

Base Treatment	Treatment 1
$n = 105$ $x = 67$ $\hat{p} = 0.638$ $\sigma^2 = .002$ $\sigma = 0.047$ $95\%\ CI\ for\ \hat{p}\ is...$ $\hat{p} \pm 2 \cdot \sigma$ $0.55 \leq p \leq 0.73$	$n = 52$ $x = 35$ $\hat{p} = 0.673$ $\sigma^2 = .004$ $\sigma = 0.065$ $95\%\ CI\ for\ \hat{p}\ is...$ $\hat{p} \pm 2 \cdot \sigma$ $0.54 \leq p \leq 0.80$
Treatment 2	**Treatment 3**
$n = 71$ $x = 18$ $\hat{p} = 0.254$ $\sigma^2 = .003$ $\sigma = 0.05$ $95\%\ CI\ for\ \hat{p}\ is...$ $\hat{p} \pm 2 \cdot \sigma$ $0.15 \leq p \leq 0.36$	$n = 52$ $x = 9$ $\hat{p} = 0.173$ $\sigma^2 = .0028$ $\sigma = 0.052$ $95\%\ CI\ for\ \hat{p}\ is...$ $\hat{p} \pm 2 \cdot \sigma$ $0.07 \leq p \leq 0.29$

Figure A1. A screenshot of the experiment taken from player 1's perspective.

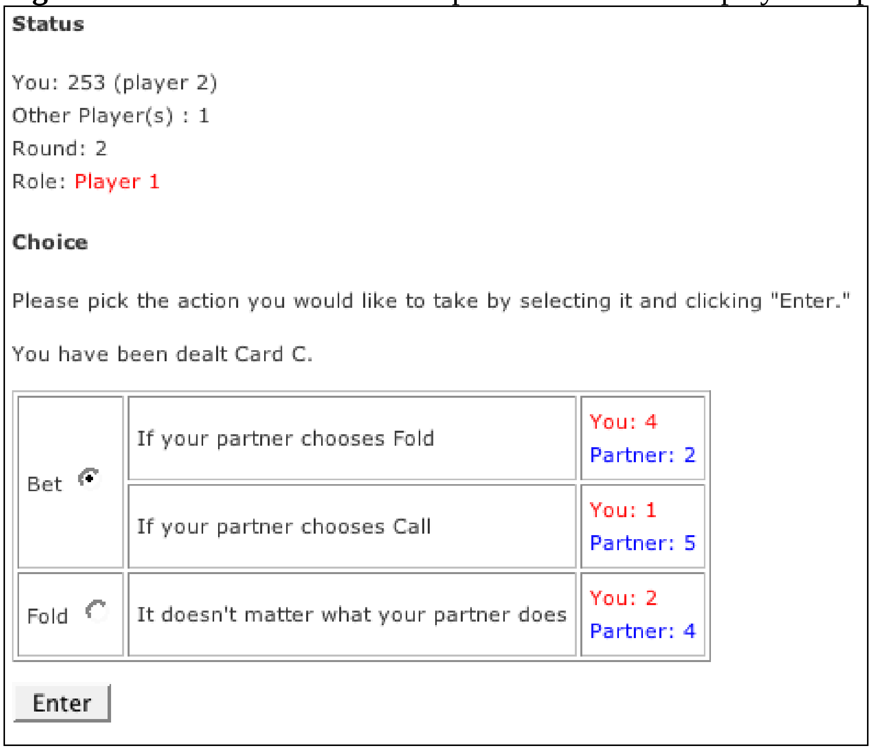
Status

You: 253 (player 2)
Other Player(s) : 1
Round: 2
Role: Player 1

Choice

Please pick the action you would like to take by selecting it and clicking "Enter."

You have been dealt Card C.

Bet	If your partner chooses Fold	You: 4 Partner: 2
	If your partner chooses Call	You: 1 Partner: 5
Fold	It doesn't matter what your partner does	You: 2 Partner: 4

Enter

Notice the bluff :)

Figure A2. A screenshot of the experiment taken from player 2's perspective.

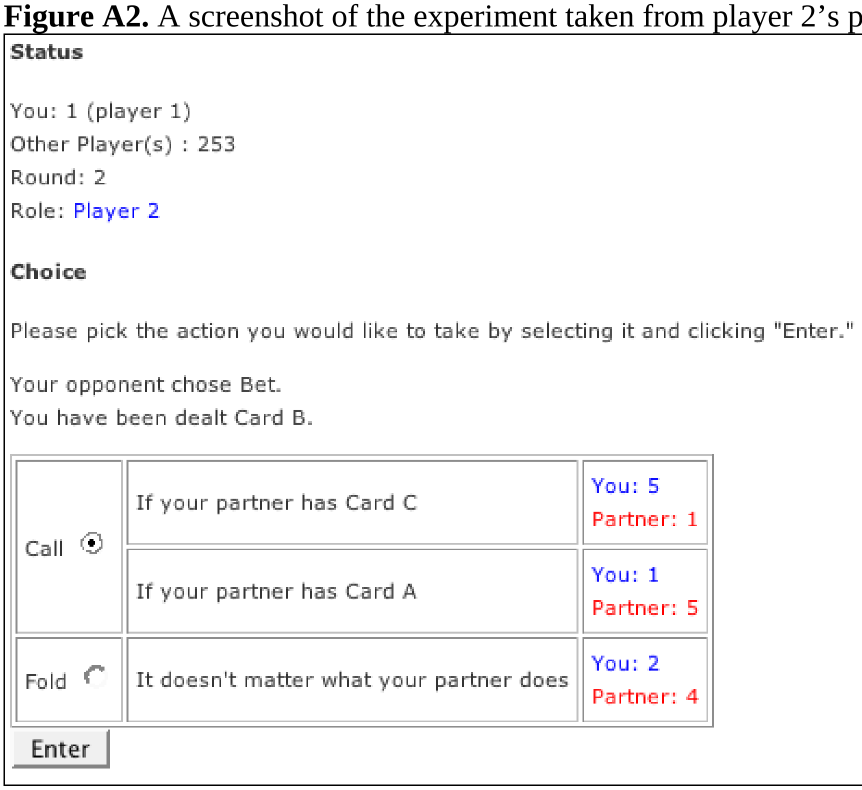
Status

You: 1 (player 1)
Other Player(s) : 253
Round: 2
Role: Player 2

Choice

Please pick the action you would like to take by selecting it and clicking "Enter."

Your opponent chose Bet.
You have been dealt Card B.

Call	If your partner has Card C	You: 5 Partner: 1
	If your partner has Card A	You: 1 Partner: 5
Fold	It doesn't matter what your partner does	You: 2 Partner: 4

Enter

…and the call of the bluff

Figure A3. Payout history from player 2's perspective.

Payout History

Rnd	Other	Role	Card	You	Opponent
1	253	Player 2		N » 4	F » 2
2	253	Player 2	C	C » 5	B » 1
Totals				9	3

Status

You: 1 (player 1)
Other Player(s) : 253
Round: 3
Role: Player 2

Choice

This round, you are the second mover. You have been dealt Card B, but it is not your turn yet.

Please tap refresh until you have a chance to move.