DISCUSSION PAPER SERIES

DP12709

COMMUNICATION IN CONTEXT: INTERPRETING PROMISES IN AN EXPERIMENT ON COMPETITION AND TRUST

Alessandra Casella, Navin Kartik, Luis Sanchez and Sébastien Turban

PUBLIC ECONOMICS

ISSN 0265-8003

COMMUNICATION IN CONTEXT: INTERPRETING PROMISES IN AN EXPERIMENT ON COMPETITION AND TRUST

Alessandra Casella, Navin Kartik, Luis Sanchez and Sébastien Turban

Discussion Paper DP12709
Published 11 February 2018
Submitted 11 February 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

COMMUNICATION IN CONTEXT: INTERPRETING PROMISES IN AN EXPERIMENT ON COMPETITION AND TRUST

Abstract

How much do people lie, and how much do people trust communication when lying is possible? An important step towards answering these questions is understanding how communication is interpreted. This paper establishes in a canonical experiment that competition can alter the shared communication code: the commonly understood meaning of messages. We study a Sender-Receiver game in which the Sender dictates how to share $10 with the Receiver, if the Receiver participates. The Receiver has an outside option and decides whether to participate after receiving a non-binding offer from the Sender. Competition for play between Senders leads to higher offers but has no effect on actual transfers, expected transfers, or Receivers' willingness to play. The higher offers signal that sharing will be equitable without the expectation that they should be followed literally: under competition "6 is the new 5".

JEL Classification: C9, D9, D64, D83

Keywords: Bargaining, cheap talk, Lying, Dictator Game, trust game, Guilt-aversion

Alessandra Casella - ac186@columbia.edu
Columbia University and CEPR

Navin Kartik - nk2339@columbia.edu
Columbia University

Luis Sanchez - las2287@columbia.edu
Columbia University

Sébastien Turban - st2511@columbia.edu
Columbia University

Acknowledgements
We thank Brian Healy and Noah Naparst for their help; Pietro Ortoleva and participants at seminars and conferences at Caltech, Columbia University and U.C. San Diego for their comments; the National Science Foundation (SES-0617934) and the Sloan Foundation (BR-5064) for financial support; and CESS at NYU for access to its lab and subjects.

Communication in Context: Interpreting Promises in an Experiment on Competition and Trust*

Alessandra Casella† Navin Kartik‡ Luis Sanchez§
Sébastien Turban¶

Abstract

How much do people lie, and how much do people trust communication when lying is possible? An important step towards answering these questions is understanding how communication is interpreted. This paper establishes in a canonical experiment that competition can alter the shared communication code: the commonly understood meaning of messages. We study a Sender-Receiver game in which the Sender dictates how to share $10 with the Receiver, if the Receiver participates. The Receiver has an outside option and decides whether to participate after receiving a non-binding offer from the Sender. Competition for play between Senders leads to higher offers but has no effect on actual transfers, expected transfers, or Receivers' willingness to play. The higher offers signal that sharing will be equitable without the expectation that they should be followed literally: under competition "6 is the new 5".

JEL Classification: C9, D9, D64, D83

Keywords: Bargaining, Cheap Talk, Lying, Dictator Game, Trust Game, Guilt-aversion.

Introduction

During the 2016 presidential campaign, in an interview that became justly famous, Anthony Scaramucci argued that the media were misinterpreting Donald Trump' statements: "No, no, no, no, don't take him literally, take him symbolically", Scaramucci—later, if only briefly, White House Communication

*We thank Brian Healy and Noah Naparst for their help; Pietro Ortoleva and participants at seminars and conferences at Caltech, Columbia University and U.C. San Diego for their comments; the National Science Foundation (SES-0617934) and the Sloan Foundation (BR-5064) for financial support; and CESS at NYU for access to its lab and subjects.

†Columbia University, NBER and CEPR, ac186@columbia.edu

‡Columbia University, nk2339@columbia.edu

§Columbia University, las2287@columbia.edu

¶Columbia University, st2511@columbia.edu

director—told MSNBC (1). Whether appropriate or not, the comment highlights an important point: judgments about trustworthiness and truth-telling rest on one's belief about the code of communication, the mapping from words to their meaning in the context in which they are used.

The theoretical literature on cheap talk communication is careful to stress the importance of the code through which a message is expressed and interpreted (2-5). Experimental studies of communication and trust, on the other hand, measure trust as believing the letter of the message, and trustworthiness as following through with the letter of the message (6-11). This paper exploits an experimental design with a rich set of messages and choices to study the impact of a change in context on communication, trust, and trustworthiness. The data lead us to conclude that context determines how a message is interpreted—the action expected after the message is sent.

The change in context we study is the introduction of competition. In social environments, the decision to trust someone is typically accompanied by the question of whom to trust–trust is naturally paired with competition. Indeed, a sizable literature studies the impact of competition on trust (9, 12-14). We find that competition carries with it a change in language: not only are new messages used, but the interpretation of the messages changes. Truthfulness has been found to be context-dependent (15). Our thesis is that such dependence is commonly understood and shapes not only the reading of messages, but also their expected reading, and thus their contextual meaning. Framing has been shown to affect beliefs about beliefs (16); our evidence can be understood as showing that competition is a change in frame that alters the shared communication code.

1 Methods

1.1 Experimental design

The experiment is a variation on classic trust games (17, 18), designed as a one-shot dictator game with an outside option and preceded by one-sided, non-binding communication (7). We ran two treatments. In the one-sender treatment (1S), two partners, a Sender and a Receiver, are matched randomly and anonymously. The Receiver can choose whether to play the game or not. If the Receiver chooses not to play, both subjects receive \$2. If the Receiver chooses to play, the Sender is given \$10 to divide freely between herself and the Receiver in any integer split. Before the Receiver decides whether to play, the Sender sends the Receiver a non-binding message of the form: "If you decide to play with me, I will give you x dollars", where x can be any integer between 0 and 10.

In the two-sender treatment (2S)–the competition treatment–one Receiver and two Senders are matched randomly and anonymously. The Receiver receives messages (non-binding and private, of the same form as above) from both Senders, identified solely as Sender 1 and Sender 2. If the Receiver chooses not

to play the game, all three players receive $2; if the Receiver chooses to play, he must also indicate with which Sender. The selected Sender is then given $10 to share with the Receiver as desired; the Sender who is not selected receives $2.

1.2 Implementation

We ran the experiment on two different platforms: in the laboratory and online, via Amazon's Mechanical Turk. The laboratory experiment was programmed in ZTree (19) and took place at CESS, the Center for Experimental Social Science at NYU, with enrolled students recruited from the whole campus through the laboratory's web site.

Each experimental session consisted of a single treatment. The terminology and the sequence of moves in the lab followed the description above, with the following caveats. After sending her message, each Sender was asked how much she would transfer if the Receiver chose to play with her, without being informed of the Receiver's actual choice. (Asking Senders' choices via this "strategy method" allowed us to reduce confounding effects of learning and to increase the number of data points.) In addition to messages, participation decisions, and transfers, we also elicited beliefs. After having decided whether or not to play, the Receiver was asked what he expected the Sender(s) to transfer, given the(ir) message(s): these are the Receiver's first-order beliefs. After having indicated her transfer, each Sender was asked what she believed the Receiver expected her to transfer, given her message: these are the Sender's second-order beliefs. We incentivized the reporting of beliefs through a procedure used in closely related studies (7, 8, and 20): in the calculation of each round's payoff, a subject earned an extra $2 if the subject's stated beliefs were within $1 of the mean of the relevant variable in the session, given the specific message. The main purpose of eliciting beliefs was to directly study whether Senders and Receivers shared a common understanding of messages, as elaborated in our analysis below.

For both treatments, we ran eight rounds without feedback, with random assignment of roles and random rematching after each round. The only information revealed during the experiment consisted of the messages sent to the Receiver (which were revealed to the Receiver only). Each subject was paid his earnings over two random rounds, in addition to a $10 show-up fee. Sessions lasted between 30 and 40 minutes, with average earnings of $20 in the 1S treatment and $19.50 in the 2S treatment. A copy of the instructions is reproduced in the Supplementary Information.

The second implementation of the experiment was via an online survey run on Qualtrics, with participants recruited from Amazon's Mechanical Turk (MTurk) service. The experimental game, its terminology, and the structure of payoffs were identical to those we used in the laboratory, but monetary payoffs were reduced and the experiment was much shorter. Each subject played only one round, and the survey took about three and a half minutes, on average, for average earnings of 110 cents. The survey included comprehension quizzes, and subjects who failed to reply correctly were prevented from proceeding. We divided the survey in four waves over an interval of three days. The first two

waves were run simultaneously and collected data from Senders, one wave for the 1S treatment and one for 2S. The second two waves of the survey, again run simultaneously, collected Receivers' data for the two treatments. Each receiver was shown one (1S) or two (2S) Senders' messages, randomly drawn from the responses to the survey's first two waves. Payoffs were calculated ex post, after all surveys were received, by randomly matching Senders and Receivers, respecting the specific messages that had been drawn for each Receiver. Both the messages and the payoff-relevant matchings were generated by sampling with replacement, which allowed us to overcome the discrepancy in the exact number of respondents between the Sender(s) and the Receiver surveys. A copy of the 2S survey for Senders is reproduced in the Supplementary Information.

The design of the experiment is summarized in Table 1. We collected more data for 2S because in 2S only one of the three partners is a Receiver. All data are available from a link in the Supplementary Information.

Treatment	# sess.	# subs.	# S	# R	# rounds	# S obs.	# R obs.
Laboratory							
1S	6	78	39	39	8	312	312
2S	7	111	74	37	8	592	296
M. Turk							
1S	1	399	201	198	1	201	198
2S	1	595	395	200	1	395	200

Table 1. Experimental design: Number of sessions, subjects, Senders (S), Receivers (R) and rounds in the laboratory and in the MTurk survey.

We organize the data collected in the laboratory in two series: first-round data only, and data aggregated over all rounds. We thus describe all results in terms of three data series: the two series collected in the lab, and the data from the MTurk survey.

Each of the three series has strengths and weaknesses: first-round lab data are more directly comparable to data collected in the experiments that are closest to ours (7-9), but they are few in number–a particular problem because our subjects face more finely-grained choices–and may reflect some confusion with the setting and the game; all-round experimental data are more numerous but are not independent–the same subjects make multiple decisions–and may be affected by learning, even in the absence of feedback; MTurk data are numerous, independent, and free from learning, but much less controlled. Most of our results are consistent across the three series, giving us confidence in their robustness.

Statistical tests for lab all-round are complicated by the lack of independence. All standard errors and test significance levels that refer to lab all-round are calculated via bootstrapping, allowing for arbitrary correlation among choices made by the same individual. The methodology is described in detail in the Supplementary Information.

In both the lab and the MTurk survey, the messages sent by the Senders were identified in the instructions as "messages". In our description below, we will refer to them interchangeably as either "messages" or "offers".

2 Results

2.1 Competition induces higher messages ...

In all three data series, Senders' messages show three regularities (Figure 1). First, there is a clear spike in the distributions at 5, the modal offer in both treatments. Second, the spike is more pronounced in 1S. Third, the empirical distribution in 2S is shifted to the right (i.e., upward), relative to 1S: competition tends to increase the Senders' offers. A Kolmogorov-Smirnov test, adjusted for discreteness, strongly rejects the assumption of equal distributions across the two treatments ($p < 0.001$ in all three series) and fails to reject the one-sided alternative that the distribution in 2S first-order stochastically dominates the one in 1S ($p = 1.000$ for lab one-round and MTurk, $p = 0.902$ for lab all-rounds).

The shift upward of the offer distribution in 2S arises because the decline in the frequency of offers 5 in 2S is accompanied by an increase in offers 6 and 7. The magnitude of the shift is less pronounced in the MTurk data, but in all three data series the changes in the frequencies of the offers have the same sign and are statistically significant (Figure 1).

As a result of these changes in the offer distribution, average offers are higher in 2S than in 1S in all three data series. In all three the difference is statistically significant. Writing values in order for lab-first round, lab-all rounds, and MTurk, mean offers go from$\{4.8, 4.9, 4.75\}$ in 1S to $\{5.6, 5.9, 5.15\}$ in 2S. A t-test rejects the hypothesis that the difference is zero ($p < 0.001$ in all three data series).

2.2 ... but not higher transfers or higher expected transfers

The noticeable shift upwards in the distribution of offers induced by competition has no parallel in the (unconditional) distribution of transfers. Transfers move up only slightly, too slightly for statistical significance (Figure 2). The two most frequent transfers in 1S–5 and 0 in all three data series–remain the most frequent in 2S in lab-all rounds and MTurk. In all three series, however, their frequency declines, with intermediate transfers becoming more common.

At the aggregate level, mean transfers are slightly higher in 2S, but the difference is not statistically significant in any series (Figures 3A and 3B, left side panels). Moreover, while competition shifted the offers upward, Receivers did not expect higher transfers on average nor were they more willing to play: in all three data series, we find no significant differences in these variables between our two treatments (Figures 3A and 3B, mid and right side panel).

How does one reconcile the higher offers with the essentially unchanged transfers and participation? Apparently, higher offers were not believed, and thus did not result into higher expected transfers. But neither did they induce higher mistrust, and thus lower participation and lower expected transfers.

2.3 "6 is the new 5"

Across all three data series, offers of 6 in 1S are rare (Figure 1A). They are also associated with lower transfers than offers of 5 (Figure 4). Together, low frequency and low transfers suggest that an offer of 6 in 1S is a signal of low trustworthiness. In 2S, on the other hand, offers of 6 are effectively indistinguishable from offers of 5: in the lab, although not in MTurk, they are almost as frequent (Figure 1A), and in all three series the average transfers following either offer are equivalent (Figure 4). Offers of 7 are absent in 1S but appear in 2S, though they remain scarcer than either 5 or 6 (Figure 1A). In all three series, offers of 7 in 2S are associated with low transfers, transfers that are statistically indistinguishable from those following offers of 6 in 1S (Figure 4).

The data indicate different follow-through for offer 6 in the two treatments, with the caveat of few such data points in 1S. In all data series more than 50 percent of offers 6 are followed by a zero transfer in 1S, versus less than 20 percent in 2S; in all data series, following offer 6, the fraction of transfers 6 is at least double in 2S relative to 1S (Figure 5A).

Is the higher follow-through of offer 6 in 2S matched by different beliefs? In all three data series and for both Receivers' and Senders' beliefs, mean beliefs after offer 6 are higher in 2S than in 1S. However, all differences are small and none is statistically significant (Figure SI.1). It is plausible that mean beliefs do not differ much because subjects hold relatively diffuse beliefs: instead of believing that a Sender will only transfer either the offer itself or zero, Receivers may assign positive probability to a variety of transfers. We can translate this perspective into a quantitative measure.

Because the term credible is ambiguous when meaning can depend on context, we use instead the word persuasive. According to the data, Senders do not send and are not expected to send more than their offer. We thus define as non-persuasive an offer of x that conveys no more information than that: it induces the Receiver to assign positive probability to any transfer between 0 and x, and zero probability to transfers above x. A persuasive offer is instead one that generates first-order beliefs in the neighborhood of the offer. Both to allow for some dispersion in beliefs and because the belief elicitation procedure rewarded subjects for being within \$1 of the realized average value, we say that an offer of x is persuasive if it induces beliefs that the transfer will only be either x or $(x-1)$.

We evaluate whether offer x is persuasive by testing whether the realized fraction of beliefs at either x or $(x-1)$ is significantly higher than the expected fraction at such values if any belief between 0 and x is equally probable, i.e. under a uniform distribution from 0 to x (Figure 6). According to this measure, in 1S only offer 5 is persuasive to Receivers in all three data series, and in all

three data series is believed to be persuasive by Senders (in MTurk only, offer 6 is also believed to be persuasive by the Senders). In 2S, however, both offer 5 and offer 6 are persuasive to Receivers in all three data series and are expected to be so by Senders. No offer besides 5 or 6 satisfies the test for either order of beliefs in either 1S or 2S.

We interpret these results as supporting the hypothesis that an offer of 6 in 2S is interpreted differently from an offer of 6 in 1S. We also know that the offer is followed by different transfers in the two treatments (Figures 4 and 5). The phrase "6 is the new 5" conveys these two points. However, the statement is stronger: it says that competition makes an offer of 6 equivalent to an offer of 5 in the absence of competition. As we noted, the hypothesis of equal mean transfers following offer 5 in 1S and offer 6 in 2S cannot be rejected in any of the three data series (Figure 4). Neither can the hypothesis of equal mean beliefs for both Senders and Receivers, although mean beliefs do not vary enough across offers to attribute much significance to the lack of rejection (Figure SI.1).

A more powerful test of beliefs compares the full distributions of expected transfers (for Receivers), and beliefs about expected transfers (for Senders), conditional on offer 5 in 1S and offer 6 in 2S (Figure 7). For both Receivers' and Senders' beliefs, formal tests confirm that the distributions are not statistically different in the lab-first round and lab-all rounds series. In both lab series, beliefs peak at 5, in response to both offer 5 in 1S and offer 6 in 2S. In the MTurk data too, there is substantive shading down of beliefs following offer 6 in 2S: only 44 percent of Receivers' beliefs are at 6, following offer 6 in 2S, versus 73 percent at 5 following offer 5 in 1S; for Senders's beliefs the numbers are very similar: 45 percent at 6, after offer 6 in 2S, versus 76 percent at 5 following offer 5 in 1S. However, in the Mturk data beliefs peak at the literal offers, and equality of the distributions is rejected for both Receivers and Senders. Although "6 is the new 5" applies to the MTurk data too when interpreted in terms of mean values (mean transfers and mean expected transfers), the histograms of beliefs suggest a more literal interpretation of offers. It is interesting to recall that the shift in offers induced by competition is also less pronounced in MTurk. (Figure 1A). Lower sensitivity to subtle changes in stimuli in MTurk, relative to laboratory experiments, has been documented in other studies (22, 23). It could be that the change in context is less salient in the less interactive, less controlled, and faster web-based environment.

3 Discussion

Our best synthetic reading of the data is that messages are used and understood differently when senders compete for trust (in 2S), relative to the baseline treatment (1S). What drives the change in the communication code?

In the economic literature, the interaction between communication, trust, and trustworthiness has been studied through many different models ((24) analyzes 24 of them), but explanations for the effectiveness of non-binding promises come down to one of two mechanisms: guilt aversion (6, 7, 25-27), the aversion to

disappointing others' expectations, and lie aversion (28, 29), the aversion to not keeping one's word. In our context, guilt arises from disappointing Receivers' expectations of the transfer they will receive; a lie amounts to transferring less than one has offered. Both mechanisms are psychologically interesting and parsimonious, and distinguishing between them can be subtle (20, 27, 29-32).

When we want to understand changes in communication codes across contexts, however, guilt aversion faces a challenge. The problem is that, akin to the standard theory of cheap talk communication (33), nothing in the theory anchors messages—a Sender does not care intrinsically about messages, only about the expectations they induce in the Receiver. Thus, while guilt aversion can explain why non-binding promises are trustworthy, there is arbitrariness in which message is used to convey which expectation. This arbitrariness applies in any particular context, and a fortiori, across contexts.

Models of lie aversion, by contrast, define lies relative to messages' literal or conventional meaning, which is exogenous to any particular context. The approach restricts the choice of messages, while allowing for a distinction between literal meaning and contextual meaning—how the message is understood. As incentives vary across contexts, not only may different messages be used, but the contextual meaning attached to the same literal message can also change. Importantly, how it changes is pinned down by context.

Consider the following minimal conceptual framework, where we impose the discipline of equilibrium reasoning—behavior optimally responds to correct beliefs about the behavior of others. We describe briefly its logic here, relegating a more detailed analysis to the Supplementary Information.

Suppose a fraction of subjects, call it θ, have high costs of lying: they always transfer what they have offered. A fraction $1-\theta$ have low cost of lying and transfer 0 no matter their offer. In the 1S game, the model predicts concentration of messages at some unique offer x, supported by the belief that any Sender offering $x\prime \neq x$ must be untruthful; and a bimodal pattern of transfers, a fraction θ of the Senders transferring x and the remainder transferring 0. The experimental data in 1S are broadly consistent with this prediction, with offer 5 being focal. The data show high concentration of offers at 5 in all data sets (Figure 1), and, conditional on offer 5, a bimodal pattern of transfers at 5 and 0 (Figure 5B).

In 2S, the data show that offers are more dispersed (Figure 1). Our barebones model can rationalize such dispersion under competition. Suppose multiple offers, all belonging to a set X, are observed in equilibrium, again supported by the belief that any $x \notin X$ is sent only by untruthful Senders. For simplicity suppose $X = \{x_1, x_2\}$ with $2 < x_1 < x_2$. An offer cannot be accepted with positive probability unless it is sent by some truthful Senders (more precisely, by enough truthful Senders to induce an expected transfer superior to the outside option of 2). With $x_1 < x_2$, a truthful Sender sends x_2 only if it is accepted with higher probability than x_1. But if x_2 is accepted with higher probability, then all untruthful Senders send x_2. Thus x_2 is sent by a mixture of truthful and untruthful Senders, while x_1 is sent by truthful Senders only. We show in the Supplementary Information that with competition both the Receiver and

truthful Senders can be indifferent between two different offers, and thus both offers can coexist in equilibrium.

The interesting finding is not merely that dispersion of offers can be supported in equilibrium; rather it is that such dispersion can be supported *only with competition*. Consider a candidate equilibrium identical to the one just discussed, but in the 1S game. Since the lower offer, x_1, is offered by truthful Senders only, it is accepted by Receivers with probability 1. But then truthful Senders have no reason to offer x_2. In the 2S game, by contrast, the lower offer of x_1 may be matched with an offer of x_2 by the competing Sender, and thus may indeed be accepted with lower probability than x_2.

These predictions match well the concentration of offers we see in the data: the single spike of offers 5 in 1S, and the larger share of offers 6 in 2S. They are also consistent with the frequencies of acceptance of offers 5 and 6 in 2S: when the two offers compete, 6 is accepted more frequently than 5 in all three data series (the ratio of the frequencies of acceptance is 2 in lab-first round, 1.25 in lab-all rounds, and 2.33 in MTurk, but in part because the data points are few only the latter is significantly different from 1 ($p < 0.001$)).

The theory also rationalizes the change in the contextual meaning of message 6. In 1S, if message 5 is the equilibrium message, then message 6 is not used. It is associated with zero expected and realized transfers. In experimental data with some noise, message 6 may appear, but we would expect it to be used rarely and to be associated with low transfers. In 2S, both messages 5 and 6 can be sent in equilibrium. The expected and realized transfer following message 5 should be 5; following message 6, it cannot be smaller than 5, but it need not be higher. Even in noisy experimental data, message 6 in 2S would then be as persuasive as message 5, and be associated with similar mean transfers. Our experimental data fit these predictions well (Figures 1, 4, 5).

Not all predictions, however, are borne out: in 2S, especially but not exclusively in the lab data, transfers are more dispersed following both offer 5 and offer 6 (Figure 5) than the theory predicts. And how should we interpret the increase in offers 7 that we observe in 2S in the lab data (Figure 1)? If it is an off-equilibrium offer, as the lack of persuasiveness (Figure 6) and low conditional transfers (Figure 4) suggest, why is it so (relatively) common?

But the model with lying aversion sketched above is truly minimal. A more ambitious theory could add lying costs that depend on the magnitude of the lie (28, 34), possibly idiosyncratically, as well as heterogeneous innate altruism—the amount a Sender would transfer in the absence of communication. Senders would then want to convince the Receiver of their high altruism and/or high lying costs, but high offers would muddle higher altruism with lower lying costs. The multidimensional signaling game (35, 36) can rationalize greater dispersion in transfers after any offer, while still predicting the desired change in the contextual meaning of messages when competition increases.

We do not pursue such richer model here partly because it is less transparent, but mostly because our goal is not to calibrate a specific model to the data. Rather it is to stress, more broadly, the possible gap between literal and contextual meaning of messages, and to highlight how models of lying aversion

can generate predictions about such a gap and help us to understand it. Our experiment is useful because it anchors the change in meaning on a structural change—competition—that is unambiguous and important. Others have found that competition can alter what is judged equitable behavior (37). We find that it can alter how communication is interpreted. The main message of this paper is that analyses of competition, communication, and trust can be made richer by incorporating the endogeneity of the language code.

We close with a question. In models of lying aversion, lying costs are sustained if behavior deviates from the literal meaning of the message. But why, if the contextual meaning is understood by the Receiver? Philosophers have debated the moral standing of lies that do not intend to deceive (38, 39). We hope that future economic research too—theoretical and empirical—will pursue this direction.

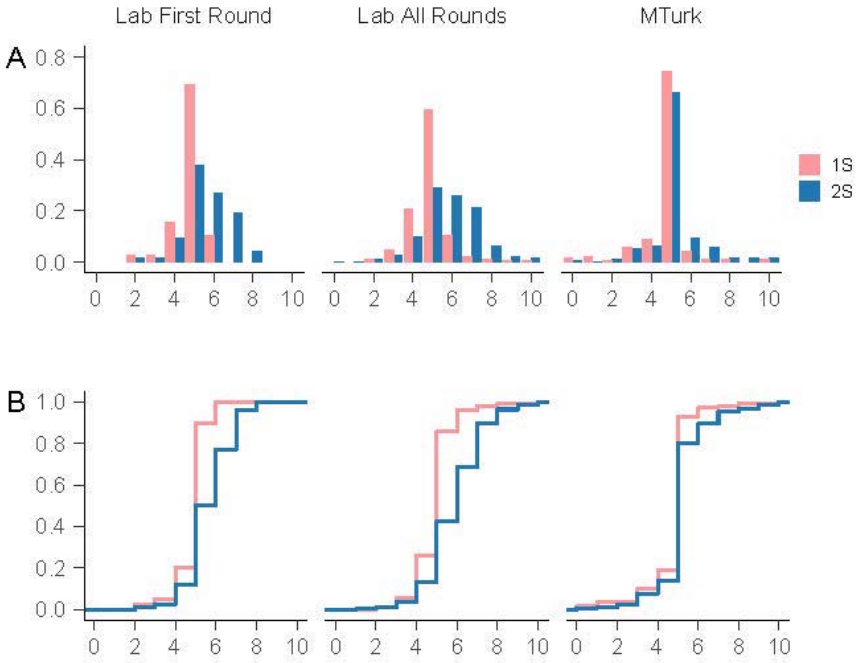

Figure 1. Offers in 1S (red) and 2S (blue): histograms of offers (Figure 1A); corresponding cumulative distribution functions (Figure 1B). In all three series the frequency of offer 5 decreases in 2S, relative to 1S, while the frequencies of offers 6 and 7 increase. Writing p-values for lab-first round, lab-all rounds, and MTurk, in order, the hypothesis of equal proportions of offer 5, against the one-sided alternative of a decline in 2S, is rejected with $p = 0.001$, $p < 0.001$, $p = 0.032$; the hypotheses of equal proportions of offer 6 and of offer 7, against the one-sided alternative of an increase in 2S, are rejected with $p = 0.034$, $p < 0.001$, $p = 0.002$, and $p = 0.005$, $p < 0.001$, $p = 0.005$, respectively (Chi-squared test of proportions for lab-first round and MTurk, bootstrapped simulations for lab-all rounds (see SI)).

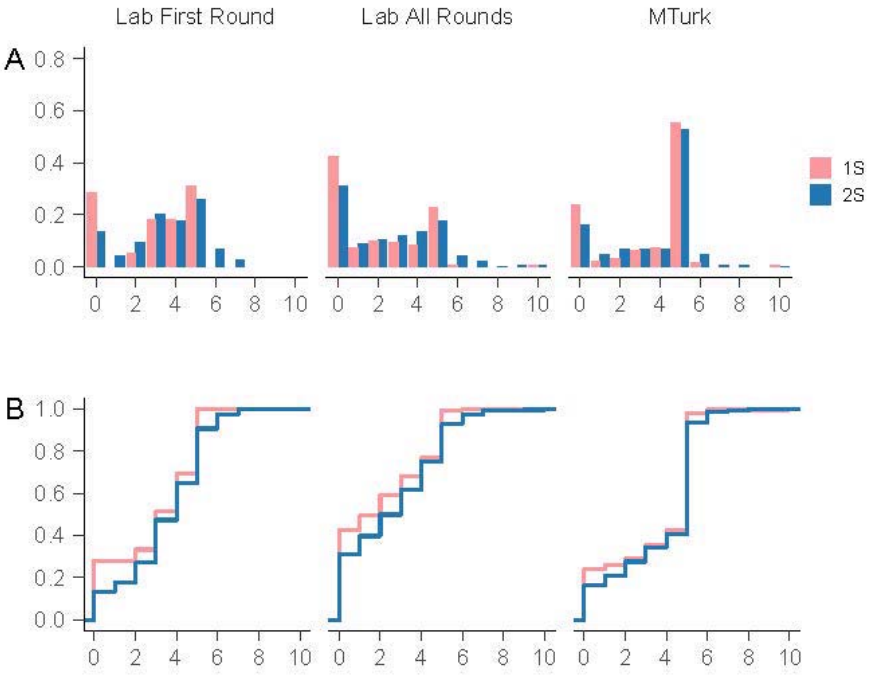

Figure 2. Transfers in 1S (red) and 2S (blue): histograms of transfers (Figure 2A); corresponding cumulative distribution functions (Figure 2B). A two-sided Kolmogorov-Smirnov test, adjusted for discreteness, cannot reject the hypothesis that the distributions are equal ($p = 0.283$ for lab-first round, $p = 0.164$ for lab-all round, and $p = 0.129$ for MTurk).

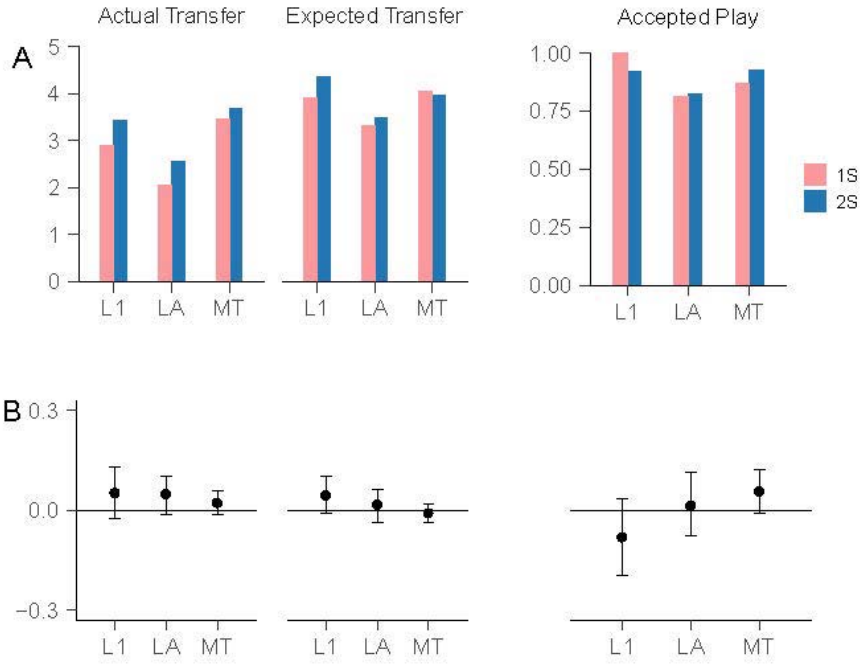

Figure 3. Transfers, expected transfers, and frequency of accepted play for 1S and 2S: point values (Figure 3A) and normalized percentage difference (2S-1S) with 95 percent confidence intervals (Figure 3B). The three data series are identified as L1 (lab-first round), LA (lab-all rounds), and MT (MTurk). The normalized difference is expressed as share of the maximal possible difference. In all cases, the 95 percent confidence interval includes zero. One-sided tests (against the alternative that the relevant variable is higher in 2S) lend marginally higher statistical significance to the differences between the two treatments. Reporting results in order (lab-first round, lab-all rounds, MTurk): $p = 0.094$, $p = 0.049$, $p = 0.114$ (for transfers); $p = 0.054$, $p = 0.215$, $p = 0.720$ (for expected transfers); $p = 0.890$, $p = 0.374$, $p = 0.046$ (for frequency of accepted play) (one-sided t-test of means and Chi-squared test of proportions for Lab-first round and MTurk, bootstrapped simulations for lab-all rounds (see SI).)

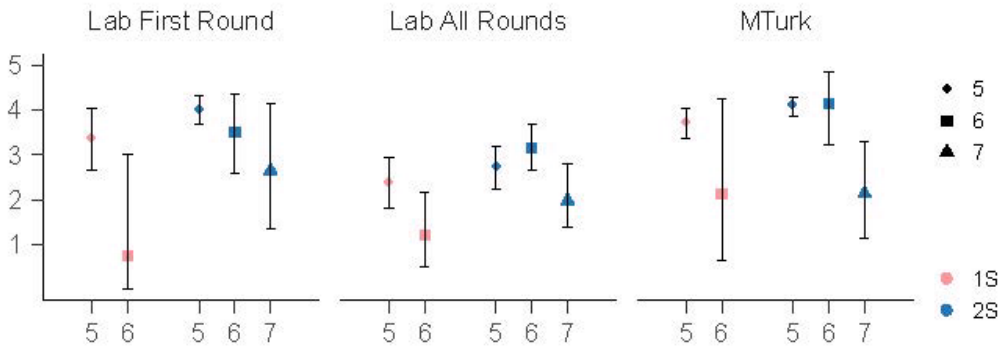

Figure 4. Mean transfers, conditional on offers 5, 6, and 7. Point values and 95 percent confidence intervals. All confidence intervals are calculated via bootstrapping to account for the non-negativity constraints and are centered on the empirical observation; in some cases the intervals are not symmetric because of the skewness of the data. We do not report transfers after offer 7 in 1S because the occurrences are too few to be meaningful: 0 such offers in lab first-round; 7 (2 percent) in lab-all rounds, and 2 (1 percent) in MTurk. The corresponding numbers in 2S are 7 (9 percent), 127 (21 percent), and 23 (6 percent). In 1S, offers of 6 are also rare (the numbers of offers 6 are reported in Figure 5). Yet, the hypothesis of equal transfers in 1S and 2S following offers 6 (against the one-side alternative of higher transfers in 2S) is strongly rejected in both lab data series, and marginally fails to be rejected in MTurk: $p = 0.011$, $p = 0.001$, $p = 0.053$ (for lab-first round, lab-all rounds and MTurk, in order; one-sided t-test for lab-first round and MTurk, bootstrapped simulations for lab-all rounds (see SI)).

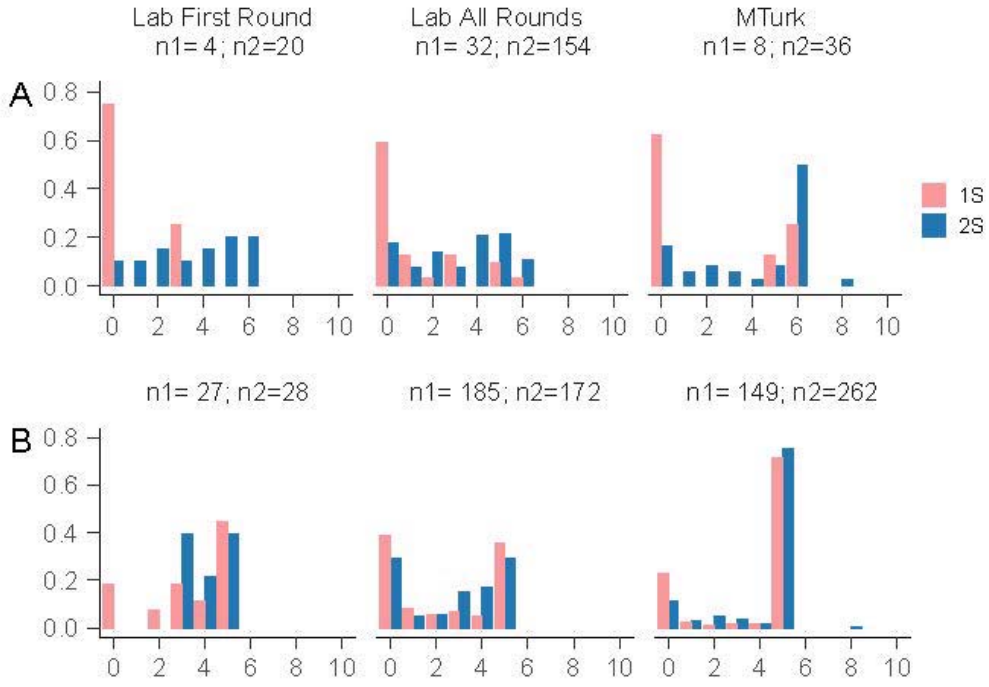

Figure 5. Transfers conditional on offer 6 (Figure 5A) and offer 5 (Figure 5B). Histograms of transfers in 1S (red) and 2S (blue) in each data series. Each panel also reports the number of data points: n1 refers to 1S, n2 to 2S. Following offer 6, the data in 1S are too few to test for equality of the distributions in lab-first round and MTurk; in lab-all rounds a Kolmogorov-Smirnov test corrected for discreteness strongly rejects the hypothesis of equal distribution in 1S and 2S ($p = 0.0032$). Following offer 5, the test cannot reject equality of the distributions in lab-first round and lab-all rounds ($p = 0.091$ and $p = 0.253$), but rejects it in MTurk ($p = 0.019$).

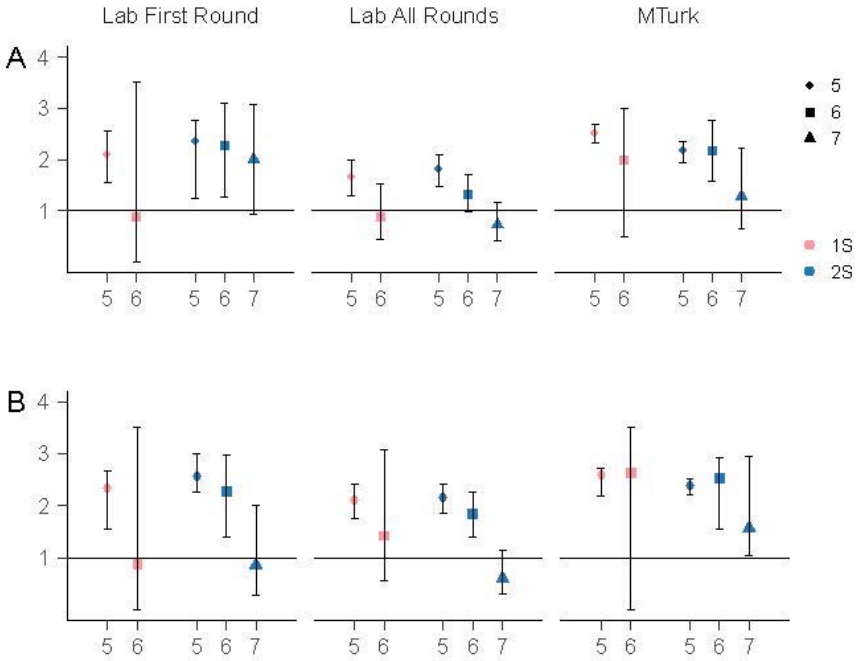

Figure 6. Receivers' (Figure 6A) and Senders' (Figure 6B) beliefs at x or $(x-1)$, following offer x. Each point is normalized by the expected mass of reported beliefs at x or $(x-1)$ if any belief between 0 and x is equally likely. An offer is persuasive if the ratio is higher than 1. For Receivers' beliefs (Figure 6A), in 1S the ratio is significantly higher than 1 only for offer 5 ($p < 0.001$ for all three data series); in 2S, for offer 5 ($p < 0.001$ for all data series) and for offer 6 ($p = 0.001$ for lab first-round, $p = 0.040$ for lab all-rounds, and $p < 0.001$ for MTurk), and for no other offer. For Senders' beliefs (Figure 6B), in 1S the ratio is significantly higher than 1 only for offer 5 ($p < 0.001$ for all data series) and for offer 6 in the MTurk data only ($p = 0.004$); in 2S, for both offer 5 and offer 6 ($p < 0.001$ in both cases for all data series) and for no other offer. See SI for details on the test.

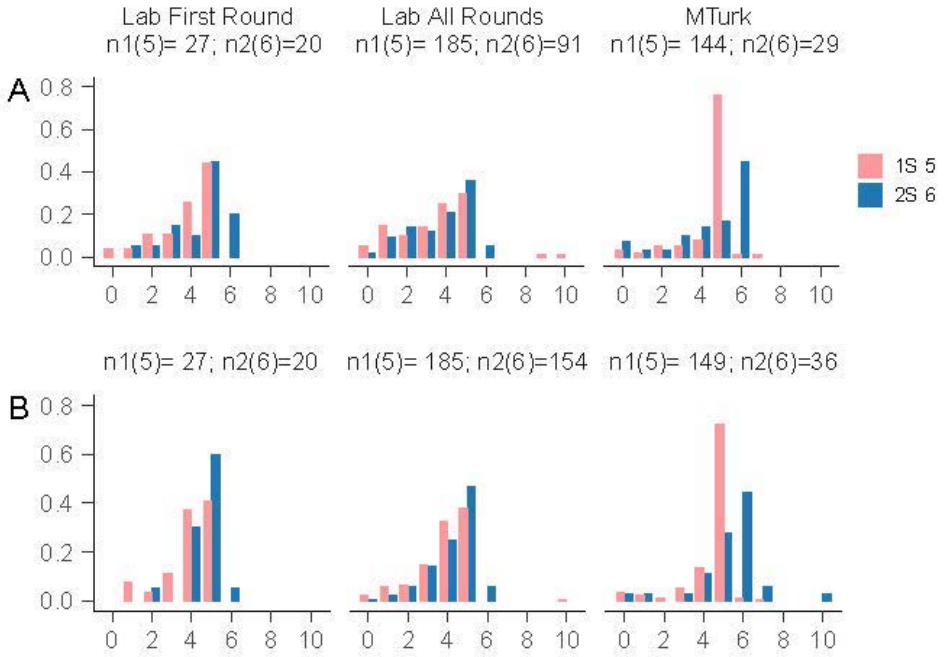

Figure 7. Distributions of beliefs conditional on offer 5 in 1S and offer 6 in 2S. Histograms of Receivers' first order beliefs (Figure 7A) and of Senders' second order beliefs (Figure 7B). A two-sided Kolmogorov-Smirnov test, adjusted for discreteness, cannot reject equality of the two distributions for lab-first round ($p = 0.269$ for Receivers, and $p = 0.1285$ for Senders), and lab-all rounds ($p = 0.631$ for Receivers, and $p = 0.168$ for Senders), but rejects it for MTurk ($p < 0.001$ for both Receivers and Senders).

References

1. MSNBC Live, December 20, 2016, Anthony Scaramucci's interview with Stephanie Ruhle. http://www.politico.com/story/2016/12/trump-symbolically-anthony-scaramucci-232848

2. Farrell, J. (1993): "Meaning and Credibility in Cheap-Talk Games", *Games and Economic Behavior* 5: 514-531.

3. Farrell, J. and M. Rabin (1996), "Cheap Talk", *Journal of Economic Perspectives* 10:103-118.

4. Demichelis S. and J. Weibull (2008), "Language, Meaning and Games: A Model of Communication, Coordination, and Evolution", *American Economic Review* 98: 1292-1311.

5. Sally, D. (2002), "What an Ugly Baby! Risk Dominance, Sympathy and the Coordination of Meaning", *Rationality and Society* 14:78-108

6. Gneezy, U. (2005), "Deception: The Role of Consequences," *American Economic Review* 95: 384-394.

7. Charness, G. and M. Dufwenberg (2006), "Promises and Partnership," *Econometrica* 74: 1579-1601.

8. Charness, G. and M. Dufwenberg (2010), "Bare Promises", *Economics Letters* 107: 281–283.

9. Goeree J. and J. Zhang (2014), "Communication and Competition", *Experimental Economics* 17: 421-438.

10. Ben-Ner, A., L Putterman and T. Ren, (2011), "Lavish returns on cheap talk: Two-way communication in trust games", *Journal of Socio-Economics* 40: 1-13 .

11. Corazzini, L., S. Kube, M. A. Maréchal and A. Nicolò, (2014), "Elections and Deception. An Experimental Study on the Behavioral Effects of Democracy", *American Journal of Political Science*, 58: 579-592.

12. Huck, s., G. Lünser and J. R. Tyran (2012), "Competition Fosters Trust", *Games and Economic Behavior* 76: 195-209.

13. Keck, S. and N. Karelaia (2012), "Does competition foster trust?", *Experimental Economics* 15: 204—228.

14. Fischbacher, U., C. Fong, and E. Fehr (2009), "Fairness, Errors and the Power of Competition", *Journal of Economic Behavior and Organization* 72: 527–545.

15. Gibson, R., C. Tanner, and A. Wagner (2013), "Preferences for Truthfulness: Heterogeneity among and within Individuals", *American Economic Review* 103: 532-48.

16. Dufwenberg, M., S Gächter, and H. Hennig-Schmidt (2011), "The framing of games and the psychology of play", *Games and Economic Behavior* 73: 459-478.

17. Berg, J., J. Dickhaut and K. McCabe (1995), "Trust, Reciprocity and Social History", *Games and Economic Behavior* 10: 122-142.

18. Dufwenberg, M. and U. Gneezy (2000), "Measuring Beliefs in an Experimental Lost Wallet Game", *Games and Economic Behavior* 30: 163-182.

19. Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental Economics* 10: 171-178.

20. Ellingsen, T., M. Johannesson, S. Tjøtta, and G. Torsvik (2010), "Testing guilt aversion", *Games and Economic Behavior* 68: 95-107.

21. Abadie, A. (2002), "Bootstrap tests for distributional treatment effects in instrumental variable models", *Journal of the American Statistical Association* 97: 284-292.

22. Bartneck, C., A. Duenser, E. Moltchanova and K. Zawieska (2015), "Comparing the Similarity of Responses Received from Studies in Amazon's Mechanical Turk to Studies Conducted Online and with Direct Recruitment", *PLoS ONE* 10(4): e0121595. doi:10.1371/journal.pone.0121595.

23. Horton, J., D. Rand and R. Zeckhauser (2011), "The Online Laboratory: Conducting Experiments in a Real Labor Market", *Experimental Economics*, 14: 399-425.

24. Abeler, J., D. Nosenzo, and C. Raymond (2016), "Preferences for Truth-Telling", IZA Discussion Paper No. 10188. https://ssrn.com/abstract=2840132.

25. Battigalli, P. and M. Dufwenberg (2007), "Guilt in Games", *American Economic Review* 97: 170-176.

26. Battigalli, P., G. Charness, and M Dufwenberg (2013), Deception: The Role of Guilt, *Journal of Economic Behavior and Organization* 93: 227–232.

27. Ederer, F. and A. Stremitzer, 2013, "Promises and Expectations," Cowles Foundation Discussion Papers 1931, Yale University, revised Mar 2016. http://cowles.yale.edu/sites/default/files/files/pub/d19/d1931.pdf.

28. Kartik, N. (2009), "Strategic Communication with Lying Costs," *Review of Economic Studies* 76: 1359-1395.

29. Hurkens, S and N. Kartik (2009),"Would I lie to you? On social preferences and lying aversion", *Experimental Economics* 12: 180–192.

30. Ismayilov, H. and J. Potters (2016), "Why do promises affect trustworthiness, or do they?", *Experimental Economics* 19: 382–393.

31. Kawagoe, T. and Y. Narita (2014), "Guilt aversion revisited: An experimental test of a new model", *Journal of Economic Behavior and Organization* 102:1-9.

32. Vanberg, C. (2008), "Why Do People Keep Their Promises? An Experimental Test of Two Explanations," *Econometrica* 76: 1467-1480.

33. Crawford, V. and J. Sobel (1982), "Strategic Information Transmission", *Econometrica* 50: 1431-1451.

34. Banks, J. (1990), "A Model of Electoral Competition with Incomplete Information", *Journal of Economic Theory* 50: 309-325.

35. Fischer, P. E. and R. E. Verrecchia (2000), "Reporting bias," *The Accounting Review* 75, 229–245.

36. Benabou, R. and J. Tirole (2006), "Incentives and Prosocial Behavior," *American Economic Review* 96, 1652–1678.

37. Schotter, A., A. Weiss and I. Zapater (1996), "Fariness and Survival in Ultimatum and Dictatorship Games", *Journal of Economic Behavior and Organization* 31: 37-56.

38. Sorensen, R. (2007), "Bald-Faced Lies! Lying without the Intent to Deceive", *Pacific Philosophical Quarterly* 88: 251-264.

39. Carson, T. (2006), "The Definition of Lying", *Noûs* 40: 284-306.

Appendix: Supporting Information

Statistical tests

We use bootstrapping techniques to correct for features of the data that may bias standard statistical tests.

Confidence intervals for means or proportions, for given treatment (1S or 2S).

Lab all-rounds data In the lab all-rounds data, multiple observations are generated by the same subject and thus cannot be assumed to be independent. Because we do not know the nature of the within-subject correlation, we calculated the statistical tests reported in the text via bootstrapping. The technique allows for arbitrary dependence across observations from a single subject, but assumes that such dependence is the same for all subjects and maintains independence across subjects. Relative to standard procedures, the only difference is that for each subject we draw the full vector of observations associated with the subject.

Call n the number of separate subjects for which we have observations related to some variable x. For each subject i, the number of observations is $k(i)$. We draw with replacement n subjects from the sample, and for each we consider the $k(i)$ observations of x associated with the subject. We then calculate the mean or proportion of the bootstrapped sample. We repeat the procedure 10,000 times and construct a distribution of means/proportions. The confidence intervals reported in the text are the intervals, centered on the empirical mean, that contain 95 percent of the means from the bootstrapped samples. It is well-known that non-parametric bootstrap confidence intervals need not by symmetric around the empirical mean.[1]

Lab first-round and MTurk data In the lab-first round and MTurk samples, we have a single observation for each subject and consider all observations independent. (Given the standardized instructions and the constancy of the physical lab and the subject pool, we discount the likelihood of session-specific effects in the lab-first round sample). However, the confidence intervals for means or proportions reported in the text were calculated via bootstrapping to account for non-negativity constraints. We followed the same procedure described for the lab-all rounds sample.

Differences in means or proportions across the two treatments.

Lab all-rounds data We apply our bootstrapping procedure to one-sided tests of differences in means or proportions across the two samples. We generate bootstrapped values for the difference under the null hypothesis that

[1] DiCiccio, T. and B. Efron (1996), "Bootstrap Confidence Intervals", *Statistical Science* 11: 189-228.

the two samples are drawn from the same population (and thus have equal means/proportions) and compare the distribution of the bootstrapped values to the difference we observe in the data. In practice, we combine the samples from the two treatments and assign to each subject a subject id. We then sample with replacement subject id's and assign each randomly to one of two groups, labelled group 1 or group 2. Call n_1 the size of the 1S sample (78), and n_2 the size of the 2S sample (111). We construct the two groups by assigning n_1 random draws to group 1 and n_2 to group 2, and we treat the samples in the two groups as if they corresponded to the two treatments. Because the assignment to the groups is random, however, subjects in the 1S treatment can be assigned to group 2, and vice-versa. When subject's id i is assigned to a group, all $k(i)$ observations of x associated with the subject are assigned to the group. We then compute the variable in question in the constructed groups, m_1 and m_2, and the difference $\Delta m = m_2 - m_1$. We repeat this procedure 10,000 times and generate a distribution of $\Delta(m)$ under the null of no difference in the population. The frequency with which the bootstrapped samples are more extreme than the observed difference in the data provides a one-sided test of difference. For example, calling $\Delta(\mu)$ the difference in means in our data, we test H_0: $\Delta(\mu) \leq 0$ v/s H_1: $\Delta(\mu) > 0$ by computing the p-value: $p = \{\#(\Delta(m) > \Delta(\mu))\}/10,000$.

Difference in distributions across the two treatments: Kolmogorov-Smirnov tests

Lab all-rounds data As above, call n_1 the number of data from the 1S treatment, with population cdf F and sample cdf F_{n_1}; call n_2 the number of data from 2S, with population cdf G and sample cdf G_{n_2}. We compute the relevant Kolmogorv-Smirnov test statistic in the data and compare it to the distribution of the test statistics obtained from the resampling procedure described earlier. For example, when testing H_0: $F = G$ versus H_1: $F \neq G$, we compute the two-sided KS test statistic $D = sup(t)|F_{n_1}(t) - G_{n_2}(t)|$ in the data. With each resampling σ, we obtain new empirical cdfs $F_{n_1'}(\sigma)$ and $G_{n_2'}(\sigma)$. Note that the subindices n_1' and n_2' will in general differ from n_1 and n_2: although the permutation procedure maintains the same number of subjects in the two groups, in the experiment the roles of Senders and Receivers are attributed randomly and thus are assumed by different subjects a different number of times. The result is that for given variable x we do not have the same number of observations per subject. Using the ks.test() function in R, we compute $D(\sigma) = D(F_{n_1'}(\sigma), G_{n_2'}(\sigma))$ on each resampling σ. To test the difference between distributions F_{n_1} and G_{n_2}, we compute the p-value $p = \{\#(D(\sigma) > D)\}/10,000$. We follow the same permutation procedure for the one-sided KS test (calculating the relevant one-sided KS test statistics).

Lab first-round and MTurk data Even when independence can be assumed, we calculate the KS tests via bootstrapping to account for discreteness of the samples. The adjustment for discreteness is based on (Abadie, 2002) and implemented with the ks.boot() function from the "Matching" R package. We

proceed as in the case of the lab-all round sample, with the only difference that since we have only one observation per subject, in each resample all subjects have the same number of observations (one), and the sample sizes are fixed at n_1 and n_2.

Senders' and Receivers' beliefs: comparison to uniform distribution

Given an offer x, the test compares the proportion of beliefs cumulated on x and $x-1$ (call it $\alpha(x, x-1)$) to the fraction that would be observed if beliefs were uniformly distributed over $[0, x]$, that is to the uniform null of $2/(x+1)$.

The test procedure is identical for all three data sets. For each offer x, we filter the data to include only subjects who received offer x (if testing Receivers' beliefs) or sent offer x (if testing Senders' beliefs). From the filtered data, we constract 10,000 resamples by bootstrapping (with replacement) the original sample. (In the case of lab-all rounds, a subject may corresponds to more than one data point if the subject has played multiple rounds in the same role at the same offer x.) From the resamples, we then construct the distribution of the proportion of beliefs cumulated on x and $x-1$. The p-values reported in the text test the null hypothesis H_0: $\alpha(x, x-1) = 2/(x+1)$ against the one-side alternative H_1: $\alpha(x, x-1) > 2/(x+1)$. They are calculated as $p = \{\#(\alpha(x, x-1) \leq 2/(x+1))\}/10,000$.

Beliefs Conditional on Offer 5 in 1S and Offers 5 and 6 in 2S

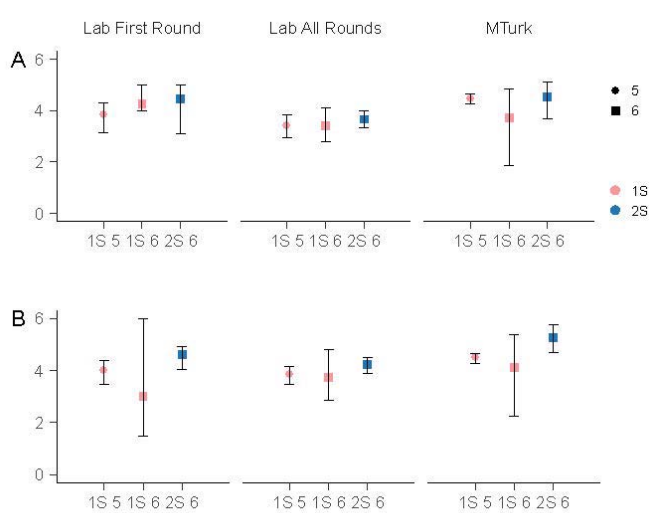

Figure SI.1. Receivers' (A) and Senders' (B) beliefs conditional on offer 5 in 1S and offers 5 and 6 in 2S. Point values and 95 percent confidence intervals. See the discussion of the statistical tests above.

A Simple Model with Binary Lying Costs.

A known fraction of subjects, θ, has high costs of lying; these subjects always transfer what they have offered. A fraction $1-\theta$ has low cost of lying; they transfer 0 no matter their offer. All subjects are risk-neutral.

The 1S game.

A single Sender (S) offers the Receiver (R) an amount x between 0 and 10. R decides whether to play or not. If R plays, S sends an amount z between 0 and 10 and keeps $10-z$ for herself. If R does not play, both R and S receive 2.

We restrict attention to non-trivial equilibria—those in which at least one offer is accepted with positive probability.

Result 1. *If $\theta \geq 1/4$, there is an equilibrium in which S makes the same offer $\tilde{x}$ regardless of her lying cost, with $\tilde{x} \in [2/\theta, 8]$, and R chooses to play. There is no equilibrium with two different offers, $x_2 > x_1 > 2$.*

The result follows from two simple observations. First, R will accept to play following an offer $\tilde{x}$ only if he expects a transfer of at least 2. Hence he plays only if $\tilde{x} \geq 2/\theta$. But a truthful S would never offer more than 8 (because that would leave her with less than 2); hence R would not accept offers higher than 8. Thus, if R accepts to play, $\tilde{x} \in [2/\theta, 8]$. Second, we can rule out the possibility of different equilibrium offers if both are strictly higher than 2. Suppose, towards contradiction, that two different offers, $x_1 > 2$ and $x_2 > x_1$ are both used and at least one induces R to play with positive probability. If R is more likely to play after x_2 than x_1, then all untruthful S would choose x_2. Thus x_1 would be offered only by truthful Senders, and because $x_1 > 2$ it would be accepted by R with probability 1. But then R cannot be more likely to play after x_2 –a contradiction. (The logic is identical if we suppose that R is more likely to play after x_1 than x_2). Hence, R must be equally likely to play after either x_1 or $x_2 > x_1$. But then a truthful S would only make the lower offer x_1, and hence R would not play after the higher offer x_2, again a contradiction.

An equilibrium with two (but not more) different offers $x_1 < x_2$ can be sustained only if $x_1 = 2$ (because in that case R can accept the offer with probability smaller than 1 even though it is made only by a truthful S). In the data, the frequency of offer 2 in in 1S is negligible: there is exactly one instance in lab-first round (out of 39) and in MTurk (out of 201), and three instances in lab-all rounds (out of 312).

The 2S game

Two Senders (S) independently and privately each offer the Receiver (R) an amount x between 0 and 10. R decides whether to play or not, and if so, with

which S. If R plays, the chosen S sends an amount z between 0 and 10 and keeps $10-z$ for herself; the rejected S receives 2. If R does not play, all players, R and both S's, receive 2.

Result 2. *If $\theta > 5/6$, there is an equilibrium in which both $x_1 = 5$ and $x_2 = 6$ are offered and accepted with positive probability. R always chooses to play. If faced with two equal offers, R accepts either one of them with equal probability; if faced with two different offers, R accepts x_1 with probability $p \in (1/4, 1/3)$, and accepts x_2 otherwise. An untruthful S offers x_2, a truthful S randomizes between x_1 and x_2.*

The specific values for x_1 and x_2, 5 and 6, stated in Result 2 are motivated by the experimental data. The result is more general (different values of x_1 and x_2 can be supported both for a specific value of θ and across different values of θ) and the number of possible equilibrium offers need not be restricted to two. Here we describe in more detail the logic underlying the result in the case of two equilibrium offers.

If both offers are accepted with positive probability, it must be that both are made by a truthful S with positive probability. Call $\sigma > 0$ the probability that a truthful S offers x_2 (and $1-\sigma$ the probability that she offers x_1). With $x_1 < x_2$, a truthful S can be indifferent between both offers only if x_2 is accepted by R with higher probability. Hence the probability that R accepts the lower offer, p, must satisfy $p < 1/2$. But if $p < 1/2$, an untruthful S will only offer x_2. Hence x_1 is only offered by a truthful S, and thus the expected transfer following offer x_1 is simply x_1. It follows that $x_1 \geq 2$. We have an equilibrium if and only if R is indifferent between accepting x_1 or x_2, and a truthful S is indifferent between offering x_1 or x_2.[2] The indifference condition for R corresponds to $x_1 = E(\text{transfer } |x_2) = x_2 Prob(S \text{ is truthful}|x_2)$, or, using Bayes rule:

$$x_1 = x_2 \frac{\theta\sigma}{(1-\theta+\sigma\theta)} \tag{1}$$

If we denote by $U_\theta^S(x)$ the expected payoff to a truthful S of making offer x, the indifference condition for a truthful S corresponds to $U_\theta^S(x_1) = U_\theta^S(x_2)$, where:

$$U_\theta^S(x_1) = \\ (10-x_1)[\theta(1-\sigma)/2 + p(1-\theta+\sigma\theta)] + 2[\theta(1-\sigma)/2 + (1-p)(1-\theta+\sigma\theta)] \tag{2}$$

$$U_\theta^S(x_2) = \\ (10-x_2)[(1-p)\theta(1-\sigma) + (1-\theta+\sigma\theta)/2] + 2[p\theta(1-\sigma) + (1-\theta+\sigma\theta)/2] \tag{3}$$

[2] In principle, $p = 0$ is possible: R could strictly prefer x_2 when faced with x_2 and x_1. Thus (1) only needs to hold as a weak inequality: $x_1 \leq x_2\theta\sigma/(1-\theta+\sigma\theta)$. With $x_1 = 5$ and $x_2 = 6$, however, no equilibrium exists with $p = 0$.

The first square bracket in each equation corresponds to the probability that the offer is accepted when the competing S offers x_1, the first term in the bracket, or offers x_2, the second term (with each event weighted by its probability). The second square bracket corresponds to the probability that the offer is rejected, again depending on the competing offer. An equilibrium is a vector $\{x_1, x_2, p, \sigma\}$ such that: $x_2 > x_1 \geq 2$; $\sigma \in (0, 1)$; $p \in [0, 1/2)$; (1) is satisfied; and (2) equals (3). Some algebra shows that, as claimed in Result 2, there is an equilibrium with $x_1 = 5$ and $x_2 = 6$ when $\theta \in (5/6, 1)$. In such an equilibrium, $p \in (1/4, 1/3)$ and $\sigma \in (0, 1)$, declining from 1 to 0 as θ increases and p decreases.

Data Sets

All data are available at http://columbia.edu/~ac186/datac4t/DATA.ZIP. Dataset S1 collects the lab data; dataset S2 the MTurk data; the readme file reports the series' definitions.

Online Appendix: Experimental instructions.

Laboratory experiment

We report below the lab instructions for the 2S treatment.

Thank you for agreeing to participate in this decision making experiment. During the experiment we require your complete, undistracted attention, and ask that you follow instructions carefully. You may not open other applications on your computer, chat with other students, or engage in other distracting activities, such as using your phone, reading books, etc. Please turn off your cell phone.

Before we begin, please read and sign the consent form, which is located at your terminal.

You will be paid for your participation in cash at the end of the experiment. Different participants may earn different amounts. What you earn depends partly on your decisions and partly on the decisions of others.

The entire experiment, including all interaction between participants, will take place through computer terminals. It is important that you not talk or communicate with others during the experiment, except as described below.

We will start with a brief instruction period. If you have any questions during the instruction period, raise your hand and your question will be answered out loud so everyone can hear. If you have any questions after the experiment has begun, raise your hand and an experimenter will come to assist you.

The experiment you are participating in will have 8 rounds.

In each round you will be randomly matched with two other participants to form a group of three. Each group will follow exactly the same rules, and what happens in one group has no effect on the other groups. You will not know, either during the experiment or afterwards, whom you were matched with in any round: all interactions are completely anonymous. In each group of

three participants, two will be assigned the role of player A and one the role of player B. In these instructions and when useful during the experiment the two A players will be called A1 and A2, but there is no meaning to the number 1 or 2; the only role of the label is to distinguish the two players. The computer will assign the roles randomly in each round, and your screen will inform you whether your role is A or B.

In each round, the two A players are each given10 dollars that he/she may split with player B. These are real dollars. Each A Player will send a message to player B indicating how he/she intends to divide the 10 dollars with B. Player B will then decide whether to Play with one of the two A players or not to Play. If player B chooses not to Play, all three players in the group receive 2 dollars each. If player B chooses to Play, for example with A2, then Player A2 splits the 10 dollars as he/she sees fit, and the A player who has not been chosen, A1 in this example, earns 2 dollars. If player B chooses to play, he/she can only play with one of the two A players in his/her group.

The messages that Players A1 and A2 send before decisions are made will be of the form: "If you decide to play with me, I will give you x dollars". Each A player will specify x, the number of dollars, by typing a number between 0 and 10 in the appropriate box on the screen.

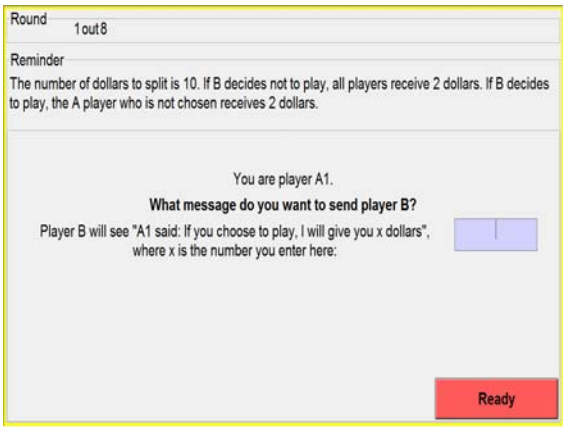

Sender's screenshot: offer.

In addition to this message, each A Player will also be asked how many dollars to give B in case B decides to Play. This decision is not constrained by the message that the A player has sent, and there is no punishment or reward for following through on the message. A1 and A2 can choose to give B any number of dollars between 0 and 10, if B decides to play. The screen of the A subjects will look like this:

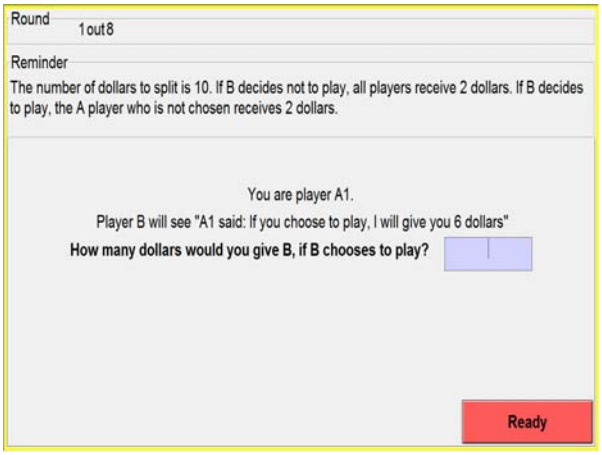

Sender's transfer.

Player B will receive the messages sent by the two A players and then decide whether to play at all, and if so with which one of the two A players. B will indicate this decision by clicking an appropriate button on the screen. The screen for Player B will look like this:

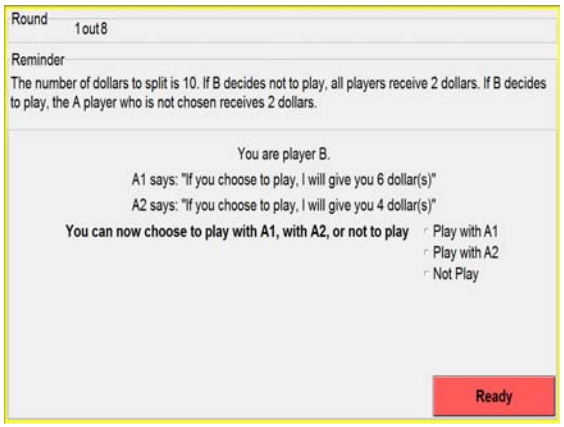

Receiver's screenshot: response to the Sender's offer.

You will not be told the decisions of the other members of your group until the end of the experiment. But keep in mind the rules of the game. If B decides not to Play, then all three members of the group earn 2 dollars for that round. If B decides to Play with one of the two A players, for example A2, than B earns the number of dollars given by A2; A2 keeps the remainder out of the 10 dollars that he/she has to split, and A1, who has not been chosen, earns 2 dollars.

After these decisions have been made, you will be asked a new question that may affect your earnings. You will see a new screen. If you are Player B, the screen will ask you how much you expect A1 and A2 to actually give you if you play the game with either of them. If you are an A Player, the screen will ask you what you expect B to guess you will give him/her if B played the game with you. You will indicate your choice by typing a number between 0 and 10 in the appropriate box on the screen. The screens will look like this: If you are B:

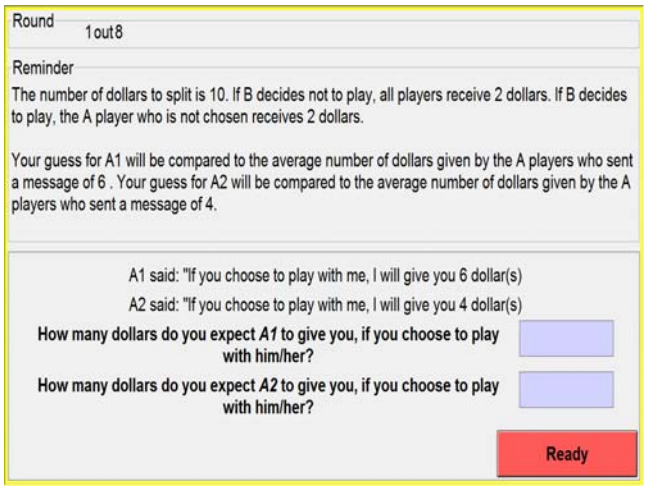

Elicitation of Receiver's beliefs.

If you are A:

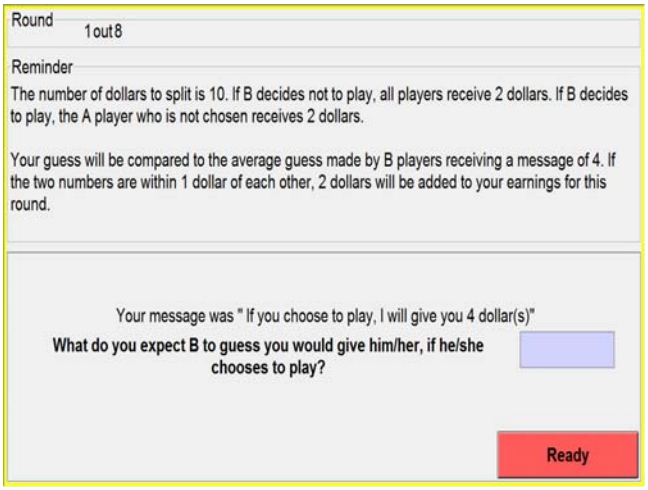

Elicitation of Sender's beliefs.

At the end of the experiment, your answer will contribute to your earnings from that round. If you are player B, you will earn 1 additional dollar from each of your two answers if the answer is within 1 dollar of being correct on average. That is, if your prediction is correct (+/- 1 dollar) when compared

to the average, over the full experiment, of what participants in role A have indicated they would give, following the specific messages that you have received in that round. If you are an A player, you will earn 2 dollars if your prediction is correct (+/- 1 dollar) when compared to the average, over the full experiment, of what participants in role B expect their partners to give following the specific message that you have sent in that round.

For example: Suppose that in the current round, Player A1's message in your group is that s/he would give B 8 dollars. If you are B, and you reply that you expect that A1 would in fact give you 6, your answer will be compared to the average of what A's have indicated they would give during the experiment after sending a message of 8 dollars. If, on average, A's who sent a message of 8 dollars have said they would give 5 dollars, then your guess of 6 is considered correct and you earn 1 additional dollar (because 6, your guess, is within 1 dollar of 5, what A players with a message of 8 have indicated on average). You will be asked the same question about player A2. If Player A2 has sent a message of 2, and you answer that you expect A2 to send 4 dollars, then you earn another additional dollar if on average A players with a message of 2 have indicated they would give either 3, 4 or 5 dollars. If your answer is more than 1 dollar away, then you earn 0 from that answer.

Exactly the same criterion applies to the A players, with the only difference that A players will answer a single question. If your role is A and your message said that you would give B 8 dollars, you earn 2 additional dollars if your guess of what B expects you to give is within 1 dollar of the average answer from B's who have received a message of 8; and 0 if you are more than 1 dollar away.

A round is concluded when all groups have answered all their questions. We will then move to a new round. New groups of three players will be formed randomly by the computer, and because the groups are formed randomly, your

partners in the group will most likely be different from round to round. Remember that all interactions are anonymous and no-one will ever know, during the experiment or in the future, the identity of the partners they are matched with in a group. Two of the members of your group will be randomly assigned to role A and one to role B, and the game will then proceed as before.

The experiment will continue in this fashion for 8 rounds. At the end of the experiment, 2 of your 8 rounds will be randomly selected by the computer program and you will be paid your total earnings for those 2 rounds, in addition to the $10 show-up fee. Remember that your earnings per round include both any earnings from decisions made about the game and any earnings from guesses. You will be paid in private and have no obligation to tell anyone how much you earned.

Are there any questions?

We will now proceed to a simple quiz that will allow you to verify that all instructions are clear. Please answer the questions on your computer; you will receive immediate feedback about your answers.

QUIZ

1. You are an A player and your message to B is "If you accept to play, I will send 4 dollars". Can you then indicate that you would send 6 dollars, should B choose to play? {Yes, No}. 2 dollars? {Yes, No}

2. A1's message to B is "If you accept to play, I will send 4 dollars". A1 then indicates that he/she would send 6 dollars. A2's message to B is "If you accept to play, I will send 6 dollars". A2 then indicates that he/she would send 5 dollars. B refuses to play. How many dollars does A1 earn? And A2? How many dollars does B earn? Suppose instead that B accepts to play with A1. How many dollars does A1 earn? And A2? How many dollars does B earn?

3. A1's message to B is: "If you accept play I will transfer 3 dollars". A1

indicates that if B agrees to play he will in fact transfer 6 dollars. B's best guess of A1's intended transfer is 5. Does B earn one additional dollar for being correct? Yes No

AFTER THE QUIZ: Are there any questions now?

If you have any questions from now on, raise your hand and an experimenter will come to assist you. We will now begin the experiment.

qualtrics FREE ACCOUNT

Instructions Mechanical Turk

In this task, you are matched with two other people. Keep in mind that these are real people, working on the task just like you.

The task is a survey. It should take you about 5 minutes. If you have not completed it within 10 minutes you forfeit your payment.

You will earn 50 cents for completing the survey plus a bonus that depends on what you and your partners respond.

We will now describe how the bonus works.

Here is how the bonus works.

Two people will be called Senders and the third person will be called Receiver.

If the Receiver chooses to work with one of the two Senders, that Sender is given 100 cents and can share them with the Receiver if he/she wants to. The other Sender receives 20 cents.

If the Receiver chooses not to work with either Sender, then all three people receive 20 cents.

The Receiver can work at most with a single Sender.

Before the Receiver makes his/her choice, each Sender sends the Receiver a message suggesting a division of the 100 cents, if chosen.

The message is not binding: a Sender, if chosen, does not need to divide the 100 cents with the Receiver in the way described in the message.

The Receiver knows that the message is not binding.

Summarizing (all numbers below are examples only)

1. First, both Senders send a message to the Receiver: “If you choose me, I will give you ...” For example:

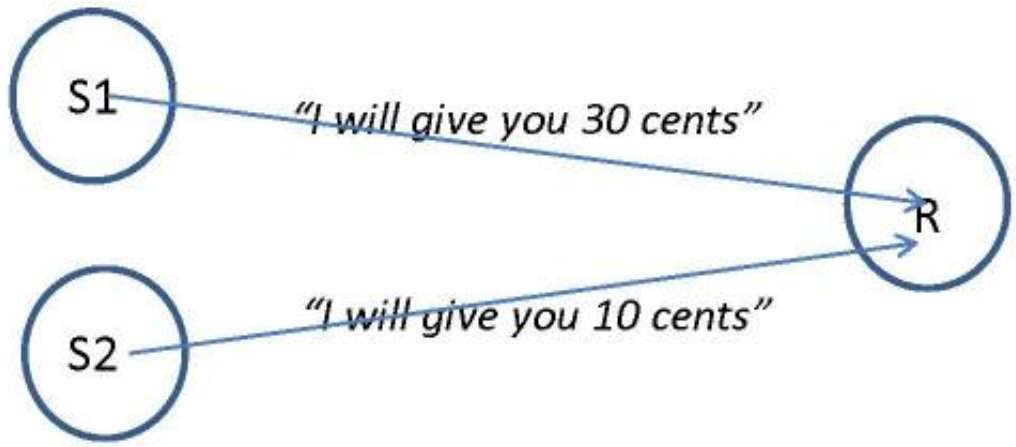

2. Then the Receiver chooses.

If R chooses S1, for example, then S1 decides how much to transfer to R, out of the 100 cents. The transfer can be any amount between 0 and 100. For example:

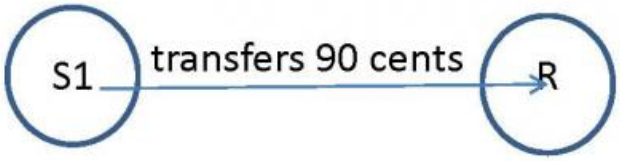

R receives a bonus of 90 cents
S1 receives a bonus of 100-90=10 cents
S2 receives a bonus of 20 cents.

If R rejects both Senders, then all three people receive a bonus of 20 cents.

ComprehensionCheck1

We want to verify that you are a real person and have read the instructions. Please answer the questions below. Feel free to check the summary instructions at the bottom of this page if useful.

If a Sender is chosen, how much (in cents) is available for him/her to share with the Receiver?

Suppose a Sender sends the message “I will give you 60 cents”. Does that Sender then need to give 60 cents?

Yes

No

Can the Receiver choose more than one Sender?

Yes

No

Show/Hide Instructions

ComprehensionCheck2

Suppose the Receiver rejects both Senders.

How many cents is the Receiver's bonus?

How many cents is the bonus for each Sender?

Suppose the Receiver chooses Sender 1 and then Sender 1 transfers 60 cents.

How many cents is the Receiver's bonus?

How many cents is the bonus for Sender 1?

And for Sender 2?

Show/Hide Instructions

Tasks

You have been assigned the role of Receiver.

The message from Sender 1 is: “If you choose me, I will give you **${e://Field/randPromiseValueA}** cents.”

The message from Sender 2 is: “If you choose me, I will give you **${e://Field/randPromiseValueB}** cents.”

You can now choose:

Sender 1

Sender 2

Reject Both Senders

Show/Hide Instructions

Extra bonus

In the previous screen, you chose ${q://QID39/ChoiceGroup/SelectedChoices}

How much do you think Sender 1 would transfer if you chose him/her (in cents)?

0 10 20 30 40 50 60 70 80 90 100

How much do you think Sender 2 would transfer if you chose him/her (in cents)?

0 10 20 30 40 50 60 70 80 90 100

Reminder: Sender 1 said he/she would give
you **${e://Field/randPromiseValueA}** cents.
Sender 2 said he/she would give you **${e://Field/randPromiseValueB}** cents.

We are asking all Senders how much they want to send, before telling them whether they were chosen. If your answers are within 10 cents of what Senders with those messages replied on average, you will earn an additional 10 cents for each answer.

Show/Hide Instructions

Verification

Thank you for participating in this survey.

We will match your answers to the Senders' answers and calculate your bonus. You will receive it shortly.

Type the text

Privacy & Terms