

Biases in Bluffing - Theory and Experiments

By Håkan J. Holm*
Department of Economics
School of Economics and Management
Lund University, Sweden

This version: October 2004.

Abstract: In some contexts people may be better at recognizing lies than truths or better at recognizing truths than lies. The implications of such detection biases are analyzed in a pure conflict of interest game with private information and communication. If the bias is modeled as a perfect signal from Nature that is sent with a certain probability conditional on the truth of a message, the bias can be incorporated into a game theoretical equilibrium analysis. Such an analysis reveals that the detection bias makes the equilibrium set shrink to a unique non-pooling equilibrium. In this equilibrium, the better a player is to detect lies the more often will the opponent player lie. This counter-intuitive result could be used in hidden information problems.

In the bluffing game experiment, subjects were telling the truth too often according to standard game theoretical predictions, but in line with one of the detection bias hypotheses. However, not all experimental observations supported the detection bias hypothesis. Other findings were a significant positive correlation between self-rated bluffing ability and actual bluffing performance. Altruism was significantly negatively correlated to bluffing performance and a corresponding positive correlation was found for cooperativeness. Gender mattered in that subjects were significantly more prone to lie to a female opponent than to a male. Furthermore, even if there was no significant difference in bluffing performance between males and females, the former rated them selves as significantly better at bluffing than the latter.

Keywords: Bluffing, Game theory, Truth detection, Lie detection, Experiment.

* I am grateful to Ola Andersson and Erik Wengström for help when conducting the experiment. I have received valuable comments from Hans Carlsson, Toshiji Kawagoe and Tetsuo Yamamori. An earlier version of this paper was presented at a session at the Economic Science Association's World Meeting in Amsterdam, in June 2004 and at a seminar at the Department of Economics at Lund University in Sweden. I am grateful for comments from the participants at these meetings. Financial support from The Bank of Sweden Tercentenary Foundation and the Swedish Competition Authority is gratefully acknowledged.

1 Introduction

*There are a terrible lot of lies going around the world,
and the worst of it is half of them are true*
Sir Winston Churchill

Lying in strategic situations has received little attention in both the theoretical and experimental literature. The standard assumption is that if a player does not want it, communication does necessarily not disclose his type or his intentions. Thus, lying is possible and costless. Without this assumption the literature on asymmetric information had to be modified and justified on partly new grounds. Furthermore, certain institutional arrangements like risky communication between firms in illegal cartels could partly be explained by psychology if lying was difficult.¹

As recognized by Crawford (2003) and Hendricks and McAfee (2003) many strategic situations involve a stage where one party has the possibility to misrepresent information. One example is where to attack in war. Since disclosure of real intentions could be exploited by the counterpart traditional game theory (see Crawford and Sobel, 1982) predicts that zero-cost messages sent in such games are not informative (i.e., is cheap talk). However, Crawford (2003) argues that there are many examples in reality where such messages are actually sent and appears to have an effect. If players are heterogeneous with respect to their reasoning capacity Crawford (2003) has demonstrated that misrepresentation in such games may matter. However, there is also a more direct and psychological explanation. Players may have different skills in lying and they may have different skills in recognizing signals that are directly observed and related to the act of lying or truth-telling. This is a related but different explanation to why certain lies have an effect and others not. In addition to this, it is not uncommon that people claim that they can see through a lie and others claiming that they are quite bad at bluffing.² Needless to

¹ Such meetings could then be motivated if firms' true intentions were communicated even if they also increased the likelihood that the authorities would caught them.

² Beliefs about such abilities might have had some historical impact. One example is when Hitler in a meeting before WWII lied to the ex-British Prime Minister Chamberlain about his intentions to invade Czechoslovakia.

say, such abilities would be an obstacle for the standard cheap-talk outcomes and even misconceptions of such abilities or involuntary actions (like e.g., blushing when lying) might disturb.

An interesting observation in psychology is that peoples' skill of detecting truths might differ from their ability to recognize lies. We will call this asymmetry a detection bias. In a review based on some 40 studies Vrij (2000) found a 67 percent average accuracy rate for detecting truths. The corresponding accuracy rate for detecting lies was only 44 percent. This difference can partly be explained by a tendency for people to judge other's statements as truthful. However, this tendency (which itself is not entirely easy to understand from a game theoretical perspective) cannot account for the entire difference, which suggest that a general truth detection bias may exist.³ In this paper it is not important if people are better at recognizing truths than lies or the contrary, the important thing is that the recognizing capacity might differ and is conditional on the truth value of the message.⁴ Because, even if there would not be any general statistical discernable detection biases, there might still be detection biases that are relationship specific. For instance, a man may know that his wife sometimes with certainty can tell when he is lying and both are fully aware about this. The reason might simply be that the man is a hopelessly poor liar and the wife a good lie detector that has learnt some observable cues associated with her husband's lies. The man is not in full control of when he sends these cues. After a while both

Chamberlain obviously with a certain confidence in his abilities to identify liars wrote to his sister: "in spite of the hardness and ruthlessness I thought I saw in his (i.e., Hitler's) face, I got the impression that here was a man that could be relied upon when he had given his word" (Ekman, 1992, pp15-16, parenthesis is mine.) In the Parliament Chamberlain said that he was convinced that Hitler did not try to deceive him. Czechoslovakia was invaded by Germany a few weeks later.

³ One possibility is that there are evolutionary explanations to the detection bias. Frank (1987) analyses a model where there are honest and dishonest types and were the players receive an imperfect signal of the other player's type. According to that model certain traits that signal honest behavior may have a survival value and it is possible that the same reasoning applies to truth detection; during the evolution it has been more important to recognize truths than lies. In any case, it is not obvious why detection abilities of truths and lies should be exactly the same.

⁴ There is some recent evidence that lying and truth-telling generates physiological different effects that are observable with modern brain-scanning techniques (see *The Economist*, 2004, July 10th, "Lie Detection, Making Windows in men's souls, p. 71-2). If lying is partially physiologically different from truth-telling it is not impossible that some of these physiological differences also generate differences in behavior that are i) observable ii) more easily recognizable for truths (or lies) than for lies (truths). Furthermore, these physiological differences also stimulates ones imagination of what can happen in the future if these techniques are made available outside research centers and court rooms.

realize this. From an economic theoretic perspective both the conceivable general tendency of detection biases and more relationship specific detection biases motivate a game theoretical analysis of the issue. Although, the issue is recognized in psychology it is (to my knowledge) new in economics and game theory.

In this paper the aim is to study the implication of a detection bias in a pure bluffing context without any disturbing mechanisms. To do this, a simple two-person bluffing game is developed. In the game an S-player receives a perfect signal of a state variable from Nature. The S-player then makes a statement of the state variable and an R-player is to make a guess if the statement is true or false. R wins a sum of money if her guess about the truth value of S's statement is right. S wins otherwise. First, it is shown that a straightforward interpretation of the detection bias is inconsistent with Nash equilibrium in this game. However, the game is possible to analyze with standard apparatus as a game with either a truth detection bias or a lie detection bias (in parentheses below) if it is assumed that the R-players' ability to detect truths (lies) better than lies (truths) manifests itself as a positive probability of her receiving a perfect signal about the statement's truth value. This signal can only be received if S is telling the truth (a lie). One somewhat unintuitive result of the equilibrium analysis is that the better the R-player is at detecting truths (lies), the higher is the probability that the S-player tells the truth (a lie).

We will also study the bluffing game experimentally to test the hypotheses of truth and lie detection biases. A number of experiments have been conducted concerning the impact of cheap talk in coordination games (see Crawford, 1998 and Camerer, 2003 for reviews on this literature). In such experiments players typically communicate by anonymous written messages (e.g., e-mails). The experiment in this paper differs in that it studies a game of pure conflict of interest and applies a face-to-face design. The reason for using a game of pure conflict of interest where the gain of one player corresponds to equally large loss to the other player is to get the players' incentives focused on bluffing and truth detection. In contrast to coordination games behavioral regularities relating to focal points and risk preferences are absent in the game studied

in this paper. The reason for using face-to-face design is to give room for truth detection cues and also and for signals aimed at feinting the counterpart.

More closely related to the experiment conducted in this paper there are some studies of bargaining games, where players are given the opportunity to communicate face-to-face and possibly misrepresent their information or intentions. In Frank et al. (1993) it was shown that players that met face-to-face before they played a prisoners' dilemma were quite good at predicting cooperative behavior of their opponent. This result may be interpreted to suggest that players involuntarily send truth signals. Furthermore, Brosig (2002) confirmed this result and got some support for the hypothesis that subjects classified as cooperative in a psychological test were better than individualistic subjects at predicting cooperative behavior. These kinds of results may have profound consequences for the understanding of the institutions surrounding bargaining if it could be substantiated with further evidence.⁵ However, somewhat contrary to this result Ockenfels and Selten (2000) found that after controlling for some objective criteria, onlookers of a bargaining process were not better at guessing if a player had a certain form of private information than could be randomly expected. These somewhat conflicting results suggest that the question deserves more empirical research. One important aspect concerns the type of game studied. The two first experiments mentioned above focus on the subjects' ability to detect the opponent's type, that is, his utility function in terms of willingness to cooperate. They do not focus on abilities to detect hidden information in general. Consequently, in the games studied players may actually not have incentives to lie. For instance, in a prisoners' dilemma, subjects with strong inequality aversion (see e.g., Fehr and Schmidt, 1999 and Bolton and Ockenfels, 2000) might prefer the cooperative outcome to all other outcomes. These subjects might tell their true intentions and would not necessarily have any motive to misrepresent their intentions. It is not inconceivable that such preferences are systematically related to the ability to lie and to the

⁵ For instance, it would then be easier to understand why bargaining partners in spite of modern communication technology and high travelling costs actually think it is worthwhile to meet.

ability of telling if someone is lying.⁶ Thus, it is not entirely apparent how to separate these preferences from pure bluffing and lie detection abilities. Related issues might affect the game studied in Ockenfels and Selten (2000). In this game subjects were to divide a sum when they with a certain probability had been exposed to a cost, which was private information. Various ideas of fairness reasoning might affect how the subjects perceive the situation in this game and if they are motivated to lie.⁷ This may also affect the interpretation of results. Contrary to this, all subjects in the bluffing game in this paper have incentive to misrepresent the private information. The opponent also has a clear incentive to outguess the player with private information.

In the psychological literature there are numerous articles on lying and lie detection (see e.g., Vrij, 2000). Typically, such studies involve S-subjects (like S-players in the bluffing game above) that are instructed to tell a lie and/or a truth, and one set of R-subjects (like the R-players in the game) that are instructed to guess whether the S-subjects are telling the truth or not. The medium through which the lie or truth is communicated varies (e.g., videotape with or without sound and tape recorders). This research provides valuable insights into the psychology of lying and lie detection cues, but it ends up focusing on situations that are partly beside the ones involved in many economic or strategic situations. In strategic situations both parties often know that they are playing a certain game with certain payoffs. This is typically not the case in the psychological studies where the outcome/payoff for the S and R-subjects is usually not even contingent on their actions. Furthermore, in strategic situations lying or truth telling (by a person in the S-player situation) is typically endogenous (in the sense that it depends on the structure of the game) and the players know this. Such an endogeneity is hard to find in psychology. This study complements the psychological approach by a game theoretical analysis of a situation, where the payoff structure is transparent and the endogeneity is present.

⁶ In fact, one result in this paper is that there is a significant negative correlation between altruism and earnings in the bluffing game.

⁷ Indeed, in the data provided in Ockenfels and Selten (2000) one type of equal of divisions clearly outnumber other divisions that are conditional on the private information. In 75 percent of the bargaining outcomes players agreed on the fifty-fifty outcome. Furthermore, see e.g., Hennig-Schmidt (2002) for a study of bargaining games where subjects refer to different principles of fairness depending on the strategic situation.

2 Theory

This section will analyze the implications of a truth detection bias and a symmetrical lie detection bias. As noted before, it should be emphasized that such an assumption would affect the analysis in many strategic situations. For instance, the basic assumption in adverse selection and signaling models is that people can lie or tell a truth without revealing any information.

The game is now introduced more formally. In period 1 Nature selects the state variable $s \in \{R, B\}$ with probability $p_N = 1/2$. The S-player then observes a perfect signal of the state variable. In period 2 the S-player makes a statement $m \in \{R, B\}$ to the R-player about the state variable. Player R is then to make a guess, $g \in \{T, F\}$, whether or not the statement is true ($g = T$) or false ($g = F$). A statement is said to be true if $m = s$ and false otherwise. R wins 1, if her guess correct, that is if $g = T$ and $m = s$, or if $g = F$ and $m \neq s$. In that case S gets zero. If R's guess is incorrect S wins 1 and R gets zero.⁸

There is an infinite number of mixed (perfect Bayesian) equilibria in this game, where $p_A \in (0,1)$ and $p_B = 1/2$, and where p_A and p_B denote the unconditional probability player S tells the truth and the conditional probability player R guesses that the statement is true, respectively.⁹ Furthermore, in equilibrium according to R's belief the probability of being in each of the nodes in the non-singleton information sets is $1/2$. There are also pure pooling equilibria, where the S-player makes a statement (R or B) with probability one and player R chooses for each statement either T or F with probability one.

Suppose now that a detection bias is introduced. Let us first interpret this bias (as psychologists appear to do) as differences in observed accuracy rates between truth and lie detection. Translating this into game theoretical terminology would mean that the probability R is

⁸ This game can be interpreted as a signalling game (see e.g., Gibbons, 1992), where Nature draws S's (i.e., the sender's) type after which he sends a message to R (the receiver) who forms beliefs about S's type and chooses an action contingent on these beliefs. The game in its full extensive form is given in Figure A1 in the appendix.

⁹ R's probability is conditional on m .

correct *ex post* varies and is conditional on if S is lying or telling the truth. Such a conception of the detection bias will cause serious analytical problems in the sense that it will not be consistent with any equilibrium. Hence, we have the following observation:

Proposition 1: A conventionally interpreted detection bias is not consistent with the existence of Nash equilibrium.

Proof: Let t be the probability that R is correct if S tells the truth and let l be the probability that R is correct if S tells a lie. A conventionally interpreted detection bias means that $t \neq l$. The expected payoff to S of telling the truth and a lie would be $1-t$ and $1-f$, respectively. For this to be a mixed equilibrium in this game the expected payoffs of the strategies should be the same, but this can obviously not be the case, since $t \neq f \Leftrightarrow 1-t \neq 1-f$. Furthermore, any combination of pure pooling strategies is not consistent with a detection bias. To illustrate take the case where S chooses B with probability one. If R then has a pure strategy we would have that $t = l$, which contradicts the assumption of a detection bias. Q.E.D.

It is not difficult to understand the intuition for this result. The assumption of a detection bias imposes an unnatural restriction on the outcome, *ex post*, that is inconsistent with endogenous rational behavior.

The question is then if it is possible to model detection bias in a more constructive way. We will start with phrasing the detection bias as a truth detection bias and then show that lie detection biases can be modeled in a symmetrical way. The suggestion of incorporating a truth detection bias to be made in this paper involves adding a stage in the game where R with a certain probability $\pi \in (0,1)$ observes a perfect signal (Y) if the statement is true.¹⁰ This stage takes place after S has made his statement and only if S tells the truth (i.e., $m = s$). In the proof of Proposition 2 we show that as soon as a truth detection bias is introduced the pooling equilibria

vanish and the game can be reduced to the one in Figure 1. Since, the decision to lie or tell the truth is the crucial aspect for S in the remaining equilibrium the analysis can do without referring to the state variable in this symmetrical game. In the reduced form of the game we simply say that S chooses between telling the truth (T) or lying (L), where T means $m = s$ and F means and $m \neq s$.

Before the equilibrium is derived some variables will be introduced. Let a denote the probability S chooses T , and let μ denote R's beliefs that she is at node R_2 in Figure 1 when in this information set. Finally, let b denote the probability R chooses T in the information set that is not singleton (i.e., consisting of nodes R_2 and R_3) and let c be the corresponding probability for the singleton node (R_1).

Proposition 2: Introducing a truth bias as described above leads to unique perfect Bayesian equilibrium in the Bluffing game. The equilibrium is characterized by $a^ = 1/(2 - \pi)$, $b^* = (1 - \pi)/(2 - \pi)$, $c^* = 1$ and $\mu^* = 1/2$.*

Proof: Let us start by ruling out all pooling equilibria. Suppose, in the non-reduced game that e.g., B is played with a certain strictly positive probability that is state (or type) independent (i.e., pooling) then R's updated belief concerning the probability she is in either of the nodes in the non-singleton information set will differ in the case she did not receive the perfect signal Y . Since, R did not receive a perfect signal, it is more probable that S has lied about the state. This will make it optimal for R to choose F with probability 1. Clearly, this could not be consistent with equilibrium, since S would expect this and react to it.

Let us now derive the mixed equilibrium that is not pooling by analyzing R's decision in the last stage.¹¹ In this stage there are two different information sets corresponding to nodes R_1 and, R_2 and R_3 in Figure 1. At node R_1 , R has received a perfect signal, which means

¹⁰ This signal could be a twinkle, a tick, a blush, a particular brainwave recognizable by EEG or a specific signal that the R player has learnt to recognize is associated with truth-telling. Furthermore, S knows this.

¹¹ This equilibrium is weakly separating in the sense that the probability S sends a certain message (B or R) in the non-reduced game is type dependent.

that it is optimal to set $c^* = 1$. At nodes R_2 and R_3 R's conditional payoff from choosing T and F will be $a(1-\pi)/(1-a\pi)$ and $(1-a)/(1-a\pi)$, respectively. Since payoff from the pure strategies must be the same in a mixed equilibrium we can conclude that $a^* = 1/(2-\pi)$. Similarly, S's payoff from the pure strategies of T and L are $(1-\pi)(1-b)$ and b , respectively and must also balance. Hence, $b^* = (1-\pi)/(2-\pi)$. Furthermore, since μ denote R's beliefs that she is at node R_2 when in this information set, consistent beliefs then require that $\mu = (1-\pi)/(a^{*-1} - \pi) = 1/2$.

Uniqueness should be clear from the following observations. First, $\mu \neq 1/2$ cannot be consistent with equilibrium since then it would be optimal for R to select a pure strategy in this information set. A pure separating strategy can obviously not form an equilibrium in this matching pennies like game. Furthermore, if $\mu = 1/2$, then any other strategies than the ones chosen would not balance the pure strategies. Hence, together with the non-existence of any pooling equilibrium this establishes the uniqueness result. Q.E.D.

To see that the reasoning behind Proposition 2 also applies to a lie detection bias, assume instead that π is the probability R observes the perfect signal (Y) if the statement is false. This stage takes place after S has made his statement if and only if S lies. Furthermore, let then a denote the probability S chooses F , and let μ denote R's beliefs that she is at the node where S has lied but no perfect signal has been emitted from Nature and let b denote the probability R chooses F in the information set that is not singleton. Finally, let c be the corresponding probability for the singleton node.

Note independently if the bias is in terms of lie detection or truth detection the introduction of it has a theoretical quality in that the equilibrium set is refined. Furthermore, it is not crucial that there actually exist a true detection signals. The important thing is that the players believe in certain detectable truth or lie signals. This has important implications as we will explain below. Let us now use Proposition 2 to make some observations concerning the implied

equilibrium behavior when a detection bias is at hand. The observations will be expressed in terms of truth detection, but the interpretation in lie detection is given in parentheses.

Observation 1: If there is a truth (lie) detection bias then $a^ > \frac{1}{2} > b^*$, and the better R is at recognizing truths (lies) the more often will S tell the truth (lie). The better R is on recognizing truths (lies) the more seldom will R guess that S is telling the truth (lie) in the non-singleton information set.*

At first sight this observation may seem counter-intuitive since one might expect S to be more wary to tell the truth the better the opponent is on tracing it. The observation is perhaps even more counter-intuitive when expressed in terms of lie detection; the better R is on detecting lies the more often will S lie. However, if it is realized that the better R is to detect lies the more informative will it also be for her not to receive a signal. Hence, the higher lie detection bias the more likely it is that S is telling the truth if no lie signal is observed. R would exploit this by guessing more often that S is telling the truth in this situation. This will make the truth strategy less profitable to S. Thus, one might say that S is squeezed by either lying with the risk of being revealed by the perfect signal or being outguessed when telling the truth.

One interesting implication of this observation (and Proposition 2) is that if R can make S believe in the detection bias game, she would command a powerful tool to improve her information about the true state even if no perfect signals actually existed.¹² Suppose, before a large invasion by an S-army, a high-ranked officer in that army is captured by the R-army. Both armies know that S will soon attack, but only the S-officers know exactly where. There are only two feasible places to attack, Redding or Blackburn and the attack and defense must be done with whole forces. R wins if and only if its officers succeed to deploy its army in the city attacked and

¹² Some professionally used lie detection methods rely on deceiving subjects into believing in the lie detection device. After the deception, physiological reactions (blood pressure, palm sweating, respiration etc.) are

the S-army wins otherwise. Now, assume that the captured S-officer only cares about the future success of his army and is brought in to be questioned by the R-officers about the place S plans to attack. Thus, the strategic situation is similar to the one in the bluffing game. However, before he is confronted with the crucial question he has carefully and successfully deceived into believing in a lie detection device possessed by the R-army and, also that the R-officers actually believe in it, which they do not. He is informed that the probability that the lie detector emits a perfect signal (Y) if he tells a lie is very high. When he is confronted with the crucial question Proposition 2 predicts that the officer will lie with high probability. Hence, if the officer's answer to the crucial question is Redding, the R-officers would know that Blackburn is the probable target for the attack.¹³

Observation 2: The players' expected payoff in equilibrium will be $\frac{1-\pi}{2-\pi}$ and $\frac{1}{2-\pi}$ for S and R, respectively. Thus, S is expected to earn less than R. S's expected payoff is decreasing in R's ability to detect truths (lies). R's expected payoff is increasing in R's ability to detect truths (lies).

This observation is hardly surprising and follows directly from Proposition 2 if the equilibrium strategies are inserted in the players' payoff functions.

Observation 3: The overall frequency for R's truth (lie) guesses is $\frac{\pi^2 - 2\pi + 2}{(2-\pi)^2} > \frac{1}{2}$.

measured when certain questions are posed (see Vrij, 2000, p.179). What is argued here is that in certain strategic situations such measurements may be unnecessary!

¹³ Clearly, this meta-lie detection mechanism is sensitive to beliefs about the game and requires that the players are rational. However, it is important that the players are not too smart. For instance, if the captured officer realizes the deception and can feign that he believes in it, he can use a meta-feinting strategy by telling the truth since he knows it will be taken as a lie. Obviously, there is no end to such a meta reasoning.

In the case of a truth detection bias the R-player chooses T when she has received a signal and is certain and also when she is uncertain with probability b . Whereas, the former situation will occur more frequently as π increases the latter will be less frequent.

3. Hypotheses to test empirically

The first natural hypothesis to test is the game theoretical prediction that there is no detection bias and that play is consistent with mixed strategy equilibrium (MSE). Testing this hypothesis is not uninteresting since previous empirical research has mainly used “small” channel communications like, e.g., e-mail messages. If subjects play consistent with MSE then the probability of sending a true or false message should be equal. This leads to our first hypothesis:

H1: Behavior is characterized by MSE : $\Pr(m = s) = \Pr(m \neq s)$.

If H1 is rejected, one reason might be the prevalence of a detection bias as described in section 2. However, as shown by Observations 1-3 a detection bias has a number of implications, which partially depends on if it is a lie detection bias or a truth detection bias. The hypotheses can be summarized into two composite hypotheses. Below H2T denotes the hypothesis in the case of truth detection and H2L denotes the lie detection hypothesis:

H2T: Behavior is characterized by truth detection bias, which means the following:

a) S-players more often tell truths than lies: $\Pr(m = s) > \Pr(m \neq s)$ (from Observation 1).

b) R-players earn more than S-players (from Observation 2).

c) For R-players’ “truth-guesses” are more frequent than “lie-guesses”: $\Pr(G = T) > \Pr(G = F)$

(from Observation 3).

H2L: Behavior is characterized by lie detection bias, which means the following:

a) S-players more often tell lies than truths: $\Pr(m = s) < \Pr(m \neq s)$ (from Observation 1).

b) R-players earn more than S-players (from Observation 2).

c) For R-players' "lie-guesses" are more frequent than "truth-guesses": $\Pr(G = T) < \Pr(G = F)$ (from Observation 3).

Our next hypotheses are motivated by aspects that might affect bounded rational subjects' bluffing behavior. This is contrasted with the more neutral assumptions of traditional game theory, which will serve as null hypotheses. Both the theoretical literature and experiments suggest that bounded rationality account for part of the deviations from the MSE. One hypothesis, put forth by Crawford (2003) is that players differ with respect to their degree of rationality. If it is assumed that there is heterogeneity with respect to rationality types among the subject pool the distribution of earnings would not necessarily be normally distributed if the game is repeated.¹⁴ Another and more psychological explanation for non-normal earnings distributions would simply be that some S-players are better than others at deception or that some R-players are better than others at lie detection.¹⁵ On the other hand if behavior adheres to MSE where types typically are homogeneous and fully rational, the distribution of earnings over individuals in the population would be the sum of i.i.d. random numbers, which would approximate a normal distribution. Taking the MSE as our null hypothesis we get the following hypothesis:

H3: The distribution of earnings is normally distributed.

There is also a question about what beliefs subjects have about their ability to bluff and to reveal bluffs. These beliefs may decide what types of strategic situations individuals prefer.

¹⁴ In a certain type of equilibrium in that model (, but not in all) rational players (so called sophisticated types) exploit bounded rational players (so called mortals) and have strictly higher expected payoffs.

¹⁵ For instance, Ekman and O'Sullivan (1991) and Ekman et al. (1999) found that some professional groups (such as members of the Secrete Service) appear to be more skilled in truth/lie detection than other.

Psychological research has only made moderate progress in establishing general methods of detecting lies. Some cues like, for instance, higher-pitched voice and decreased movement in arms and legs appear to be statistically linked to lying (see Vrij, 2000). Hence, it is conceivable that some (boundedly rational) people are more prone to send away these cues than others and are therefore less able to bluff. On the other hand, people may differ in their perceptiveness of such signals and may thus differ in their ability to reveal bluffs. If this is the case it is quite possible that subjects have varying beliefs of their ability to bluff and/or to reveal bluffs. Furthermore, if that is the case and they are correct in their evaluations of themselves one should observe a relationship between such skills and bluffing game earnings. One additional consequence is that if R-players are convinced that they are good at detecting lies and they can choose between one situation where the S-player is uninformed about m and one where he is informed they would prefer the latter situation. The reason is that in the first situation the S-player would not know if he is telling a truth or a lie. If B is sufficiently convinced in her ability she would be prepared to forego money to change from the first to the second situation. In principle, it is also possible that the S-player believes that he is good at feinting. If this is the case it is conceivable that the S-player is willing to forego money in order to get information about m . However, from a traditional game theoretical perspective this reasoning is somewhat foreign. Thus, the composite hypothesis about these beliefs is the following:

H4. Hypothesis about beliefs about bluffing and detection skills.

- a) There is no relationship between beliefs about capacities to bluff or to detection detect truth/lies and earnings in the bluffing game.*
- b) If S-players can avoid the signal (i.e., m). R-players would not be prepared to pay for forcing S to receive the signal.*
- c) If S-players do not receive the signal (i.e., m) they will not be prepared to pay for getting information about the signal.*

4. Design of the experiment

The experiment consisted of three similar sessions with 14 subjects in each during the Spring of 2003. Subjects were recruited from the introductory course in Economics at the School of Economics and Management at Lund University. As the subjects arrived they received general information about the experiment. They were informed about practical aspects (such as the expected duration) and the general purpose of the study (to study decision making).

The subjects received one ID-note and two forms containing questions. ID-notes were used to increase the subjects' anonymity and to serve as a receipt when collecting their earnings. The first form contained questions about personal characteristics (e.g., age and gender), self-evaluation about their capacity to bluff, reveal lies, trust, trustworthiness and honesty.¹⁶ The second form contained questions aimed at measuring personality traits. The reason for including these questions was to study if bluffing and lie detection was associated with certain personality traits. A number of questions from the International Personality Item Pool were used. These questions are designed to measure certain psychological personality traits.¹⁷ We measured six personality traits, the degree of impulse control, altruism, cooperativeness, risk-taking, extraversion and excitement-seeking.¹⁸

All tasks before the Bluffing game subjects were conducted individually, where talk to anybody else than the experimenters was explicitly forbidden. The Bluffing game took place in a room where each S-player was seated according to his/her ID-note number at one side of a table. Each R-player was then seated in front of one S-player according to a pre-specified order (also determined by their ID-note numbers). Players kept their roles throughout all repetitions of the game. Subjects were then given both verbal and written instructions to the game.¹⁹ They were also given the opportunity to ask questions before the game started.

¹⁶ These questions are available upon request by the author.

¹⁷ See <http://ipip.ori.org/ipip/>.

¹⁸ The degree to which a subject is considered to have a certain personality trait is measured by the subject's answer to ten different yes or no questions. We mixed questions for different traits to make it less obvious to the subjects, which trait the questions concerned. The questionnaire used is available upon request by the author.

¹⁹ See Appendix for instructions.

Each round started by the experimenter distributing cards to S-players and asking them to look at their cards. The S-player was then asked to state the color of his card and R to state whether S's statement was true or false. After that, S showed the true color of the card to R and filled in a paper with a scoring table. If R's guess about S's statement was correct R earned SEK15, otherwise S earned the same amount. R then checked that S had filled in the table correctly and signed with his ID-note number.²⁰ End of round. Before the next round R-players rotated to a new table and a new S-player. This was repeated in ten rounds of which the first seven were with a new S-player. After the tenth round was finished the papers with scoring tables were collected.

To measure the value R-players put on their ability to reveal lies (see hypothesis H4b) when the S-players made statements a new treatment was added in five additional rounds in the two first sessions. In this treatment S was not to look at the card before he made the statements, which meant that S did not have more information about the color of the card than R. However, R could bid on the right of having S to look at the card before the statement was made. In this way R revealed her valuation of her ability to reveal lies/truths in the statements made by S. The eliciting method was incentive compatible and the maximum amount R was allowed to bid in each round was SEK 10.²¹ In the last session S-players (and not R-players) were to bid on the right to look at their cards before making the statement. In this way S's valuation of his ability to bluff or feint R was obtained (see H4c).

5. Results

In all we observed the behavior of 42 subjects. The average age of the subjects was 22 years and 55 percent were males. All parts of the experiment took somewhat less than one hour and the

²⁰ SEK15 corresponded roughly to USD2 when the experiment was conducted.

²¹ The method used was a variant of the mechanism used in Becker, DeGroot and Marschack (1964). For details about instructions see the instructions in Appendix. An exercise round was conducted to get the subjects familiar with the method.

average earning was SEK110, which is more than the hourly wage for this group. Below the results relating to the hypotheses in section 3 are presented.

5.1 The hypotheses of MSE and Truth Detection Bias

Our first hypothesis concerns whether behavior is characterized by MSE. Figure 2 displays the average proportion of S-players that told the truth in each round. According to the MSE hypothesis this figure should be around 0.5, since the probability of telling the truth ought to be the same as the probability of telling a lie. It should be clear from the figure that on average S-players were in each one of all ten periods more inclined to make a true statement than to make a false. Out of 210 statements made by the S-players 137 was truthful. Consequently, a statistical test strongly rejects hypothesis H1, namely that the probability of a truthful statement is the same as the probability of a false ($p < 0.0001$).

The result that truths are more frequent than lies is clearly consistent with hypothesis HT2a. At the same time, S-player behavior strongly rejects that there is a lie detection bias (i.e., H2La). Let us now look at the R-players. The S players' behavior could in principle be exploited by R-players; they could earn more by guessing that S-players were telling the truth more often than they are lying. However, they did not exploit this. Instead R-players guessed that S-players were telling the truth only 46 percent of the time. However, the difference is too small to make it possible to reject H2Tc at conventional levels of significance. In Figure 2b, the average proportion of R-players guessing that the statements were true in the different periods is given. Hence, R-player behavior is more consistent with the mixed strategy equilibrium and the lie detection bias (H1 and H2Lc) than with truth detection bias. There is also the prediction of earnings. According to both the truth and lie detection bias hypotheses, R would on average more often be correct than wrong in her guesses, which would make her earn more than A. However, in the ten rounds S-players earned about the same amounts as the R-players. S-players won 52.4 percent of the games and the R-players won 47.6 percent. Consequently, it is not possible to reject hypotheses H2Tb or

H2Lb. To sum up, the hypotheses of MSE and lie detection bias are rejected due to the S-players' behavior. The a-part of the truth bias hypothesis is supported by the data and it is not possible to reject any specific part of it. However, one cannot say that H2T is strongly and consistently supported by the data. Hence, neither MSE nor the theory of detection biases is strongly and consistently supported by the data.

5.2 Individual earnings

The next question concerns the individual differences with respect to bluffing and lying detection abilities. As noted previously if there are no differences in skills concerning lie detection or bluffing, differences in total earning should only be randomly generated and thus normally distributed. However, if there are skill differences this might make the earnings distribution non-normal.²² A Kolmogorov-Smirnov normality test does not reject the null-hypothesis of normally distributed total earnings. These tests are illustrated in Figure 3a-b with a line describing the expected frequencies of earnings given normality and the dots representing the actual observations. From the figures it can be verified that the dots do not deviate markedly from the line, which suggests that the data is approximately normally distributed. Hence, this observation does not suggest that there are large and consistent skill differences in bluffing or lie detection. Hypothesis 3 is supported by the data.

5.3 Beliefs about bluffing and lie detection capacities

Before the bluffing game each subject was asked to rate their abilities by confronting them with the following statements in the questionnaire:

S1: "I am good at bluffing"

S2: "I can tell when other people are lying"

S3: "Other people can tell when I am lying"

²² However, it should be stressed that skill differences might be present even if the distribution of total earnings is normal.

For each statement the subject could answer “Completely wrong”, “Fairly wrong”, “Somewhat wrong”, “Somewhat right”, “Fairly right” and “Completely right”. The alternatives were rated from 1 to 6 with the first getting 1 and so on to the last alternative getting 6. The average ratings were 3.4, 3.7 and 3.8 on S1, S2, S3 respectively. The distributions of how the subjects rated themselves are given in Figure 4a-c. First, it should be noted that these beliefs are hardly consistent with the standard assumption in economics that bluffing does not pose a problem. In fact, the majority of subjects say that S2 or S3 is “Somewhat right”, “Fairly right” or “Completely right.” Furthermore, the subjects rate themselves as somewhat better at detecting lies than at bluffing. In the bluffing game context this could suggest the subjects’ average belief is that the R-player has an advantage over the S-player. This belief is consistent with the hypothesis of a detection bias. Let us now see if these self-evaluations are related to their performance in the game. For S-players S1 and S3 are the relevant statements, since S-players had the opportunity to tell a truth or a lie. We find that both S1 and S3 are significantly related to their earnings at the 10 percent level with their expected signs. The correlation coefficient for S1 and S3 are 0.40 ($p = 0.07$) and -0.43 ($p = 0.05$). For R-players S2 is the relevant statement. Here, we find a weaker insignificant relationship, but with the expected sign. The correlation coefficient is 0.10 ($p = 0.65$). Together these results suggest that H4a should be rejected.

H4b concerns if R-players are prepared to forego money in order to have the S-player looking at the card if S was uninformed about the color of the card. We observed 14 R-players that had the opportunity to bid on the right to get the S-player look at the cards in five rounds. The average bids in the five rounds were 0.9, 1.5, 2.0, 0.7 and 1.0, which can be considered low. The distribution of all bids is given in Figure 5. Most subjects bid 0. The number of players making a strictly positive bid ranged from 3 to 5 in the periods. Often the same group of players made the positive bids.

H4c has to be evaluated based on only 7 subjects’ behavior in five rounds. However, there is an indication that the null hypothesis is correct for most subjects. Only one

subject's behavior deviated from the null hypothesis. This subject bid 10, 2, 5, 5 and 5 in the different periods.

To sum up, there appears to be a relationship between self-perceived abilities and performance, at least when it comes to bluffing. People who believe they are good at bluffing and who thinks that others cannot tell when they are telling the truth perform better than others in the bluffing game. Furthermore, although the vast majority of subjects are not prepared to pay money to get themselves into situations where they can feint or detect lies, some people are.

5.4. Personality traits and gender

The opportunity will also be taken to report on how personality traits and gender are related to the behavior in the bluffing game. These observations are not directly related to any hypothesis, but might be of some exploratory interest, since the experiment conducted in this paper appears to be new in the literature.

First, we ask if there are any personality traits that characterize a person that is good at bluffing. The correlation coefficients between the traits and the S-player and R-player earnings are reported in Table 1. It can be verified from the table that the relationships are not very strong. However, we find a significant negative correlation between altruism and bluffing performance. Possibly, people with strong altruistic feelings may have problems with bluffing in a convincing way. Another significant correlation is between cooperativeness and lie and truth detection. This positive correlation also makes sense. The ability to distinguish bluffs from truths ought to be more valuable for individuals oriented toward cooperation than for others.

Table 1. here

Another question concerns conceivable gender differences in bluffing. On average the subjects chose to bluff about 34.8 percent of the time. The corresponding percentages for males

and females were 34.6 and 35.0, respectively. The lack of gender differences also holds for the R players' inclination to guess that the S player was bluffing. Males were guessing that S was bluffing 55.8 percent of the time. The corresponding percentage for females was 53.3. Although, males were somewhat more inclined to suspect a lie, the difference is small and not significant.

The next question is then if there are gender differences in abilities to bluff and to detect lies. According to some researchers there might be reasons to expect this. For instance, Campbell (2002) writes that women are better than men "to conceal facial and bodily movements of emotions – for example when asked to pretend that a foul-tasting drink is pleasant or to express delight when they lose a competitive game." (p. 85). We did not detect any notable gender differences in the winning probabilities. Male S-players won 53.6 percent of the time and the figure for female S-players was 51.0. Male R-players were on average somewhat better in guessing the truth value of S's message. Thus, male R-players won 51.7 percent of the time and female R-players won 42.2 percent. This difference is however, not statistically significant. More interesting is that S-players are more likely to bluff if the R-player is a woman than a man. When the R-player was a woman S-players were bluffing 43.3 percent of the time. When the R-player was a man only S-players bluffed only 28.3 percent of the time. A chi-square test rejects the hypothesis that R-player females encounter the same proportion of lies as male R-players ($p = 0.024$). This finding appears to be new and can have important implications for gender differences in bargaining (e.g., about wages and positions).²³

Despite the fact that there is no significant gender difference in bluffing or lying detection abilities there is an interesting notable gender difference in the subjects' beliefs concerning their ability to bluff. Male subjects' average self-rating on the statement "I am good at bluffing" (see section 5.3) was 4.2 while the females' average rating 2.5. A Mann-Whitney test strongly rejects the null hypothesis that the male and female distributions of beliefs stem from the

²³ Although, the result in the bluffing context is new it is in the same spirit as the finding reported in Holm (2000), namely that subjects in a battle of the sexes experiment played significantly more "hawkish" against female co-players.

same underlying distribution. ($p = 0.002$).²⁴ Thus, while there is no significant gender difference in actual bluffing ability males are significantly more prone to believe that they are good at doing it. Some would not be surprised by this result.

6. Concluding Remarks

This paper aims to contribute to the study of bluffing in several ways. First, the implications of a possible lie or truth detection bias in bluffing are studied theoretically. This is done by developing a pure bluffing game with conflicts of interest. In such a game it is shown that the conventional interpretation of such a detection bias made by psychologists is inconsistent with standard game theoretical notions of equilibrium. However, if the bias is modeled as a perfect signal from Nature that is sent with a certain probability conditional on the truth value of the message, the bias can be incorporated into a game theoretical equilibrium analysis. One of the results from such an analysis is that the equilibrium set shrinks to a unique non-pooling equilibrium. In this equilibrium, the better a player is to detect lies the more often will the opponent player lie. This somewhat counter-intuitive prediction could, in principle, be used to reveal hidden information if the uninformed is able to make the informed believe in the detection bias.

Bluffing behavior is also studied in an experiment. In the bluffing game S-players choose to tell the truth significantly more often than they lied, something that is inconsistent with mixed strategy equilibrium, but consistent with the truth detection bias hypothesis. However, R-player behavior did not support the truth bias hypothesis. For instance, according to this hypothesis R-players should earn more than S-players, which was not the case. It is however, premature to make any general conclusions on the detection bias hypothesis from this single experiment. It needs to be tested in different contexts, where bluffing is endogenous. The context used in this experiment should be considered as relatively detection bias unfriendly. For instance,

²⁴ Males also believe that people are less likely to be able to tell if they are lying (see statement S3). However,

subjects did not know each other and the psychological cost in terms of e.g., guilt when lying in this very symmetrical game of pure conflict of interest with play cards ought to be very small. To induce truth detection bias it may be necessary that people know each other fairly well (so they know what cues to look for). It may also be that if lies or truths do not generate any strong emotions, no detectable signals will be emitted. Furthermore, one way to recognize lies is because it is cognitively more demanding to lie.²⁵ In this experiment there were no obvious cognitive difference between the task of lying and telling the truth. To sum up, the experiment in this paper does not provide a definitive answer to the question of truth detection bias. Instead the experiment can hopefully be considered as complementing psychological research on the issue and as a starting point of research into the question based on game theoretical foundations.

The experiment also generated results that might be of general interest to the study of bluffing behavior. One set of results concerns subjects' beliefs about their bluffing and lie detection abilities. Subjects had varying beliefs about their abilities in this respect. These beliefs are hard to reconcile with the standard assumption that bluffing is uncomplicated. The majority of subjects stated that there were some truth in the statements "I can tell when other people are lying" and "other people can tell when I am lying." A significant relationship between self-rated bluffing ability and actual performance in the game was detected. This suggests that people may know if they are good liars or not. Also, some personality traits were correlated with bluffing and lie detection. A significant negative correlation was found between altruism and bluffing performance. Furthermore, a significant positive relationship was found between cooperativeness and lie detection performance. Hence, altruists are not good liars and people who are cooperatively oriented are better than others in telling if a person is bluffing.

We also report gender effects. One finding was that subjects were significantly more prone to lie to a female opponent than to a male. Furthermore, even if there was no significant

this difference is not statistically significant ($p = 0.11$ for a corresponding Mann-Whitney test).

²⁵ Montaigne appears to have had this insight when he wrote "He who is not very strong in memory should not meddle with lying."

difference in bluffing performance between males and females, the former rated them selves as significantly better at bluffing than the latter.

Literature

- Bolton G. E., and Ockenfels, A., (2000). "A Theory of Equity, Reciprocity and Competition," *American Economic Review*, vol. 90(1), pp. 166-93.
- Brosig, J., "Identifying Cooperative Behavior: some experimental results in a prisoner's dilemma game," *Journal of Economic Behavior and Organization*, 47(2), pp. 275-290.
- Camerer, C., (2003). "Behavioral Game Theory, Experiments in Strategic Interaction," Princeton: Princeton University Press.
- Campbell, A., "A Mind of Her Own, The Evolutionary Psychology of Women," Oxford: Oxford University Press.
- Crawford, V., and Sobel, J., (1982), "Strategic Information Transmission", *Econometrica*, 50(6), 1431-51.
- Crawford, V., (1998), "A Survey of Experiments on Communication via Cheap Talk," *Journal of Economic Theory*, 78(2), 286-98.
- Crawford, V., (2003), "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 93:1, 133-149.
- Ekman, P. and O'Sullivan, M. (1991). "Who Can Catch a Liar?," *American Psychologist*, 46, 913-920.
- Ekman, P., (1992). "Telling lies: clues to deceit in the market place, politics and marriage." New York: W.W. Norton.
- Ekman, P., O'Sullivan, M., and Frank, M. (1999), "A Few Can Catch a Liar", *Psychological Science*, 10, 263-266.
- Fehr E., and Schmidt, K. M., (1999). 'A Theory of Fairness, Competition, and Cooperation', *Quarterly Journal of Economics*, vol. 114(3), pp. 817-68.
- Frank, R., (1987), "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?" *American Economic Review*, 77:4, 593-604.

- Frank, R., Gilovich, T., and Regan, D., (1993), "The Evolution of One-Shot Cooperation: An Experiment," *Ethology & Sociobiology*, 14, 247-256.
- Hendricks, K., and McAfee, P., "Feints", (2003), University of Texas, unpublished manuscript.
- Hennig-Schmidt, H., (2002) "The Impact of Fairness on Decision Making," in (eds.) Andersson, F., and Holm, H., "Experimental Economics: Financial Markets, Auctions, and Decision Making," Boston: Kluwer Academic Publishers.
- Holm, H., (2000), "Gender Based Focal Points," *Games and Economic Behavior*, 32, 292-314.
- Ockenfels, A., and Selten, R., (2000), "An Experiment on the Hypothesis of Involuntary Truth-Signalling in Bargaining," *Games and Economic Behavior*, 33, 90-116.
- Vrij, A., (2000), "Detecting Lies and Deceit, The Psychology of Lying and the Implications for the Professional Practice," John Wiley & Sons, Chichester.

Appendix

Written information and instructions to the subjects in the experiment

General information

Dear Participant,

You have been given the opportunity to participate in a study. The study includes experiments and questionnaires. The purpose of the study is to expand the knowledge of economic decision making.

You have the opportunity to earn money in the experiments. The money you will earn will be paid to you after a lecture next week. If you do not attend to that lecture you can receive your money at Håkan Holm's office (the Alpha building, 4th floor, Department of Economics). To pay your money we will need the ID-note that you have received.

Your answers on the questionnaires will be confidential. Furthermore, it will not be required of you to write your name on any of these questionnaires. To identify yourself without using name, you have received a number on an ID-note which you should keep as a receipt. **You will need this receipt when you collect your money. Thus, keep your ID-note to yourself and do not return it when you hand in the questionnaires.**

Yours sincerely,

Ola Andersson, Håkan Holm, Erik Wengström.

Information about the experiment

You will now participate in an experiment. Half of the group will be called A-players and the other half B-players. In each round, A will look at a play card that has either a red or black color. (Dames and hearts are counted as red and the spades and cloves are counted as black.) After this A is to tell the color of the card and B's task is to determine if A has told the true or false color. If B's guess is correct, he will earn SEK15 and if B's guess is incorrect A will earn SEK15. Hence, A will only earn money if B's guesses incorrectly and B will only earn money if B guesses correctly. The rules can be summarized by the following situations:

1. A tells the true color of the card and B guesses that A has stated the true color. B wins and earns SEK15 in this case.
2. A tells the true color of the card and B guesses that A has stated the false color. A wins and earns SEK15 in this case.
3. 1. A tells the false color of the card and B guesses that A has stated the true color. A wins and earns SEK15 in this case.
4. A tells the false color of the card and B guesses that A has stated the false color. B wins and earns SEK15 in this case.

In each round A will fill out a form with scores, which will be used when each player's earnings is to be calculated. The B-player will confirm with his ID-number that the scores are correct.

Form for Scores

Your ID-number: _____

To be able to determine the winner of the rounds you will here be asked to write down the true color of the card and if you have stated the true or false color of the card. You are also to state if your opponent has guessed that your statement was true or false. After this, your opponent will (with his own pencil) write his ID-number and thereby confirm the score.

Round	The true color of the card. (<i>black</i> = cloves, spades, <i>red</i> = dames, hearts).	Write <i>true</i> if you stated the cards true color. Otherwise write <i>false</i> .	Write <i>true</i> if the opponent guessed that you told the truth, write <i>false</i> otherwise.	The opponent's ID-number (written with his/her own pencil.)
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				

Information about the experiment

You are now asked to participate in an additional experiment. Half of the group will be called A-players and the other half B-players. The rules are the same as before, but with one difference. A is not allowed to look at the play card before A states the color of the card. However, in each round B has the possibility to pay for having A to look at the card before A tells the color of the card. To get this possibility B must state a maximum price between SEK0 and SEK10 that she/he is prepared to pay for having A to look at the card. The maximum price is not necessarily the price B will pay, it is the price B will pay at most. The price B will pay will be determined by a random draw in the following way. A random number between 1 and 10 will be drawn in each round. If the random number is smaller or equal to B's maximum price then B will pay the random number (in SEK) to the experimenters. In this case A must look at the card before she/he states the color of the card. However, if the drawn random number is larger than the maximum price then no transaction will occur and B will not pay any money to the experimenters. In this case A will not look at the card before she/he states the color of the card.

Note, if B is not interested in buying the possibility of having A to look at the card, then B should state a maximum price of 0. If B wants to be sure to have A to look at the card, he should state a maximum price of SEK10.

As in the previous experiment A will fill out a form with the scores, which will be used when each player's earnings is to be calculated. The B-player will confirm with his ID-number that the scores are correct.

Form for Scores

Your ID-number: _____

To be able to determine the payoffs you will here be asked to write down the price your opponent is prepared to pay for you to look at the card and the price drawn. You are also asked to write down the true color of the card and if you have stated the true or false color of the card. You are also to state if your opponent has guessed that your statement was true or false. After this, your opponent will (with his own pencil) write his ID-number and thereby confirm the score.

Round	The price the opponent is prepared to pay	The randomly drawn price	Write <i>yes</i> if you had to look at the card and write <i>no</i> otherwise.	The true color of the card. (<i>black</i> = cloves, spades, <i>red</i> = dames, hearts).	Write <i>true</i> if you stated the cards true color. Otherwise write <i>false</i> .	Write <i>true</i> if the opponent guessed that you told the truth, write <i>false</i> otherwise.	The opponent's ID-number (written with his/her own pencil.).
<i>Exercise round</i>							
1							
2							
3							
4							
5							

Figures




Figure 1: Reduced bluffing game with detection bias.




Figure 2a. Average proportion truth telling by S-players in the different rounds.




Figure 2b. Average proportion truth guessing by R-players in the different rounds.

Normal Probability Plot

EES 2004 : Experiments in Economic Sciences - New Approaches to Solving Real-world Problems




Figure 3a. A normality test. The line displays the expected frequencies for normally distributed earning and the dots the actually observed distribution of total earnings for A players.

Normal Probability Plot




Figure 3b. The line displays the expected frequencies for normally distributed earning and the dots the actually observed distribution of total earnings for B players.




Figure 4a. Distribution of ratings on S1.




Figure 4b. Distribution of ratings on S2.




Figure 4c. Distribution of ratings on S3.




Figure 5. Histogram of R-player bids.

Figures in Appendix




Figure A.1: Bluffing game without detection bias.




Figure A.2: Bluffing game with detection bias.

Tables

Earnings	Impulse Control	Altruism	Cooperativeness	Risk-taking	Extraversion	Excitement-seeking
S-player	-0.29	-0.40*	-0.13	0	-0.15	-0.09
R-player	-0.11	-0.25	0.38*	-0.08	0.05	-0.16

Table 1: Correlation coefficients between personality traits and earnings in the bluffing game. * indicates that the correlation coefficient is significantly different from zero at the 10-percent level.