

Lundquist, Tobias; Ellingsen, Tore; Gribbe, Erik; Johannesson, Magnus

Working Paper

The cost of lying

SSE/EFI Working Paper Series in Economics and Finance, No. 666

Provided in Cooperation with:

EFI - The Economic Research Institute, Stockholm School of Economics

Suggested Citation: Lundquist, Tobias; Ellingsen, Tore; Gribbe, Erik; Johannesson, Magnus (2007) : The cost of lying, SSE/EFI Working Paper Series in Economics and Finance, No. 666, Stockholm School of Economics, The Economic Research Institute (EFI), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/56266>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Cost of Lying

Tobias Lundquist, Tore Ellingsen, Erik Gribbe and Magnus Johannesson*
Department of Economics, Stockholm School of Economics

SSE/EFI Working Paper Series in Economics and Finance
No. 666
June, 2007

Abstract

We experimentally investigate the effect of cheap talk in a bargaining game with one-sided asymmetric information. A seller has private information about his or her skill and is provided an opportunity to communicate this information to a buyer through a written message. Four different treatments are compared; one without communication, one with free-form communication, and two treatments with pre-specified communication in the form of promises of varying strength. Our results suggest that lying about private information is costly and that the cost of lying increases with the size of the lie and the strength of the promise. Freely formulated messages lead to the fewest lies and the most efficient outcomes.

Key words: Deception, communication, lies, promises, experiments.

JEL Codes: C91, D82.

Acknowledgement: We are indebted to Tomas Östlund, Helena Samuelsson and Jonas Öberg for constructing the software application *Xgame* with which we conducted the experiment. We also thank Lars Gribbe, Daniel Kjellander, Alf Thiel Metelius, Jon Tillegard and Jakob Stolpe for research assistance. Erik Lindqvist, Eva Ranehill, Jens Sandblom, Robin Stenskog and Herman Stål provided helpful comments. Finally we thank the Jan Wallander and Tom Hedelius Foundation and the Swedish Research Council for financial support.

* Corresponding author: Department of Economics, Stockholm School of Economics, Box 6501, SE-113 83 Stockholm, Sweden; e-mail: magnus.johannesson@hhs.se

1. INTRODUCTION

In this paper we examine a number of issues concerning the credibility of communication. More specifically, we are interested in the role of promises and lies. According to Immanuel Kant (1785), lying is always wrong, since we can not justify the permissibility of lying as a universal law. This view is in stark contrast to the conventional assumption in economic theory that lying is permitted, costless and often the rational choice. As lying is costless, people will behave opportunistically and lie at will. Since everybody is aware of this, nobody has any reason to believe what is conveyed in communication. Communication is consequently referred to as “cheap talk” (Crawford and Sobel 1982).

Introspection and casual observations suggest that neither of these extreme positions is a true description of human behavior. Experimental evidence firmly rejects that communication is just cheap talk (see the overviews of experimental results in Ledyard 1995; Crawford 1998; and Camerer 2003). In a meta-analysis of prisoners’ dilemma experiments, Sally (1995) found that communication has a strong positive effect on cooperation and that communication was the single most effective factor in increasing cooperation. He furthermore found that elicited promises had a strong additional effect on cooperation over and above the effect of communication per se. An effect of communication has also been found in different bargaining games. For example Radner and Schotter (1989) and Valley et al. (2002) found that non-binding pre-play communication increase trade in a double auction.

The fact that non-binding communication does matter must spring from the fact that it is credible to a certain extent. This, in turn, suggests that there exists a cost of misrepresenting the truth as argued by several authors (e.g. Frank 1987, 1988; Hirshleifer 1987; Ostrom et al.

1992; Klein and O'Flaherty 1993; Braver 1995; Kerr 1995).¹ If the cost of lying is sufficiently high defection will no longer be optimal in a prisoners' dilemma for an individual that has communicated that she will cooperate. However, it is difficult to draw firm conclusions about the cost of lying from the experimental literature on social dilemmas. Other factors than a cost of lying may also explain why communication increases cooperation in social dilemmas. With social preferences cooperation may be an equilibrium and communication may help to coordinate on this equilibrium (Sen 1967). The communication in itself may also affect social preferences, such as the degree of empathy towards the other player.

Gneezy (2005) recently carried out a more direct test of the cost of lying. He studied a simple sender-receiver game where lies increase the payoffs to the liar and decrease the payoff to the other party. He found that lying is costly (i.e. not everybody lied), and that people's inclination to lie increases the more they stand to gain from the lie and decreases the more others lose from it.

In this experimental study we provide further evidence on the cost of lying along two dimensions not studied by Gneezy (2005). We test whether the cost of lying depends on the size of the lie, i.e. that the lie is more costly the further you deviate from the truth. We furthermore test if the cost of lying depends on the degree of the lie, i.e. the strength of the promise. Previous work suggests that promises increase the credibility of communication (Sally 1995; Ellingsen and Johannesson 2004; Charness and Dufwenberg 2006).

In the experiment a seller has private information about her talent and can send a written message to the buyer. After receiving the message, the buyer decides whether or not to offer the seller a fixed-payment contract. If a contract is signed, the buyer makes a profit if

¹ The idea that communication can create commitment is related to the broader literature in psychology, which emphasizes people's desire for consistency, see e.g., Heider (1946), Newcomb (1953), Festinger (1957) and Cialdini (1993).

the seller's talent is above a certain threshold and a loss otherwise. Sellers with a talent below the threshold have an incentive to lie about their talent to ensure a contract. Four different treatments are compared: one without communication, one with free-form communication and two treatments with pre-specified communication in the form of promises of varying strength. Consistent with Gneezy (2005) we find that lying is costly. Our results also suggest that the cost of lying increases with the size of the lie and the strength of the promise. Furthermore, free-form communication leads to the fewest lies and the most efficient outcomes.

2. EXPERIMENTAL DESIGN

Gneezy (2005) tested if the propensity to lie depends on the potential gain of the liar and the potential loss of the other party. In our set-up we hold constant for the monetary consequence of lying and focus on two different dimensions of the cost of lying: the size of the lie and the strength of the promise that goes with it.

2.1 Subjects and procedure

Four experimental treatments were carried out. In the first treatment (NC) no communication was allowed. This treatment was included as a control group to measure the efficiency effects of communication. In two treatments (P1 and P2) the seller could send a pre-specified message to the buyer. The degree of promise was varied between these two treatments (see below). In the final treatment (FC) the seller could send a freely formulated message to the buyer.

In all treatments, an equal number of subjects were placed in two adjacent rooms, without the possibility to communicate visually or verbally with the other room, and asked to

take a place at a computer. The rooms were labeled 3 and 4.² The experiment was completely computerized.³ In Figure 1 we depict the timeline in the experiment. The experimental instructions are provided in the Appendix. [insert Figure 1 about here]

All subjects first performed a test consisting of 55 questions of mixed general knowledge, word power and logical reasoning character.⁴ The results of the subjects in each room were ranked separately and the subjects were given a score between 1 and 100 based on this rank. The score followed a uniform distribution where subjects in the lowest tenth percentile were given a random score between 1 and 10, the second lowest tenth percentile a random score between 11 and 20, and so on up to a hundred. A larger amount of correct answers always gave a higher score. The distribution was communicated to the subjects on their screens. The participants were then informed that they were to decide whether or not to enter into a contract with an anonymous participant in the other room, whose identity they would never know, and presented with the payment structure in Figure 2. In order to be sure that subjects understood the payment structure, they could not move on to the negotiations before correctly answering a number of control questions. [insert Figure 2 about here]

In the no communication treatment, subjects simply checked either a box reading “enter into contract” or a box reading “do not enter into contract”. In the first treatment with pre-specified messages (P1), the subjects in room 3 (the sellers) were given the option either to send a message to their counterpart in room 4, claiming that “my score is X”, or not to send a message, in which case the subject in room 4 received the message “A has chosen not to send a message”. In the second treatment with pre-specified messages (P2), the available message read “I solemnly promise that my score is X”. In the free-form communication the seller

² These are the actual names of the computer labs. We felt it would be confusing to rename the rooms 1 and 2.

³ The program, entitled *XGame*, was created especially for this experiment.

⁴ The test was devised by us.

could send a written message (or refrain from sending any message). The only restriction on the message was that they do not reveal anything that could jeopardize anonymity. After the message was sent and read, subjects chose whether to enter into contract or not as in the no communication treatment.

To increase the number of observations collected in the experiment the roles were then reversed, but with new randomly assigned partners, the subjects in room 3 now acting as buyers rather than sellers. The procedures above were then repeated in the same way as for the first round. The participants were not informed of the results from the first round until the entire session was completed.

After the second round was completed, the subjects answered a number of follow-up questions including questions about beliefs about the propensity of liars. In the belief questions subjects were asked to guess how many percent of those with a score below the cut-off they thought lied about their score and how many percent of those with a score above the cut-off they thought lied about their score. If subjects guessed these fractions correctly, within five percentage units, they were paid an additional SEK 50.⁵ The follow-up questions also included a number of questions concerning the cost of lying to shed further light on our hypotheses. After the follow-up questions subjects were paid and allowed to leave. The payment to the subjects in room 3 was placed into sealed envelopes by the instructor in room 4 and vice versa, to ensure anonymity towards the experimenters.

We carried out seven sessions of the experiment within a week in January and February 2005: one session of the control group and two sessions each of the three communication treatments. Each session consisted of 40 subjects (20 in each room), with the exception of one treatment with free-form communication with 38 subjects (as some subjects failed to show

⁵ They were also asked to guess what everybody else answered to these two questions, but these responses were almost perfectly correlated with their own beliefs.

up).⁶ Thus, in total 278 subjects participated in the experiment, and they were given a show up payment of 50 Swedish kronor (about \$7 at the time of the experiment). The subjects were recruited from The Stockholm School of Economics, Stockholm University and The Royal Institute of Technology.⁷ Students from the Stockholm School of Economics were undergraduates in business and economics and the other subjects were undergraduates from various departments such as economics, philosophy, law, political science etc. The gender distribution and the subject affiliation of the subjects are shown in Table 1. Slightly more than 50% of the subjects were men, and the subjects were about evenly divided between the Stockholm School of Economics and other affiliations.

2.2 Design considerations

We chose to let the participants score depend on their performance in a test, rather than just letting them receive a score randomly. We wanted the subjects to feel that they had “earned” their score to add realism to the situation. Had the test not been included, and the scores had been distributed at random, we believe subjects would have been more detached and more likely to lie. Konow (2000) has shown that when exogenous factors, i.e. factors that are beyond the individual’s influence, in this case chance, affect the allocation of income, this is perceived as unjust. A seemingly unfair proceeding when distributing the scores might prejudice the participants towards behaving deceitfully.

Several experiments show that varying the degree of anonymity significantly changes outcomes (see e.g. Hoffman et al. (1996), and Charness and Gneezy (2003)). In order to avoid

⁶ With 20 subjects in each room the fraction of buyers above the cut-off of 70, will be exactly 30%. In the session with 19 subjects in each room the probability for a buyer to be matched with a seller above the cutoff was slightly different from 30 %. It was 32 % in turn 1 and 26 % in turn 2.

⁷ They were recruited via e-mail, posters and presentations in connection to classes.

the possibility of soliciting responses that are motivated by social proximity or reputational concerns, we made the procedure anonymous towards both other subjects and the experimenters. The fact that computerized communication was used should further increase anonymity, in that subjects are unable to infer anything from handwriting.

We wanted to have a cut-off for the test score (the score for which a contract is beneficial for both sellers and buyers), such that the expected payoff from a contract is negative for the buyer (as this would give some scope for communication to improve the outcome). This implies a cut-off of over 50. As we are primarily interested in deceptive lies we also wanted to ensure that a large fraction of subjects have a pecuniary argument to lie. We therefore set the cut-off at 70; where 70% of subjects were below the cut-off.

2.3 Hypotheses and tests

Following the classification of lies by Gneezy (2005) we distinguish between two types of lies in our data. A seller with a score below the cut-off may lie and claim to have a score above the cut-off. We call this a deceptive lie as the aim is to increase the payoff of the liar and decrease the payoff of the receiver. A seller with a score above the cut-off may also lie about the score to increase the likelihood of a contract (if a specific score above the cut off is perceived as more plausible than the actual score). This is called a white lie as the aim is to increase the payoff of both parties.

Our first hypothesis to be tested is that the cost of lying depends on the degree of lying, i.e. the degree of promise that goes with the lie. The hypothesis implies that a stronger promise increases the cost of lying and therefore decreases the likelihood of lying. This is tested by comparing the fraction of lying between treatment P2 and P1, with the hypothesis that the fraction of lying is lower in the treatment with stronger promises (P2). The fraction of

lying is also compared between free form communication (FC) and the two treatments with fixed communication (P1 and P2), but here we do not have an a priori hypothesis concerning the direction of the effect. To compare the fraction of liars between treatments we use a Pearson contingency table chi-square test (D'Agostino et al. 1988). To further evaluate the first hypothesis we also test if the beliefs about lies differ significantly between the three communication treatments. We use bootstrap techniques to compare mean beliefs, as this variable was not normally distributed.^{8,9}

Our second hypothesis is that the cost of lying depends on the size of the lie, i.e. on how much the lie deviates from the truth. This hypothesis implies that sellers that are below the cut-off should be less likely to lie the further below the cutoff they are. To test if the size of the lie affects the propensity to lie we use logistic regression analysis, with lying (yes/no) as a function of the test score, controlling for the experimental treatment. We also control for the turn (whether the subject is the buyer in the first or the second turn) and gender and subject affiliation (Stockholm School of Economics versus other affiliations). Gender and student affiliation has been shown to be important variables in other experiments.

If lying is costly, the value of communication will increase as sellers below the cut-off are less likely to lie and claim that they have a score above the cut-off. Knowing this, the buyer will be more inclined to offer a contract. We therefore also test the hypothesis that communication increase efficiency. This third hypothesis is tested by comparing the efficiency between the communication treatments and the no communication treatment. If a stronger promise increases the cost of lying it also follows that strong promises increases

⁸ According to a Kolmogorov-Smirnov test the null assumption of normality can be rejected at the 5% level for both the beliefs about deceptive lies and the beliefs about white lies in all four treatments. The only exception is the beliefs about deceptive lies in the free-form communication treatment ($p=0.129$).

⁹ Bootstrap techniques make it possible to conduct statistical testing without imposing normality, i.e. by inferring the underlying distribution from which the data has emerged. For an introduction to bootstrap methods,

efficiency more than weak promises. This fourth hypothesis is tested by comparing the efficiency between treatment P2 and treatment P1. Efficiency is also compared between the free form communication (FC) and the two treatments with fixed communication (P1 and P2), but here we do not have an a priori hypothesis concerning the direction of the effect. Efficiency is analyzed on a pairwise data level. The fraction of contracts (pairs in which a contract is agreed upon) and the fraction of successful contracts (pairs in which a contract is agreed upon and the test score of the seller is over 70) is compared between treatments with a Pearson contingency table chi-square test. The mean payoff to the pair is compared between treatments using bootstrap techniques, as this variable was not normally distributed.¹⁰ The significance levels for all bootstrap tests have been obtained by generating 2,099 bootstrap replications (Davidson and MacKinnon 2000). All reported p-values are two-sided.

3. RESULTS

The results section is divided into a section that looks at lying (individual level data) and a section that reports the results about efficiency (pairwise-data). The section about lying is furthermore subdivided into a section about the behavioral results and a section about the beliefs and the questionnaire data. We refer to sellers with a score on the test below the cut-off as “low-talent” and sellers with a score on the test above the cut-off as “high-talent”.

3.1 Lies and deception

see e.g. Efron and Tibshirani (1993). Ellingsen and Johannesson (2004) discuss the choice of test statistics in more detail.

3.1.1 Experimental results

In Figure 3 we show a scatterplot of all the observations. The observations on the 45-degree line in the scatterplot in Figure 3 are senders who state their true value, above the line are senders who lie and state a higher value than what is true. Many low-talent sellers choose to lie and state a higher value than what is true, but many also refrain from sending a message; these are shown on the horizontal line. Looking more closely at the deceptive lies we can see that many liars state a value that is around 0-15 points over the cutoff.

Only three high-talent sellers lied about their score, limiting the possibility to find a difference between the treatments in white lies. All these three liars stated a lower value than they had. However, high-talent sellers only have a pecuniary incentive to exaggerate their test scores if a higher score is more likely to be accepted as a true score by the buyer. To test this we ran a logistic regression analysis between the test score of high-talent sellers and the probability that the buyer would offer a contract (including only those high-talent sellers that communicated a specific score). The variable for the test score was positive, but not significant ($p=0.149$ without a control for treatment and $p=0.228$ with controls for treatment). The point estimate therefore suggests a pecuniary argument for white lies, but the effect is not significant and we cannot reject the null hypothesis of no relationship. The insignificant relationship may explain the low fraction of white lies. [insert Figure 3 about here]

The deceptive lies of low-talent sellers are further illustrated in Figure 4, which shows the points sent of all low-talent sellers who lied. In the free-form communication treatment some subjects who lied did not specify a specific score. These are illustrated separately in the Figure and are of the form “we will both gain from signing a contract”. The figure shows that

¹⁰ The Kolmogorov-Smirnov test rejects the normality assumption in all four treatments ($p<0.001$ in all treatments).

most liars state a value close to 70, while very few state a value above 85. The fraction of deceptive lies is highest with weak promises, followed by strong promises and lowest in the free-communication treatment. As can be seen in the figure some low-talent sellers who lie send a score below the cut-off (5 subjects). These observations are difficult to classify, as they are not deceptive lies according to the definition of deceptive lies above (as they do not send a score above the cut-off). However, in the analysis below these observations have been included as deceptive lies in order not to lose any observations and because the subjects are clearly lying. Excluding these five observations does not change the conclusions reported below.¹¹ [insert Figure 4 about here]

In Table 2 we report the number of white and deceptive lies in the three communication treatments. First it should be noted that the fraction of white lies and deceptive lies significantly differ from 100% in all treatments, consistent with the idea that lies are costly.¹² As there are only three white lies across all three communication treatments we cannot expect any significant differences here. However, for deceptive lies the fraction of lies varies substantially between the three communication treatments. The fraction of deceptive lies with free-form communication (40%) is significantly lower than with pre-specified weak promises (77%) or prespecified strong promises (64%). We had no prior hypothesis concerning the direction of this effect. However, the results suggest that it is more costly to lie in free-form communication. To test the hypothesis that the cost of lying increases with the strength of the promise, we compare the fraction of lies between the treatment with strong promises and the

¹¹ Three of these subjects are in the strong promises treatment and two are in the weak promises treatment. Excluding these observations decreases the difference in the fraction of liars between free-form communication and pre-specified communication somewhat, but the difference is still significant at the 5% level. It also strengthens the association between the “points on test” variable and the probability of a deceptive lie in the logistic regressions in Table 3 somewhat (the p-value decreases to 0.076 in regression 3 in Table 3).

¹² We use a one-sample t-test to test if the fraction of liars differs from 100%. The p-value is below 0.001 in the t-test in all cases; for white lies with free-form communication a formal p-value cannot be estimated as there are no lies and therefore no standard deviation to compute the t-value.

treatment with weak promises. The effect goes in the expected direction, but does not quite reach statistical significance ($p=0.147$). [insert Table 2 about here]

Our second hypothesis to be tested is that the cost of lying increases with the size of the lie. With the setting of our experiment this implies that low-talent sellers should be more inclined to lie the higher their score (as a smaller lie is needed to communicate a score above the cut-off). Figure 5 gives some first tentative evidence. Figure 5 shows the fraction of liars in each decile of the test score distribution. There is a tendency for the fraction of liars to increase when the test score gets closer and closer to the cut-off. [insert Figure 5 about here]

In the logistic regressions in Table 3 we evaluate this hypothesis more formally. In regression 1 in Table 3 we test if the points on the test for low-talent sellers affect the propensity to lie, controlling for the experimental treatment. The test score is significant at the 5% level, consistent with the hypothesis that the cost of lying depends on the size of the lie. According to the marginal effect, the probability of lying increases with 0.4 percentage units for a unit increase in the test score. The dummy variables for the weak and strong promises treatments are also significant at the 5% level consistent with the non-parametric test results. As before the difference between the two promises treatments is not significant ($p=0.144$). [insert Table 3 about here]

In the second regression we add a control for if the low-talent subject was a seller in the first or the second turn. This has little effect on the other variables, but the turn variable itself is significant with a positive sign. This implies that the probability to lie is significantly higher in the second turn. In the third regression we also add controls for gender and subject affiliation. This lowers the effect of the test score somewhat, but the variable is still significant at the 10% level. Women are less likely to lie than men, but the effect is not significant ($p=0.161$). Subjects from the Stockholm School of Economics (SSE) are more

likely to lie than other subjects, and this effect is significant at the 10% level. To explore this effect further we added interactions between subject affiliation and all other variables. Only one of these interactions was significant at the 10% level; the effect of turn was significantly higher among SSE subjects. However, the five interactions were jointly significant at the 5% level according to a likelihood ratio test.¹³ In Table 3 we therefore also present separate regression equations for the two subject pools. Among SSE subjects the test score variable is not significant, and the turn variable is highly significant with a marginal effect of 30 percentage units. For other subjects the results are more encouraging. The test score variable is significant at the 5% level. In this sample the probability to lie is also significantly lower in the strong promises treatment than in the weak promises treatment, consistent with the first hypothesis. The turn variable has no significant effect in this sample. Overall the regression results support the hypothesis that the cost of lying increases with the size of the lie. The effect is significant for the combined sample and is significant among non SSE subjects.

In Figure 6 we show the predicted effect of the test score on the propensity to lie in the three communication treatments based on the third regression equation.¹⁴ For free-form communication the estimates predicts that the propensity to lie is almost twice as high for a test score of 70 (the highest possible score below the cut-off) compared to a test score of 1 (the lowest possible score). [insert Figure 6 about here]

3.1.2 Beliefs and questionnaire data

¹³ The chi-square value of the likelihood ratio test was 11.537 (5 DF), and the critical value on the 5% level is 11.07.

¹⁴ We also tested adding interactions between the test score variable and the experimental treatments, but these interactions were not significant individually or jointly at the 10% level in any of the regression equations in Table 3.

To complement the behavioral results in the experiment we also collected data on beliefs and questionnaire data about attitudes to deception. The beliefs about deceptive lies and white lies are presented in Table 4. Subjects expected few white lies in all three communication treatments; the average estimate is about 10% in all three treatments.¹⁵ For deceptive lies subjects believed that 79% of low-talent sellers lied with pre-specified weak promises, which is close to the actual fraction of deceptive lies in that treatment of 77%. With pre-specified strong promises subjects believed that 70% lied, which is also relatively close to the actual fraction of 64% in that treatment. The difference in beliefs between the weak and strong promises treatments is significant at the 10% level consistent with our first hypothesis. In the free-form communication subjects believed that 68% of low-talent sellers lied, which is substantially higher than the actual fraction of 40%. This fraction is significantly lower than the beliefs in the weak promises treatment, but does not differ significantly from the beliefs in the strong promises treatment. [insert Table 4 about here]

The beliefs about deceptive lies are strongly correlated with the propensity to lie for low-talent sellers. The Pearson correlation is 0.475 and the Spearman nonparametric correlation is 0.444. If the beliefs about deceptive lies are added to the regression equations in Table 3 it is also strongly significant with a positive sign. This pattern is consistent with that the cost of lying is a function of the perceived fraction of other individuals that would lie in the same situation. It is also consistent with the concept of guilt aversion of Charness and Dufwenberg (2006), i.e. that we do not want to let down other people's expectations. However, it is difficult to infer any causation from this result. Subjects may simply form beliefs based on their own actions or rationalize their decision to lie by reporting a high belief

¹⁵ These results may be somewhat inflated by the incentive used to elicit beliefs. Subjects were paid SEK 50 if their beliefs were within five percentage units of the correct estimate. A subject that wants to maximize the chance of winning the incentive payment should therefore never reply less than 5%, even if she thinks that less than 5% actually lied (as that will lower the upper bound of the "payment interval").

about the fraction of liars. This is referred to as the “false consensus effect” in psychology (Ross et al. 1977).

Table 5 presents answers to several questions about deception in the questionnaire. A majority of subjects seem to agree that white lies are positive and a majority of subjects claim that they would lie to help someone else even if it hurts them. About 80% of subjects disagree to the statement “I could lie to get revenge even if it hurts me”, indicating that lies that hurt both parties should not be all too common. When it comes to deceptive lies, 64% agree that they are more inclined to lie, the more they stand to gain from it and 82% also agree that they are less inclined to lie the more the other person has to lose. The two latter results are consistent with Gneezy’s (2005) findings. [insert Table 5 about here]

Nearly all subjects are less inclined to lie the greater is the risk of discovery. About 80% of subjects believe that there are degrees of lying and 88% agree that they are less inclined to lie the greater the lie has to be in order to be believed. This is consistent with the hypothesis that the cost of lying increases with the size of the lie. A great majority, 90%, agree that having given a promise makes it more difficult to lie, and an equal proportion agree that there are degrees of promises. This supports our hypothesis that promises increase the cost of lying and that different promises affect this cost to different degrees.

Three questions are also concerned with to what extent the subjects have lied in different circumstances in real life. About a third of the participants say that they have lied at least once in applications, and a somewhat higher fraction say that they have lied when selling something. Less than 20% say that they have lied in their income tax returns. These results further reinforce that lying is not costless.

3.2 Efficiency

The results concerning lies and deception in the experiment suggest that communication may improve efficiency, as many low-talent subjects refrain from lying. In Table 6 we report the efficiency results.¹⁶ Somewhat surprisingly the fraction of contracts is 30% in the no communication treatment. This is surprising as the expected profit from accepting a contract is negative for the buyer. However, if the buyer cares about the payoff of the seller, i.e. wants to maximize efficiency, it is rational to accept the contract. Some experimental results suggest that such efficiency/altruism concerns may be important (Charness and Rabin 2002; Engelmann and Strobel 2004). Possibly some buyers also find it more fun/exciting to accept the contract to see what happens in the experiment. [insert Table 6 about here]

With weak promises the fraction of contracts is 22% and with strong promises the fraction of contracts is 30%. In the free-form communication the fraction of contracts is somewhat higher at 37%. This difference is significant compared to the weak promises treatment, but not compared to the other two treatments. The fraction of successful contracts where the seller is a high-talent seller is 25% without communication. This is close to the expected 30% rate (as 30% of sellers are above the cut-off). This fraction increases somewhat to 39% and 42% with weak and strong promises respectively, but this effect is not significant. With free-form communication the fraction of successful contracts is 55%, which is more than twice as high as without communication. This effect is significant at the 10% level.

Looking at the mean payoff to the pair combines the effect of communication on the fraction of contracts and the fraction of successful contracts. The lowest possible mean payoff to the pair is SEK 50 (as the buyer earns SEK 50 with no contract), and the highest possible mean payoff to the pair is SEK 80 (full efficiency). The payoff is lowest in the no communication treatment (SEK 57.5), and is only marginally improved with pre-specified

¹⁶ Note that as we investigate efficiency at the pairwise data level, individual characteristics like gender and subject affiliation cannot be analyzed here.

communication with weak promises. With strong promises the payoff increases somewhat to SEK 62.5, but this is not significantly higher than with weak promises or no communication. The payoff is highest in the free-form communication (SEK 70.5), and this is significantly higher than with no communication or communication with weak promises. We therefore find support for the third hypothesis that communication improves efficiency. However, for the fourth hypothesis, that strong promises improve efficiency more than weak promises, we cannot reject the null hypothesis.

4. DISCUSSION

Even if people are not as adamantly opposed to lying as Kant (1785), they are not completely opportunistic either. When subjects are given a pecuniary incentive to lie about their test score in our experiment, between about 40% and 80% lie depending on the form the communication takes. These figures are clearly below 100%, indicating that at least some subjects find it costly to lie. This result is in line with the recent work of Gneezy (2005).

Our results also suggest that the cost of lying increases with the size of the lie. Subjects are more likely to lie when a small lie is sufficient to claim a score above the cut-off, than when a large lie is necessary. It may be argued that the test will lead to selection-bias of participants with larger or smaller inclination to lie into certain score brackets, i.e. that there is a correlation between aptitude on this particular test and the tendency to lie. Although this cannot be ruled out, we find it unlikely. We are not aware of any previous evidence that people with more general knowledge or a larger vocabulary should be more (or less) prone to lie. The population is also fairly homogenous; the students all belong to faculties where a high grade point average is required for admittance.

When it comes to the form of the communication we find strong evidence that lying is less likely in free-form communication than in pre-specified communication. Consistent with this, free-form communication also had the strongest effect on efficiency. As most communication in real life is in free-form rather than in terms of restricted pre-fabricated messages, this result implies that using pre-fabricated messages in experiments will not capture the full effect of communication.

Our evidence on the effect of the strength of promises on the cost of lying is less strong. The results suggest that the cost of lying increases with the strength of the promise, but for the overall sample this effect do not quite reach statistical significance. However, the effect is significant for the sample of students that are not from the Stockholm School of Economics (non-SSE students). The beliefs about deceptive lies also differed significantly between weak and strong promises. The hypothesis is also backed up by taking a closer look at the messages sent in the free-form communication treatment. A total of four senders explicitly used the expressions “I promise” or “I swear” and all four of these were truthful messages. Thus, it would seem that such expressions are not used readily when lying. Formulations similar to the one used by us in the strong promises treatment are also used in practice to induce truth telling. A large number of institutions make you promise upon “your honor and your conscience” that you are presenting the true state of affairs, for instance when submitting your tax returns, when testifying in court, or when applying for reimbursement from the state health insurance. If these statements did not increase commitment, they would be completely superfluous. We thus find some support for the idea that this might actually work.

Overall our results are much stronger among non-SSE students. These students lied less, which may be the principle explanation for the difference in results. If the propensity to lie is very high, it becomes more difficult to find significant differences. That the SSE students lie more often suggests that these students are closer to “homo economicus”. This is in line with

the results of Carter and Irons (1991) who found that economics students are more selfish than other students due to a selection effect. However, other studies comparing student populations present mixed results (see the overview in Camerer 2003). For SSE students the propensity to lie also increased significantly between the first and the second turn; even though roles were reversed with new partners and the outcome of each round was not revealed until after the completion of the second turn, thus leaving no scope for learning (apart from the repetition of instructions). The motivation for letting subjects swap roles (with new partners) was to double the number of observations. The follow-up questions indicate that some subjects seem to be convinced that they were cheated in the first round, and therefore lie in the second round. This suggests that beliefs about the fraction of other subjects that lie in the same situation are important for the cost of lying. It also cannot be ruled out that some SSE subjects tried to act strategically in the first round by not lying; even though nothing in the instructions hinted towards the possibility of reputation effects being important. On the contrary it was explicitly stated before the experiment started that there would only be two rounds of the experiment and that each subject would be randomly matched with another subject in each round.

REFERENCES

Bolton, G. and Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review* 90, 166-193

Braver, S.L. (1995). Social contracts and the provision of public goods. In: Schroeder, D. (Ed.). *Social Dilemmas: Perspectives on Individuals and Groups*. New York: Praeger.

Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton: Princeton University Press.

Carter, J.R. and Irons, M.D. (1991). Are economists different and if so, why? *Journal of Economic Perspectives* 5, 171-177.

Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117, 817-869.

Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica* 74, 1579-1601.

Charness, G. and Gneezy, U. (2003). What's in a name? Anonymity and social distance in dictator games and ultimatum games. Working paper, Department of Economics, University of California at Santa Barbara.

Cialdini, R.B. (1993). *Influence: Science and Practice*, 3rd edition. New York: Harper Collins.

Crawford, V. and Sobel, J. (1982). Strategic information transmission. *Econometrica* 50, 1431-1451.

Crawford, V. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory* 78, 286-298.

D'Agostino, R.B., Chase, W. and Belanger, A. (1988). The appropriateness of some common procedures for testing the equality of two independent binomial populations. *American Statistician* 42, 198-202.

Davidson, R. and MacKinnon, R.G. (2000). Bootstrap tests: how many bootstraps? *Econometric Reviews* 19, 55-68.

Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior* 47, 268-298.

Efron, B. and Tibshirani, R.J. (1993). *An Introduction to the Bootstrap*. Monographs on Statistics and Applied Probability, No. 57. New York: Chapman and Hall.

Ellingsen, T. and Johannesson, M. (2004). Promises, threats and fairness. *Economic Journal* 114, 397-420.

Engelmann, D. and Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review* 94, 857-869.

Fehr, E. and Schmidt, K. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114, 817-868.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford: Stanford University Press.

Frank, R.H. (1987). If homo economicus could choose his own utility function, would he want one with a conscience? *American Economic Review* 77, 593-604.

Frank, R.H. (1988). *Passions within Reason*. New York: W.W. Norton.

Gneezy, U. (2005). Deception: the role of consequences. *American Economic Review* 95, 384-394.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology* 21, 107-112.

Hirshleifer, J. (1987). On the emotions as guarantors of threats and promises. In: Dupré, J. (Ed.). *The Latest on the Best: Essays in Evolution and Optimality*. Cambridge MA: MIT Press.

Hoffman, E., McCabe, K. and Smith, V. L. (1996). Social Distance and Other Regarding Behavior in Dictator Games. *American Economic Review* 86(3), 653 -660.

Kant, I. (1785). *Fundamental Principles of the Metaphysics of Morals*. Translated by Thomas Kingsmill Abbott. Available on http://www.phatnav.com/books/pdf/fundamental_principles_of_the_metaphysic_of_morals.pdf (2005-01-14)

Kerr, N.L. (1995). Norms in social dilemmas. In: Schroeder, D. (Ed.). *Social Dilemmas: Perspectives on Individuals and Groups*. New York: Praeger.

Klein, D.B. and O'Flaherty, B (1993). A Game-theoretic rendering of promises and threats. *Journal of Economic Behavior and Organization* 21, 295-314.

Knack, S. and Keefer, P. (1997). Does social capital have an economic payoff? A cross-country investigation. *Quarterly Journal of Economics* 112, 1251-1288.

Konow, J. (2000). Fair shares: accountability and cognitive dissonance in allocation decisions. *American Economic Review* 90, 1072-1091.

Ledyard, J.O. (1995). Public goods: a survey of experimental research. In Kagel, J.H. and Roth, A.E. (Eds.) *Handbook of Experimental Economics*. Princeton: Princeton University Press.

Mas-Colell, A., Whinston, M.D. and Green, J.R. (1995). *Microeconomic Theory*. New York: Oxford University Press.

Newcomb, T. (1953). An approach to the study of communication acts. *Psychological Review* 60, 393-404.

Ostrom, E., Walker, J.M. and Gardner, R. (1992). Covenants with and without the sword: self-governance is possible. *American Political Science Review* 86, 404-417.

Radner, R. and Schotter, A. (1989). The sealed-bid mechanism: an experimental study. *Journal of Economic Theory* 48, 179-220.

Ross, L., Greene, D. and House, P. (1977). The false consensus effect: an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13, 279-301.

Sally, D. (1995). Conversation and cooperation in social dilemmas: a meta-analysis of experiments from 1958-1992. *Rationality and Society* 7, 58-92.

Sen, A.K., (1967). Isolation, assurance and the social rate of discount. *Quarterly Journal of Economics* 81, 112-124.

Valley, K., Thompson, L., Ginnons, R. and Bazerman, M.H. (2002). How communication improves efficiency in bargaining games. *Games and Economic Behavior* 38, 127-155.

TABLES

Table 1. Descriptive statistics of the sample.

Treatment:	Number of subjects	Gender		Subject affiliation	
		Men	Women	Stockholm School of Economics (SSE)	Other*
No communication (NC)	40	19	21	21	19
Pre-specified communication (P1)	80	45	35	41	39
Pre-specified communication (P2)	80	44	36	40	40
Free-form communication (FC)	78	43	35	38	40
Total	278	151	127	140	138

*Stockholm University and Royal Institute of Technology.

Table 2. Experimental results: deception (individual level data).

Treatment:	White lies		Deceptive lies		Total lies	
	Number	%	Number	%	Number	%
Pre-specified communication: weak promises (P1)	2/24	8.33	43/56	76.79	45/80	56.25
Pre-specified communication: strong promises (P2)	1/24	4.17	36/56	64.29	37/80	46.25
Free-form communication (FC)	0/23	0.00	22/55	40.00	22/78	28.21
p-value of difference:*						
P1 versus P2	0.551		0.147		0.206	
P1 versus FC	0.157		<0.001		<0.001	
P2 versus FC	0.322		0.010		0.019	

* A Pearson chi-square test is used.

Table 3. Results of logistic regression analysis on the probability of a deceptive lie. Subjects with a test score ≤ 70 .

Variable		All subjects			SSE subjects	Other subjects
		(1)	(2)	(3)	(4)	(5)
Constant	β	-1.033	-1.498	-1.423	-0.414	-1.562
	se	0.419	0.475	0.525	0.844	0.683
	p-value	0.014	0.002	0.007	0.624	0.022
Points on test	β	0.017	0.018	0.015	-0.0041	0.024
	se	0.0085	0.0086	0.0088	0.0085	0.012
	p-value	0.040	0.039	0.091	0.787	0.049
	marginal effect	0.0041	0.0042	0.0035	-0.0008	0.0058
Treatment: P1*	β	1.656	1.729	1.751	0.953	2.379
	se	0.427	0.439	0.446	0.661	0.627
	p-value	<0.001	<0.001	<0.001	0.149	<0.001
	marginal effect	0.391	0.407	0.410	0.177	0.591
Treatment: P2*	β	1.030	1.085	1.117	1.989	0.898
	se	0.399	0.409	0.420	0.777	0.567
	p-value	0.010	0.008	0.008	0.010	0.113
	marginal effect	0.243	0.254	0.262	0.370	0.223
Turn (1=seller in second turn)	β		0.843	0.812	1.625	0.309
	se		0.348	0.355	0.598	0.482
	p-value		0.015	0.022	0.007	0.522
	marginal effect		0.198	0.190	0.302	0.077
Woman	β			-0.500	-0.426	-0.511
	se			0.357	0.574	0.488
	p-value			0.161	0.458	0.295
	marginal effect			-0.117	-0.079	-0.127
SSE subject	β			0.706		
	se			0.361		
	p-value			0.051		
	marginal effect			0.165		
$\beta P1 = \beta P2$	p-value	0.144	0.137	0.153	0.182	0.014
Number of observations		167	167	167	75	92
Chi-square		20.724	26.775	33.040	14.580	24.007
p-value (chi-square)		<0.001	<0.001	<0.001	0.012	<0.001
Log-likelihood		-101.699	-98.673	-95.540	-38.093	-51.679
McFadden pseudo-R ²		0.092	0.119	0.147	0.161	0.188
% individual prediction		66.467	69.461	70.060	80.000	73.913

*Baseline category: free-form communication (FC)

Table 4. Beliefs about lies.

Treatment:	Beliefs about % white lies		Beliefs about % deceptive lies	
	Number of observations	Mean (std)	Number of observations	Mean (std)
Pre-specified communication: weak promises (P1)	80	12.56 (17.88)	80	79.28 (19.89)
Pre-specified communication: strong promises (P2)	80	11.88 (17.99)	80	70.34 (26.18)
Free-form communication (FC)	78	13.18 (23.68)	78	67.96 (28.74)
P-value of difference:				
P1 versus P2	0.799		0.019	
P1 versus FC	0.838		0.004	
P2 versus FC	0.681		0.575	

Table 5. Questionnaire answers about deception.

<i>Question</i>	<i>Strongly disagree (%)</i>	<i>Slightly disagree (%)</i>	<i>Slightly agree (%)</i>	<i>Strongly agree (%)</i>	<i>Σ (Agree)</i>
A white lie that makes someone else feel good is only positive.	3	20	60	17	77
I could lie to help someone else even if it hurts me.	4	28	57	11	68
I could lie to get revenge even if it hurts me.	34	48	16	3	19
I am more inclined to lie, the more I have to gain from the lie.	8	28	44	20	64
I am less inclined to lie, the more others have to loose from the lie.	8	10	44	38	82
I am less inclined to lie, the greater the risk of discovery.	2	5	26	67	93
You either lie or you don't, there are no degrees of lying.	30	49	17	4	22
I am less inclined to lie, the greater the lie has to be for me to be believed.	4	8	40	48	88
If I promise someone to tell the truth, that makes it very difficult for me to lie to that person.	1	8	41	49	90
There are degrees of promises, some formulations contain a greater promise than others.	2	8	51	39	90
	Never (%)	On singular occasion (%)	A few times (%)	Many times (%)	
Have you ever lied in an application – in writing or in an interview - for example when applying for work, membership, school or scholarships?	69	23	8	0	
Have you ever lied when selling something?	59	31	8	3	
Have you ever consciously reported false information in your income-tax return?	83	12	3	1	

Table 6. Experimental results: efficiency (pairwise data).

Treatment:	Contracts		Successful contracts		Payoff to pair in SEK
	Number	%	Number	%	Mean (std)
No communication (NC)	12/40	30.00	3/12	25.00	57.50 (26.67)
Pre-specified communication: weak promises (P1)	18/80	22.50	7/18	38.89	58.75 (28.43)
Pre-specified communication: strong promises (P2)	24/80	30.00	10/24	41.67	62.50 (33.28)
Free-form communication (FC)	29/78	37.18	16/29	55.17	70.51 (40.64)
P-value of difference:					
NC versus P1	0.371		0.429		0.725
NC versus P2	1.000		0.326		0.312
NC versus FC	0.438		0.078		0.036
P1 versus P2	0.281		0.856		0.375
P1 versus FC	0.044		0.278		0.031
P2 versus FC	0.339		0.328		0.143

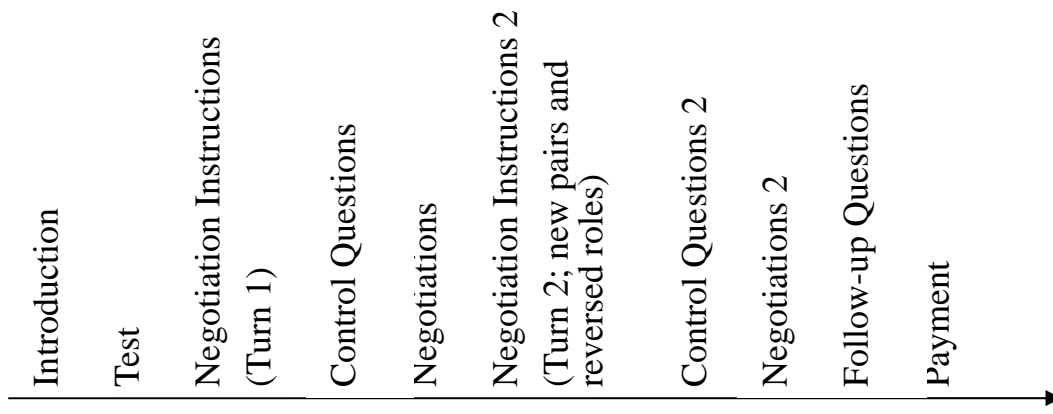


Figure 1: Experimental timeline.

	No contract	Contract, A's score > 70	Contract, A's score ≤ 70
A's payment	0	50	50
B's payment	50	100	0

Figure 2: Payment structure.




Figure 3: Scatterplot of the relationship between the points on the test and the points sent in the communication. Observations on the 45-degree line are sellers who stated their true points on the test. Sellers who choose to not send a message are shown on the horizontal line.




Figure 4: Points sent in the communication by low-talent sellers who lied. The low-talent sellers in the free-communication treatment that lied without communicating an exact score are denoted in the “FC>70” bar.

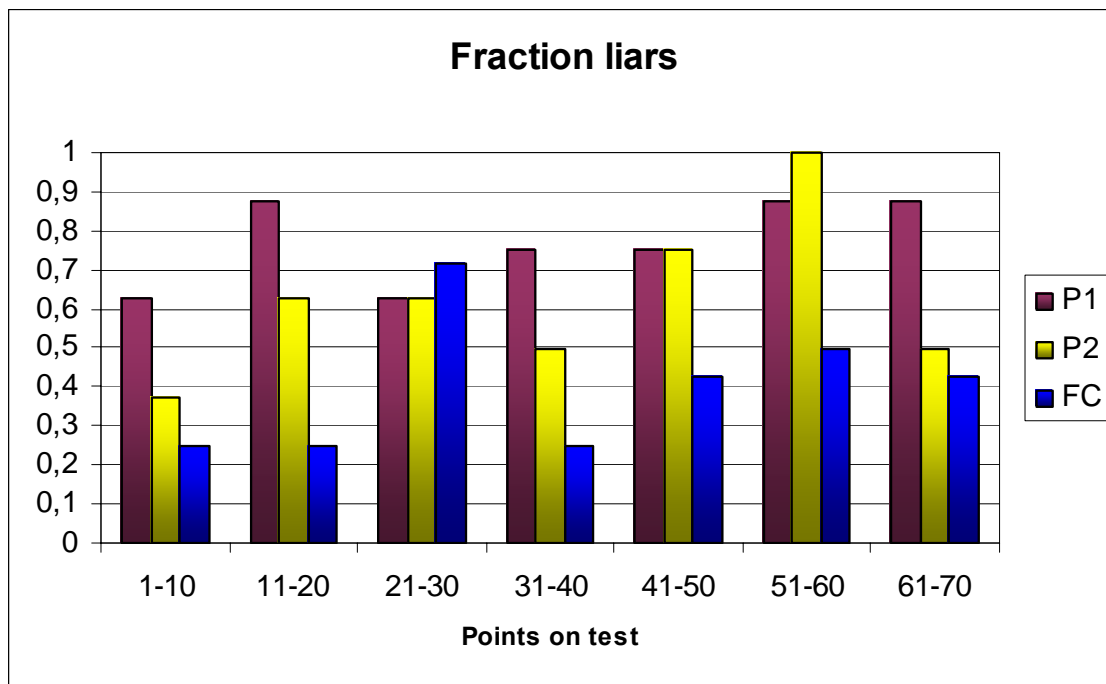


Figure 5: The fraction of low-talent sellers that lied in each decile of the test score distribution.




Figure 6: *The estimated probability of lying as a function of the test score in each of the three communication treatments. Based on the logistic regression equation in Table 3 (equation 3).*

Appendix

This is a translation of the instructions to the experiments, which were conducted in Swedish.

Introductory instructions

Hello and welcome.

You are about to participate in an economic experiment, in which you will take a test, make decisions in two stages and answer a few follow-up questions. The purpose of the experiment is to examine the decisions people make under certain conditions.

Your choices and your answers will only be used for scientific purposes. You have drawn a random identification number and we will not be able to link your person to your identification number, and thus not know which participant has made what specific choice.

The participants have been divided into two equally sized groups in two rooms which we call room 3 and 4. During the experiment, you will randomly be matched with participants in the other room, one person in round one and another in round two. You will not be informed with whom you have been matched, neither during nor after the experiment. After round two is completed, you will answer a few questions. When that is done, the experiment is over.

Depending on yours and others' actions, you will receive between 0 and 200 kronor in addition to the 50 kronor you have already received. Everyone will receive a sealed envelope at the end of the experiment, containing potential payment. The payment to participants in room 3 is sorted into envelopes by the experimenter in room 4 and then distributed by the experimenter in room 3, who does not know their contents, and vice versa.

During the entire experiment, you will receive instructions telling you what to do. If something is unclear, please feel free to raise your hand and we will come round to answer questions individually. During all parts of the experiment, all participants will be informed both of the conditions that apply to themselves, as well as to the person they are matched with. You will thus know what instructions and choices the person you are matched with faces.

There may be moments where you will have to sit and wait, and we ask you to be patient.

You may not talk to each other during the experiment and we ask you not to discuss the experiment with others afterwards.

Kind regards,

Erik & Tobias

Test instructions:

You will now answer a number of different types of questions. The test is conducted within a time limit, and your result may come to affect the payment you get when the experiment is over. The test is not an iq-test and thus does not reflect such characteristics.

The test is performed under great pressure against the clock and consists of 55 questions of varying nature. You have 10 minutes to complete the test from the moment you start with the first question. This will give you about 10-12 seconds per question. You may skip questions but you can not go back to answer them later.

The test will be corrected and everybody will be given a score between 0 and 100. The score is not equal to the number of correct answers, but a standardization that depends on how many questions each individual in your group has answered correctly. The person with the most correct answers in the room gets the highest score, second to most gets the second to highest score and so on. If two participants have an equal number of correct points, the person who has answered the most questions gets the higher score. 10 % of the participants in each group get a score between 1 and 10, 10 % between 11 and 20, 10 % between 21 and 30 and so on up to 100 points. Exactly what score you get in each interval is decided randomly, but a person with more correct answers will always get a higher score than a person with less. The test is thus corrected separately for each room and the scores are uniformly distributed. Your score will not be revealed to any other participant.

Please do not start until we give the word

[The test]

Negotiation – introduction:

You have received [...] points out of a possible 100.

You have now randomly been matched with a participant from the other room. You will be called [A/B] and the person with whom you have been matched [B/A]. You will not know with whom you were matched, neither now nor after the experiment.

You will now decide whether to enter into contract with each other or not. In order for the contract to hold, you must both accept. The payment will be the following:

You have role [A/B]

Your payment		B's payment
No contract:	0 kronor	50 kronor
Contract:	50 kronor	0 kronor if your score is 70 or less 100 kronor if your score is over 70

[or]

Your payment		A's payment
No contract:	50 kronor	0 kronor
Contract:	0 kronor if your score is 70 or less 100 kronor if your score is over 70	50 kronor

To ensure that you understand the setup we would like you to answer a few hypothetical questions:

If [I/A] have 30 points and we both choose to accept the contract, I get ___ and [A/B] gets ___.

If [I/A] have 45 points and I choose not to accept the contract, I get ___ and [A/B] gets ___.

If [I/A] have 75 points and we both choose to accept the contract, I get ___ and [A/B] gets ___.

If [I/A] have 95 points and I choose not to accept the contract, I get ___ and [A/B] gets ___.

The information that follows is identical to A and B:

[NC] (Nothing)

[P1] A will now be given the opportunity to send a message to B, where A can fill in a statement regarding his or her score. A can fill in a number between 0 and 100. The statement reads "My score is X", where A fills in a number for X. This is the only message A can send. If A chooses not to send any statement, B will receive the message "A has chosen not to send any message". A and B will then independently decide whether or not to accept the contract after B has received the message.

[P2] A will now be given the opportunity to send a message to B, where A can fill in a statement regarding his or her score. A can fill in a number between 0 and 100. The statement reads "I solemnly promise that my score is X", where A fills in a number for X. This is the only message A can send. If A chooses not to send any statement, B will receive the message "A has chosen not to send any message". A and B will then independently decide whether or not to accept the contract after B has received the message.

[FC] A will now be given the opportunity to send a message to B. The message may not contain anything that reveals A's identity. If A chooses not to send a message, B will receive: "A has chosen not to send any message". A and B will then independently decide whether or not to accept the contract after B has received the message.

(for P1) **Negotiation 1a:**

You have role [A/B]

[payment structure]

In order for the contract to hold, both must accept.

A total of 30 % of the participants in [your/A's] room have a score over 70.

[A] You may now, if you wish, send a message to B in which you fill in a statement about your score. This message will be communicated to B before s/he makes his or her decision.

Check the corresponding box to decide whether you want to enter into contract or not. If you choose to send a message enter the score below.

My score is ____

No message

[B] Message from A (wait till message is shown below).

Negotiation 1b:

Do you want to enter into contract?

Yes

No

Negotiation 2:

You will now be matched with a new person in the other room. You will be [Y/Z] and the person with whom you are randomly matched [Z/Y]. You will not find out each other's identities in this round either.

[The rest of negotiation 2 is the same as above, but with the roles reversed.]

Follow-up questions:

We would like you to briefly explain the reasoning behind your decisions.

In stage one.

In stage two.

Beliefs:

We also want you to guess how honest the participants have been.

How many percent of those with less than 70 points do you think lied about their score? ____%

What do you think all others have answered to the above question on average? ____%

How many percent of those with more than 70 points do you think lied about their score? ____%

What do you think all others have answered to the above question on average? ____%

If you answer these four questions correctly within five percentage points, you will get an additional 50 kronor.

Of all those who participate in the experiment, how many do you think lie about their score conditional on their true score and what score (on average) do you think they claim that they have?

True score	How big a fraction lies? (0-100 percent)	They claim (on average) that they have...(0-100 points)
1-10	%	points
11-20		
...		
91-100		

Questionnaire:

[The questionnaire]