

IZA DP No. 9704

Gossip and the Efficiency of Interactions

Dietmar Fehr
Matthias Sutter

February 2016

Gossip and the Efficiency of Interactions

Dietmar Fehr

Wissenschaftszentrum Berlin

Matthias Sutter

*University of Cologne, University of Innsbruck
and IZA*

Discussion Paper No. 9704
February 2016

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Gossip and the Efficiency of Interactions*

Human communication in organizations often involves a large amount of gossiping about others. Here we study in an experiment whether gossip affects the efficiency of human interactions. We let subjects play a trust game. Third parties observe a trustee's behavior and can gossip about it by sending a message to the trustor with whom the observed trustee will be paired (for the first time) in the next round. While messages are non-verifiable and sometimes also incorrect, the possibility of gossip is highly efficiency-increasing compared to a situation without any gossip. In two further control treatments, we show that the mere fact of being observed by third parties cannot explain the efficiency-increasing effect of gossip, and that noisy gossip (where information transmission from third parties to trustors can fail) still increases efficiency, but less so than if information transmission is undisturbed.

JEL Classification: C72, C92

Keywords: gossip, communication, trust game, efficiency, experiment

Corresponding author:

Matthias Sutter
Department of Economics
University of Cologne
Albertus-Magnus-Platz
50923 Köln
Germany
E-mail: matthias.sutter@wiso.uni-koeln.de

* We thank seminar participants at the ESA-meeting in Santa Cruz, UIBK-MPI-Workshop in Innsbruck, HeiKaMax Workshop in Mannheim and IfW Kiel for helpful comments. We are grateful to Lorenz Kurrek for programming and assistance in running the sessions as well as discussions at an early stage of this project. Dietmar Fehr acknowledges financial support from the Deutsche Forschungsgemeinschaft (DFG) through CRC 649 "Economic Risk".

1. Introduction

Communication is pervasive in human interactions. While communication naturally serves a plethora of purposes, an important function is the exchange of information. Information about other individuals can, for example, be valuable in situations that involve collective action or trust as it helps to make judgements about others' willingness to cooperate or repay trust. Indeed, the more complex the social environment is, the more information is needed about potential interaction partners (Dunbar 1994). Communication allows individuals to gather the necessary information directly or through third parties, i.e., gossip. Gossip, as we define it here, is the exchange of information (positive or negative) about absent third parties (see also Foster 2004).¹ As such, talking about absent individuals increases the amount of information a person can absorb and it reduces monitoring costs because it is not necessary to observe norm violations directly to find out about who does not abide by the rules.²

In fact, gossip seems a widespread phenomenon not only in everyday life, but also in organizations (Noon and Delbridge 1993, Burt and Knez 1995). While it often has a negative connotation, gossip – for instance, at the coffee machine – may serve as social glue in organizations.³ For example, gossip may create bonds between co-workers or may be used to disseminate information and to communicate experiences, norms as well as formal and informal rules. Thereby, gossip plays a considerable role for an organization's culture and for the coordination of activities within organizations (Noon and Delbridge 1993). Gossip can also be used to influence the reputation of workers (Coleman 1990). First, knowing that potential norm violations are likely discussed by co-workers or communicated to alert others may actually induce norm compliance if workers care about their reputation.⁴ Second, in the same way as gossip is used to point out norm violations, it can be utilized to praise behavior to reinforce norm compliance. This dual purpose (positive and negative) of gossip raises the question how it affects efficiency within organizations.

¹ Foster (2004) provides an extensive discussion about definitions of gossip. The definition of gossip we adopt in our study captures the main essence of talking about other people (and implicitly includes evaluative content through positive and negative messages) and is tractable for our experimental setup. It is also different from rumor, which is characterized by unsupported information or opinions with no discernible source that is spread through the network (see e.g., Bloch, Demange and Kranton 2015).

² This, of course, does not preclude that the exchanged information is not true (e.g., Fehr and Fischbacher 2003).

³ The negative connotation may arise because gossip is commonly associated with idle talk, distracting people from work tasks, and with pointing out misdeeds, norm violations or derogatory comments (e.g., Dunbar 2004, Foster 2004).

⁴ Knez and Simester (2001) conjecture, for example, that the success of firm-wide team bonuses at Continental Airlines was in part due to mutual monitoring and the possibility that norm violations were discussed within work teams (see also Hamilton, Nickerson and Owan 2004).

In this paper, we study the role of gossip for interactions between individuals. In the field, it is typically difficult to observe gossip and its effect on behavior because outside observers often lack the context of conversations and it is costly to track and observe subsequent behavior. Therefore, we turn to evidence from a controlled laboratory experiment to study whether and how gossip affects behavior in situations that involve trust. Trust is an important ingredient for the functioning of organizations as it may, for example, reduce monitoring costs, increase efficiency and facilitate decentralized decision making (e.g., LaPorta et al 1997, Bloom, Sadun, van Reenen 2012). More precisely, we study a standard trust game (Berg et al 1995) in an environment where players are randomly matched in every round and modify this game by introducing a third party. The third party – who has no material interest in the trust game – can observe a trustee’s behavior and then gossip about it by sending a message to the trustor with whom the trustee will be paired (for the first time) in the next round. However, whether trustors receive gossip depends on the third party’s willingness to pay to gossip. Importantly, we do not restrict gossip content and thus explicitly allow positive and negative gossip. Moreover, gossip is non-verifiable for the trustor and it may even be incorrect.

Despite these (realistic) limitations of gossip, we find that gossip is highly efficiency-increasing, compared to a situation without any gossip. Trustors send higher amounts as first movers, and trustees return more money as second movers in the trust game, resulting in an efficiency increase – measured by the total available pie – of 28 percent. Our findings indicate that this increase in efficiency is mainly driven by the mere possibility of gossip. First, trustees already increase their returns even before gossip can take place. Second, we observe that gossip rarely takes place. That is, only about every fourth message is transmitted from gossipers to trustors. These findings provide direct evidence that norm compliance in work groups can be enforced by the possibility that misbehavior, but also compliance, may be discussed among co-workers (see e.g., Knez and Simester 2001, Hamilton, Nickerson and Owan 2003).

In two additional control treatments, we show that the efficiency-increasing effect of gossip cannot be explained by the mere fact of being observed by a third party, and that efficiency is attenuated if there is a small chance that gossip is not relevant to the current interaction. First, we find that pure observation of trustees does not induce higher returns compared to situations where gossip is possible or situations without third parties. Second, we provide evidence that noisy gossip – where information transmission from third parties to trustors can fail by providing information

about the wrong trustee – increases efficiency as well, but less so than if information transmission is undisturbed.

Our experiment on the effects of gossip is obviously related to the growing literature in economics on how feedback systems, reputational concerns, third-party interventions or communication affect the efficiency of interactions. In the case of feedback systems, this is often motivated by the case of online trading platforms where sellers and buyers may get feedback based on their transaction experience that, in turn, affects prices and trading volume (see, e.g., Duffy and Feltovich 2002, 2006, Bolton, Katok and Ockenfels 2004, Bohnet and Huck 2004, Bohnet et al 2005, Brown and Zehnder 2007, Bracht and Feltovich 2009, Charness, Du and Yang (2011), Huck, Lünser and Tyran 2012, Duffy, Xie and Lee 2013). Typically, in these types of experiments the transaction history of trustees (i.e., sellers) is publicly available. In contrast to our study, these experiments thus exogenously provide information of actual behavior that is based on the buyer's own experience ("hard evidence"). Such unambiguous feedback has been shown to increase efficiency in trust games. It is less clear whether endogenously provided and non-verifiable gossip has similar effects.

Another strand of literature discusses the effects of indirect reciprocity, which captures the idea that human subjects treat others nicer when those others have been nicer to yet another person (see, e.g., Seinen and Schram, 2006, Sommerfeld et al 2007, Engelmann and Fischbacher, 2009). Similar to the previously mentioned studies on feedback systems, subjects in these studies usually carry an image score on which interaction partners can base their helping decision. As such, information about a partner's reputation is always available and, importantly, this information is based on previous behavior and thus always correct. In contrast, information about previous behavior of an interaction partner in our study is only available if a third party gossips and we deliberately do not require that gossip conveys correct information (although it turned out that only a small fraction of gossip contains wrong information).

Our experiment also contributes to the literature on third-party interventions, which are essential for the enforcement of social norms. Most of these studies focus on punishment (and rewards) of third parties in the context of dictator or trust games (e.g., Fehr and Fischbacher 2004, Charness, Cobo-Reyes and Jiménez 2008, Almenberg et al 2011, Lerner et al 2014, Nikiforakis and Mitchell 2014). Third-party interventions in these papers have immediate and direct material consequences, for example, when third parties inflict punishment on norm violators and thereby reduce their payoffs. In contrast, a third-party intervention in our experiment ("gossip")

has no immediate material consequences for the target of this gossip. It only provides a positive or negative signal about a trustee's past behavior for the receiver of gossip, which may induce this person to trust or withdraw trust in future interactions. Thus, the third-party intervention we study is weaker than in the previous literature and we particularly consider the consequences of both positive and negative interventions.

From a more general perspective, our paper is related to the broad literature on the effects of communication on cooperative behavior (see, e.g., Isaac and Walker 1988, Crawford, 1998, Bochet, Page and Putterman 2006, Ben-Ner and Putterman 2009, Oprea, Charness and Friedman 2014). Typically, communication has been found to increase cooperation levels, also in trust games (Charness and Dufwenberg, 2006, 2010).⁵ In most of this literature, communication takes place between the interacting parties, though, while in our approach we let (unaffected) third parties communicate with the trustor only, and we are then particularly interested in how this communication affects a trustee's behavior who is *not* involved in communication.

2. Experimental Design

The design of our experiment is based on a standard trust game (Berg et al., 1995) where two players, a trustor and a trustee, interact with each other. Both players receive an endowment of 8 Experimental Currency Units (ECU). First, the trustor decides how much of his endowment $x \in [0, 8]$ to send to the trustee. The invested amount x is tripled by the experimenter and after observing the transferred amount the trustee has to decide about a back transfer of $y \in [0, 8+3x]$. Given standard selfish preferences, the subgame-perfect Nash equilibrium in this game is for the trustor to send nothing and for the trustee to return nothing. Of course, this equilibrium outcome is inefficient, and total payoffs would be higher if trustors sent a positive amount x to the trustee. Notice that investing $x = 8$ and returning twice the invested amount results in a Pareto efficient outcome that shares the total pie equally.

In total, we have four different treatments as explained below. In each treatment, subjects play the trust game 10 times with feedback after each round. Subjects keep their randomly assigned roles throughout the experiment and in each round they are anonymously and randomly matched

⁵ An important insight of this literature is that the communication protocol is crucial for enhancing cooperative outcomes. While free-form (chat) or face-to-face communication greatly improves cooperation, other forms of communication such as signaling actions through numerical values lead to mixed effects (see e.g., Wilson and Sell 1997, Duffy and Feltovich 2002, Bochet, Page and Putterman 2006 or Bochet and Putterman 2009).

with another subject. We also ask for trustors' beliefs about the trustees' back transfers in each round. Belief elicitation is not incentivized, though, in order to keep the decision environment as simple as possible and because our main focus is on the economic effects of gossip on the efficiency of interactions, but not on the influence of gossip on beliefs.

Our *Baseline* treatment is a replication of the standard trust game, repeated for 10 rounds with random matching. In all other treatments, we introduce a third party. The third parties receive an endowment of 16 ECU in each round and they are not affected by the actions of the trustors and trustees in material terms.⁶ The main treatment of interest is the *Gossip* treatment. In this treatment, the third party (gossiper) can gossip about trustees' behavior. More specifically, a gossiper observes the return ratio ($y/3x$) of a trustee in round t and can then write a message to the trustor who is going to be paired with the observed trustee in the next round $t+1$. Messages from gossipers to trustors are limited to 300 characters and free-form, but gossipers are instructed to not identify themselves or use offensive language.⁷ Note that, while the return ratio is informative about a trustee's behavior, it represents only minimal information about the behavior of a trustee. That is, gossipers are not informed about the absolute transfer that the respective trustee had received, and thus it remains unclear to the gossiper how much money both the trustor and the trustee earned. As a consequence, relative payoff comparisons cannot influence the decision to gossip.

Whether a message of a gossiper is transmitted to the trustor depends on the gossiper's willingness to pay (WTP). Using the Becker-DeGroot-Marschak (1964; hereafter BDM) mechanism, the gossiper has to state a price $P \in [0, 8]$ which he would be willing to pay for transmitting the message to the trustor. If this price P is smaller than a randomly drawn variable $r \in [0, 8]$, the message is not transmitted and the gossiper does not have to pay anything. If P is larger or equal to r , the message is transmitted and the gossiper pays a price r that is subtracted from the gossiper's endowment of 16 ECU per round. Correspondingly, we speak of gossip if messages are transmitted to a trustor. The trustor observes a box on his decision screen that displays the gossiper's message before he has to make his decision about the investment x . Like in experiments on third party punishment (Fehr and Fischbacher, 2004; Leibbrandt and Lopez-Perez, 2012) where third parties incur costs to punish norm-deviant behavior, we expect that (many)

⁶ The endowment of 16 ECU for the third party ensures that in the most favorable case where a trustor sent the full endowment of 8 ECU and the trustee returned twice this amount, the two players in the trust game and the third party earn the same amount.

⁷ We did not monitor messages in real time, but an ex-post analysis revealed that all subjects complied with the instructions.

gossipers have a willingness to pay to transmit their message to trustors in order to inform them about the matched trustee's behavior. Even more so, we expect a higher willingness to pay for negative messages to warn trustors about non-cooperative behavior of the trustee. Similar to the cooperation-increasing effect of third-party punishment (Mussweiler and Ockenfels, 2013), we expect that gossip increases cooperative behavior of trustees – making them more responsive to the trustors' investments – and, as a consequence, higher investments of trustors.

In addition to *Baseline* and *Gossip*, we have two further treatments that serve as controls. In *Observation*, the third party can only observe a trustee's behavior (more precisely the return ratio), but cannot gossip about the trustee. This control treatment allows us to check whether the mere fact of being observed in *Gossip* changes a trustee's (and as a consequence the trustor's) behavior. If this is the case, we expect to find no significant differences in behavior between *Observation* and *Gossip*. If gossip has an additional effect – on top of the potential effect of being observed (that can be measured by comparing *Observation* to *Baseline*) – then we expect to see a significant difference in behavior between *Observation* and *Gossip*.

Finally, treatment *Noisy Gossip* is identical to *Gossip*, except that a transmitted message can get confused with a message that was intended for a different trustor. More specifically, there is a 20 percent chance that a trustor receives a message that refers to the behavior of a trustee with whom he is *not* matched in the current round. This means that, conditional on the message being transmitted through the BDM-mechanism, the trustor receives the message that the gossiper has written about his current interaction partner only in four out of five cases. This information is public knowledge among gossipers, trustors and trustees. Given that the informational value of gossip is lower in *Noisy Gossip* than in *Gossip*, we expect weaker effects of gossip in the former treatment.

-- Table 1 about here --

We ran the experiment in the experimental laboratory of a large public university using z-Tree (Fischbacher, 2007). In total, we recruited 528 students from various fields from a database (ORSEE, Greiner, 2015) where students can register for participation in economic experiments. For the breakdown of the number of participants per treatments and for the details of the treatments, see Table 1. For treatments with a third party (*Observation*, *Noisy Gossip* and *Gossip*) we ran eight sessions with 18 subjects each. Because of the random matching of subjects, this results in eight independent observations for each of these treatments. For the *Baseline* we ran sessions with either

12 subjects or 24 subjects. In sessions with 24 subjects we randomly assigned subjects to matching groups of twelve in the beginning. Again, this yields eight independent observations for the *Baseline*. Subsequently, if not specified otherwise, we ran all our non-parametric tests based on the eight independent observations per treatment. Sessions lasted about one hour and average earnings were 14.2 Euro.

Note that the third parties can write messages in both *Gossip* and *Noisy Gossip*. Since the content of these messages is of potential interest, we recruited additional 20 students, who were not familiar with the purpose of the study, to rate and code all messages (in total we observe 717 messages in the two treatments). More specifically, we randomly divided all messages in four sets of about 180 messages and assigned five coders to each set of messages.⁸ Accordingly, each message was evaluated by five independent coders who had to indicate (i) whether a message was positive, neutral or negative (with respect to the description of the trustee's behavior), (ii) whether a message was detailed (i.e., a message indicated the return ratio), vague (i.e., a message did not explicitly state the return ratio) or wrong, and (iii) whether a message evaluated a trustee's behavior or not. In the analysis we will use the median of the answers to these three questions provided by the coders. In addition, the coders had to indicate one or more pre-specified categories describing the content of a message (see the appendix for the list of categories and the coder instructions).

3. Results

We start by comparing the *Gossip* treatment to the *Baseline* treatment in order to examine the effects of gossip on the efficiency of interaction. First we present data on trustors' and trustee's behavior both with and without an opportunity of a third party to gossip to the trustor. After having analyzed behavior in the trust game, we investigate the third party's willingness to gossip and the content of their messages. In section 3.2 we will then use our two control treatments (*Observation* and *Noisy Gossip*) to shed further light on how the presence of third parties affects behavior.

⁸ We divided the messages in smaller subset to not overstrain coders by evaluating all 717 messages. The evaluation took on average 43 minutes (for about 180 messages) and coders were paid 15 Euro.

3.1 The effects of gossip – *Baseline* vs. *Gossip*

3.1.1 Trustors' and trustees' behavior

Figure 1 shows the evolution of trustors' investments in the left panel and trustees' returns in the right panel in each of the ten rounds in *Baseline* and *Gossip*. It is evident that average investments are higher with gossip than without gossip in each round. While the difference is small and statistically insignificant in the first round (5.7 vs. 5.1; t-test: $t = 1.40$, $p = 0.165$), the gap widens over time. The difference becomes significantly different already in the second round according to a Mann-Whitney U-test (MWU-test, $z = 2.33$, $p = 0.02$), and remains so until the last round. Over all rounds the average investment in *Baseline* is 4.6 (58 percent of the endowment), but 6.4 (80 percent of the endowment) in *Gossip*. This corresponds to an increase of efficiency – measured by the total available pie – by 28 percent. Using session averages of investments as independent observations, we can clearly reject the hypothesis of equal investments in the *Baseline* and *Gossip* treatment (MWU-test: $z = 3.36$, $p = 0.01$).

-- Figure 1 about here --

Turning to trustees' behavior shown in the right panel of Figure 1, we observe that their returns are higher in *Gossip* than in *Baseline* in each round. Importantly, the difference in returns is already large and significantly different in the first round (9.4 vs. 6.1, t-test: $t = 3.27$, $p < 0.01$). This suggests that trustees anticipate that their reputation is at stake, and thus they increase their returns in *Gossip* even before gossip can take place and even though gossip is costly for the gossiper. The difference in returns increases further in later rounds. Averaging over all rounds, returns in *Baseline* are 5.0, whereas they are twice as high in *Gossip* (10.2). Again, we can reject the hypothesis of equal returns in the *Baseline* and *Gossip* treatment using an MWU-test ($z = 3.26$, $p < 0.01$).

Notice that trustors initially do not anticipate the higher returns of trustees in *Gossip* than in *Baseline*, as they expect an average return of 8.8 in *Gossip* and of 8.5 in *Baseline*. As a consequence of the different returns in round one, however, trustors' beliefs start diverging as well across treatments, with a second-round expectation of 9.6 in *Gossip* and of 7.8 in *Baseline*. Over all ten rounds, trustors expect returns of 10.6 in *Gossip*, but only 7.3 in *Baseline* (see Table 2) (MWU-test, $z = 2.31$, $p < 0.021$).

-- Table 2 and Table 3 about here --

In Table 3 we present a random-effects GLS regression that confirms the main insights presented so far. In column (1), we regress the trustor's investment on the expected returns of trustees, the round and the treatment dummy for *Gossip*. The latter is clearly significant, indicating higher investments in this treatment. Also expected returns are significant, meaning that trustors send more money when they expect higher returns. In column (2) we take the trustee's return as the dependent variable and show that returns are higher in *Gossip* even when controlling for investment levels (also significantly positive) and for the round (significantly negative).

-- Figure 2 about here --

Turning to profits, we see that trustors benefit from higher investments in *Gossip* because of the higher returns by trustees. On average, they earn 11.8 per round in *Gossip*, while average earnings are only 8.5 in *Baseline* (MWU-test, $z = 3.15$, $p < 0.01$). The latter earnings are only slightly higher than the initial endowment of 8. Profits do not differ across treatments for trustees. On average they earn 16.6 in *Baseline* and 17.1 in *Gossip*. This indicates that the return ratios ($y/3x$) must differ substantially across the two treatments. Indeed, while in *Baseline* the return ratio is on average 0.31, it is about 0.51 in *Gossip* (MWU-test, $z = 2.73$, $p < 0.01$). This means that trustees in *Baseline* return, on average, slightly less than the initial (and then tripled) investment of the matched trustor, thus keeping a large share of the generated pie for themselves. This is in line with previous results from trust games that show that investments are in most cases not profitable for trustors (Sutter and Kocher, 2007). In contrast, trustees are much more generous in *Gossip* where they return about half of the received investment. Figure 3 displays the cumulative distribution of return ratios. It is apparent that the distribution in *Gossip* stochastically dominates the distribution in *Baseline* (Kolmogorov-Smirnov test, $D = 0.88$, $p < 0.01$). We summarize our results so far as follows:

Result 1: *Gossip* leads to higher returns of trustees and higher investments of trustors. This means that gossip increases overall efficiency.

3.1.2 Contents of gossip and its effects on trustors

We now turn to the question whether gossipers use the opportunity to send messages to the trustors. The left panel of Figure 3 reveals that gossipers write a message in 359 out of 432 cases (83 percent) in the *Gossip* treatment. Moreover, gossipers have on average a strictly positive willingness to pay

($P = 2.23$) for sharing a message (gossip), as evidenced in the right panel of Figure 3. More detailed, the willingness to pay for a message with content is on average $P = 2.58$. Somewhat surprisingly, we note also a positive willingness to transfer an empty message ($P = 0.51$). Yet, gossipers' willingness to pay for an empty message is zero ($P = 0$) in the vast majority of 64 out of 73 cases. Eight of the remaining nine cases are due to one subject who stated a positive willingness to pay for an empty message in all but the first two rounds. Overall, the willingness to pay for sending a message is significantly higher in the first half than in the second half of the experiment, indicating that gossipers consider the opportunity to send a message as more important earlier on in the experiment.⁹

We observe a considerable degree of heterogeneity in the gossipers' willingness to pay. About 26 percent of price statements are $P = 0$, in the range $0 < P \leq 1$ we observe 21 percent, and the remaining 53 percent satisfy $1 < P \leq 8$. The share of subjects who always state $P = 0$ is ten percent, while 29 percent have an average willingness to pay of less than 1 over all ten rounds. The remaining 61 percent of subjects state on average $P > 1$.

Interestingly, there is no relation between a gossipers' willingness to pay and whether the observed return ratio is relatively high or low. For example, for return ratios smaller than or equal to one third, we observe $P = 2.3$ and for return ratios higher than one third $P = 2.2$. In general, the correlation between WTP and observed return ratio is low and insignificant (Spearman's $\rho = 0.05$, $p > 0.32$).

We turn next to a more detailed analysis of the content of messages. First, a large majority of written messages report actual observed behavior in a detailed way (75 percent). That is, these messages always include the observed return ratio and almost every third of these messages contains some kind of judgment about the observed behavior. More specifically, such judgments are used to warn trustors about malicious behavior of trustees (in 40 percent of messages with judgements; e.g., "no repayment, don't send anything") or to encourage investments in case of observing positive behavior (54 percent; e.g., "totally fair, he repaid 2/3!"). Second, a fraction of written messages is vague (18 percent), which means that there is some leeway in how they can be interpreted by the trustor (e.g., "he paid back reasonably" or "perfect gameplay").¹⁰ Third, there is

⁹ Regressing the willingness to pay for sending a message on round reveals a significant negative coefficient estimate (-0.113 , *std. err.* 0.037 , $p = 0.018$).

¹⁰ There is some heterogeneity in fairness views of gossipers, which makes the interpretation of messages possibly harder, in particular if gossipers do not state the observed return ratio. For example, while some gossipers find a return ratio of 0.5 fair, others find return ratios below 0.66 unacceptable.

a small fraction of written messages (7 percent) that either contains wrong information or pointless information (e.g., “Participant A”), which we label as wrong messages. While the WTP for detailed messages (2.73) and vague messages (2.50) is significantly higher than for wrong messages (0.55) (t-tests, $p < 0.05$), there is no significant difference between the WTP for detailed and vague messages (t-test, $t = 0.75$, $p = 0.45$).

The positive willingness to pay for sharing information is a precondition that trustors do receive gossip. In fact, messages are transferred in 26 percent of cases (in the other cases, the BDM-mechanism did not lead to the transfer of the message). How do trustors react to gossip? If a message is transmitted and negative, trustors typically react with low investments. In such cases, they invest on average only 3.2. Notice that this is even lower than the average investment in *Baseline*, albeit the difference is not statistically significant. In strong contrast, investments are clearly higher after receiving a neutral (6.0) or a positive message (7.2). Moreover, subjects clearly interpret receiving no message (i.e., in 74 percent of cases) as good news. That is, investments average to 6.7 when trustors receive no message. While investments are not significantly different after receiving a neutral, positive or no message (all $p > 0.6$, MWU-tests), the difference between investments after a negative message and after a positive, neutral or no message, respectively, is always significant (all $p < 0.03$, MWU-tests). Moreover, investments in cases of positive, neutral or no messages are significantly higher than investments in *Baseline* (all $p < 0.02$, MWU-tests). Thus, gossip can affect behavior and efficiency in two ways. Gossip about negative behavior clearly hurts efficiency, whereas positive, neutral or no gossip leads to higher efficiency compared to a situation where gossip is not possible.

Result 2: Subjects display a positive willingness to pay for gossip, and therefore 26 percent of messages are communicated to the trustors. Trustors react very sharply to negative gossip, but also very positively by increasing investments when they receive positive, neutral or even no messages.

3.2 Control Treatments

3.2.1 Third-party observation without gossip

The previous section has illustrated the impact of gossip on the efficiency of interactions in our trust game. The possibility of gossip seems to discipline trustees through reputational concerns, which results in higher returns than in a situation where no gossip is present. Interestingly, this effect prevails despite the fact that only about every fourth message is transmitted from a gossip

to a trustor. The mere opportunity that some third party can gossip with a future interaction partner is sufficient to trigger higher returns.

However, a necessary condition for gossip is the observation of behavior. This raises the question to what extent the mere observation of a trustee changes their behavior rather than the possibility of gossip from gossipers to trustors. To address this question we run a control treatment (*Observation*) where a third party can only observe the return ratio of a trustee in round t , but has no possibility to send a message to a trustor. All other details are identical to the *Gossip* treatment, including that third parties receive an endowment of 16 ECU per round.

-- Figure 4 about here --

Figure 4 displays the evolution of average investments and returns in *Observation*. It is apparent that trustees return substantially less when the third party can only observe trustees' behavior compared to the case where third parties can also gossip about trustees' behavior. Already in the first round, returns are (weakly) significantly lower in *Observation* (7.6) than in *Gossip* (9.4) (t-test, $t = 1.68$, $p = 0.097$). Across all rounds, trustees return on average 6.9 in *Observation* and we can clearly reject the hypothesis of equal returns in *Gossip* (with an average of 10.2) and *Observation* (MWU test, $z = 2.42$, $p = 0.016$).

Importantly, the lower returns are not driven by lower first-round investments in *Observation*. Although the first-round investments in *Observation* are, on average, slightly lower than in *Gossip* (5.3 vs. 5.7), the difference is not significant (t-test, $t = 0.97$, $p = 0.33$). Across all rounds, the average investments are, however, significantly different between *Observation* (with 5.6) and *Gossip* (with 6.4; MWU test, $z = 2.1$, $p = 0.036$).

Looking at the return ratio provides an alternative view on trustees' returns that takes the different investment levels in the two treatments into account (see also Figure 2 for the cumulative frequency distribution). The average return ratio in *Observation* is 0.35, which is significantly different from the average return ratio of 0.51 in *Gossip* (MWU test, $z = 2.84$, $p < 0.01$).

Although the possibility of gossiping is essential for higher returns and investments, it is still conceivable that third-party observation affects trustees' behavior because they care about how they are seen by others even though there are no direct consequences in our setting. Comparing trustees' behavior in *Observation* and *Baseline* shows, however, that pure observation has no statistically significant effect on trustees' behavior. While returns are, on average, higher with

observation (6.9) than without third-party observation (5.0), the difference is not statistically significant (MWU test, $z = 1.37$, $p = 0.17$). It is noteworthy, though, that investments are higher in *Observation* (5.6) than in *Baseline* (4.5; MWU test, $z = 2.1$, $p < 0.04$). Overall, the findings suggest that pure observation of trustees is not enough to change their behavior.

The evidence so far shows that the possibility to gossip is essential for changing trustees' behavior, probably because of reputational concerns. Arguably, reputational concerns are characterized by strategic considerations and higher complexity, which may also be reflected in longer reaction times.¹¹ Indeed, if third parties have an opportunity to gossip, trustees take significantly more time to make their decision. In *Gossip*, trustees need 17.5 seconds on average to reach a decision, while they need only 13 seconds in *Baseline* and 13.1 seconds in *Observation*. The differences are statistically significant according to pairwise MWU-tests between *Gossip* and the other two treatments ($z > 2.4$ and $p < 0.02$ in both comparisons).¹² Reaction times are not significantly different between *Baseline* and *Observation*, lending further support to the interpretation that pure third party observation does not induce reputational concerns.

Result 3: The mere observation of trustees through third parties (who are unaffected by the behavior of trustors and trustees) cannot explain the effects of gossip that are summarized in Result 1. Both trustors' and trustees' behavior is significantly different between *Observation* and *Gossip*.

3.2.2 Noisy gossip

We have seen that the possibility of gossip greatly enhanced efficiency of interactions, despite the fact that gossipers and trustees were strangers and gossip was non-verifiable. Although gossip barely took place, it is important to understand the role of the situational context. Previously, we have considered a situation where gossip was non-verifiable, but where trustors knew that a received message was about their present interaction partner. This assumption is possibly met in groups of manageable size, but arguably not in large-scale groups where the information flow and interaction frequency grow exponentially with group size and where it is thus harder for potential

¹¹ There is an emerging literature arguing that fair behavior is intuitive, resulting in faster reaction times (e.g., Rubinstein 2007, Rand, Green and Nowak 2012, Cappelen et al. 2015 among others, but see e.g., Tinghög et al. 2013 for conflicting evidence). Note, that our interest here is not whether intuitive processes favor fair decisions, but to see whether reaction times (as a proxy for attention and deliberation) are longer when one's reputation is at stake, i.e., in a more complex decision situation.

¹² Notice that trustors in *Baseline* (10.8 seconds) and *Observation* (10.0 seconds) reach faster decisions than in *Gossip* (15.8), too (MWU-tests, $p < 0.01$ in both comparisons). While these differences can stem from the fact that trustors have to read and process messages in *Gossip*, the difference in reaction time is still (weakly) significantly different if we only consider decisions in *Gossip* where trustors do not receive a message (MWU-test, $z = 1.79$, $p = 0.074$).

gossipers to track behavior of others and easier to confuse information. That is, there may be situations where it is not clear that gossip is perfectly reliable, even when there is no reason to believe that the shared information is wrong.

To address this issue, we implement another control treatment. In *Noisy Gossip* we introduce the possibility that a transmitted message is confused with another message (i.e., the shared information is about some other trustee) with a probability of 20 percent. This is public knowledge for the three players (trustor, trustee and gossiper). All other details are the same as in the *Gossip* treatment.

There are two noteworthy features of this treatment. First, in addition to the non-verifiability of gossip for the trustor, we explicitly manipulate the reliability of gossip for trustors as transmitted messages can be misleading even though the content may accurately describe some other trustee's behavior. Second, the link between gossipers and trustees is exactly the same as in *Gossip*, i.e., gossipers observed the return ratio of the trustee. However, while in *Gossip* there is only uncertainty about whether third parties engage in costly gossip, it is now more salient that a trustor may not get information about the current interaction partners' past behavior.

Figure 4 shows investments and returns in the *Noisy Gossip* treatment alongside the other three treatments. As with *Baseline* and *Observation*, first-round investment behavior does not differ statistically between *Noisy Gossip* (5.3) and *Gossip* (5.7) (t-test, $t = 1.03$, $p = 0.3$). Although investments increase in the first half of the experiment, they remain always below those in *Gossip*. The average investments over all rounds are 5.5 in *Noisy Gossip* and 6.4 in *Gossip* (MWU test, $z = 2.26$, $p = 0.024$).

Interestingly, trustees' returns are very similar to the previously observed returns in *Observation*, as can be seen in the right-hand panel of Figure 4. On average, trustees return 6.7 in *Noisy Gossip* (see also Table 2), which is significantly lower than the 10.2 in *Gossip*, though (MWU test, $z = 2.94$, $p < 0.01$). Again, lower returns are not the result of initially low investments as there are no differences in first-round investments in *Gossip* and *Noisy Gossip*. Rather, the return ratio of 0.36 is comparable to the return ratio in *Observation* and significantly different from the return ratio (0.51) in *Gossip* (MWU test, $z = 2.94$, $p < 0.01$). These observations suggest that gossip loses its potential for increasing efficiency when it becomes less reliable, i.e., we observe a reduction of efficiency by 14 percent compared to *Gossip*. But since gossip always involves the observation of behavior through a third party we observe similar effects on trustee behavior in *Noisy Gossip* and *Observation*.

It is unlikely that the lower return rates in *Noisy Gossip* are due to less careful considerations of decisions or inattention because of the noisiness of gossip. Taking trustees' reaction times as a proxy for attention and deliberation of decisions indicates little difference between *Noisy Gossip* (16.4 seconds) and *Gossip* (17.5 seconds) (MWU test, $z = 0.74$, $p = 0.46$).¹³

Even though noisy gossip is less effective than gossip, it increases efficiency compared to *Baseline*. From Figure 4 it is apparent that average investment levels are clearly higher throughout all rounds than in *Baseline*. The average investment level in *Noisy Gossip* is 5.5. Accordingly, the total pie in the trust game increases by 16 percent compared to *Baseline* (4.6). The difference in average investments between *Baseline* and *Noisy Gossip* is significant using a MWU test ($z = 2.00$, $p = 0.046$).

It is noteworthy that the 20 percent-likelihood of confusing messages does not lead to less written messages in *Noisy Gossip*, compared to *Gossip*. We observe that gossipers write only slightly less messages (81 vs. 83 percent of all possible cases) and have only a slightly lower WTP for sharing this information (2.53) in *Noisy Gossip* than in *Gossip* (2.58). Moreover, the transmission rate is similar across the two treatments (25 vs. 26 percent). Neither transmission rates nor WTP are significantly different across treatments (MWU test, $p > 0.4$ in both cases). There are only small and statistically insignificant differences in message contents. While the share of neutral messages is about the same in both *Gossip* and *Noisy Gossip* (60 percent), there is a slight shift from positive to more negative messages from 14 percent in *Gossip* to 18 percent in *Noisy Gossip* (Fisher's Exact test, $p = 0.249$). Messages in *Noisy Gossip* are also slightly less detailed (68 percent) and thus more often vague (23 percent) or false (9 percent) (Fisher's Exact test, $p = 0.12$).

Despite these small differences in content, trustors react in a less pronounced way than in *Gossip* upon receiving a message. That is, investments are higher after receiving a negative message and lower after no message, neutral or positive message compared to *Gossip*. More specifically, after a negative message, they invest on average 4.1, which is slightly lower than investments after a neutral message (4.6). Investments are higher after receiving a positive (5.2) or no message (5.8). However, only the difference between receiving no message and a negative messages (MWU test, $z = 2.32$, $p = 0.021$) but not between receiving a positive and negative message (MWU test, $z = 1.16$, $p = 0.25$) is statistically significant. This is probably a consequence

¹³ The same holds for decision times of trustors (15.3 versus 15.7 seconds).

of the higher uncertainty whether the message content really refers to the currently matched trustee or (mistakenly) to someone else.

Result 4: If gossip about trustees' behavior becomes more unreliable, gossip loses a considerable part of its efficiency-increasing effect. Efficiency is therefore lower in *Noisy Gossip* than in *Gossip*.

4. Conclusion

Gossiping about others occurs frequently. However, gossip is typically associated with idle chat or negative behavior. Often, it is second-hand information because it is not based on the gossiper's own experience with the gossip target. Despite these limitations of gossip, and despite its often negative connotation, gossip may have positive effects on the efficiency of interactions. According to Dunbar (1994, 2004), language has evolved to facilitate social bonding in large social groups through the exchange of social information within groups. In this sense talking about other people or gossiping allows to keep track of behavior or changes in the social network and to maintain efficient relationships. This view suggests that gossip may prevent free riding or malicious behavior of individuals whose behavior may become the substance of gossip.

In this paper, we have investigated how gossip may affect the efficiency of interaction between two interaction partners, a trustor and a trustee, in a trust game. We find that if a third party can gossip about the behavior of the trustee, returns are much higher, and, as a consequence, also the investments of the trustor increase significantly, compared to a baseline condition without any third party that may gossip. These effects of gossip increase the efficiency of interaction significantly by raising the total pie to be distributed in the trust game by 28 percent. Interestingly, there is more gossiping going on in the earlier part of the experiment than in the later part, which may indicate that gossip is particularly useful to help establishing an efficient level of interaction early on, which then stays high until the end of our experiment. Trustors react in particular to negative gossip about the trustee with whom they are paired in a particular round. As expected, negative gossip leads to a withdrawal of trust and thus lower investments of trustors. On the other hand, investments react positively to positive, neutral or even no messages, as investments are higher than in a baseline condition without a gossiper. Accordingly, the possibility of gossip fosters trust by selectively trusting interaction partners, which reflects findings that the ability to choose partners can greatly enhance cooperation in social dilemma situations (e.g., Coricelli, Fehr and Fellner, 2004, Page, Putterman and Unel 2005, Charness and Yang 2014, Kamei and Putterman 2015). Remarkably, the

efficiency-increasing effects of gossip hold despite the fact that only about one quarter of messages is successfully transmitted to the trustor. This shows that the availability of gossip itself has a strong effect on the efficiency of interactions.

One straightforward candidate to explain the latter finding is the fact that gossip requires the observation of trustees' behavior through the gossiper. Therefore, it might simply be the fact of being observed that triggers higher efficiency with gossip. In a control treatment (*Observation*) where third parties can only observe trustees, but not gossip about their behavior, we find that mere observation of trustees does not trigger the strong effects of gossip in the main treatment. Therefore, observation can at best explain parts of these effects. It is essential that observers have the possibility to gossip about the observed behavior in order to add to the efficiency-increasing effect. Because it is sufficient that someone may gossip about trustees' behavior to deter norm violations, it turns out that the threat of gossip is a cheap and efficient instrument of social sanctioning and, at the same time, circumvents the second-order public good problem of punishment (Coleman 1990). This provides, therefore, another explanation for why costly punishment of norm violations is rarely observed outside the laboratory (e.g., Balafoutas and Nikiforakis 2012, Balafoutas et al 2014).

Our findings highlight that gossip can be a powerful force encouraging trust and trustworthiness despite that gossipers and trustees are strangers and despite that the potentially shared information does not originate from their common experience. Possibilities that a third party may share information about one's behavior with a future interaction partner are ubiquitous in daily and organizational life. However, it seems crucial that gossip – whether correct or not – pertains to the current interaction partner. Indeed, in a second control treatment (*Noisy Gossip*), we demonstrate that gossip has to be reliable to unfold its full effects. When it is possible that trustors receive gossip about trustees with whom they are *not* paired, then efficiency is reduced by 14 percent, compared to a situation where gossip always refers to the trustee with whom an interaction takes place. While this is not due to less frequent gossiping (compared to *Gossip*), it is rather because trustees reduce their returns in response to the reduced reliability of gossip, and consequently trustors lower their investments as well. Notwithstanding these possible restrictions of gossip, our results emphasize its potential for increasing the efficiency of interactions. Gossip can, in particular, serve as cost-effective instrument to enhance trust and trustworthiness in settings where it is difficult to keep directly track of behavior of potential interaction partners. In such cases,

gathering information about others through intermediaries may be useful to assess the trustworthiness or willingness to cooperate of others.

References

- Almenberg, J., Dreber, A., Apicella, C. L., & Rand, D. G. (2011). Third party reward and punishment: group size, efficiency and public goods. In *Psychology and punishment*. New York: Nova Publishing.
- Balafoutas, L., Nikiforakis, N., and Rockenbach B. (2014). Direct and indirect punishment among strangers in the field. *Proceedings of the National Academy of Sciences*, 111(45): 15924-15927.
- Balafoutas, L., and Nikiforakis, N. (2012). Norm enforcement in the city: A natural field experiment. *European Economic Review*, 56(8): 1773-1785.
- Becker, Gordon M., DeGroot, Morris H. and Marschak, Jacob (1964). Measuring utility by a single-response sequential method. *Behavioral Science*, 9(3): 226-232.
- Ben-Ner, A., and Putterman, L. (2009), Trust, communication, and contracts: An Experiment. *Journal of Economic Behavior and Organization*, 70(1-2): 106-121.
- Berg, J., Dickhaut, J., McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior*, 10(1): 122-142.
- Bloch, F., Demange, G., and Kranton, R. (2015). Rumors and Social Networks, *Working paper*.
- Bloom, N., Sadun, R. and Van Reenen, J. (2012). The organization of firms across countries. *Quarterly Journal of Economics*, 127(4): 1663-1705
- Bochet, O., Page, T., and Putterman, L. (2006). Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior and Organization*, 60(1): 11-26.
- Bochet, O., and Putterman, L. (2009). Not just babble: opening the black box of communication in a voluntary contribution experiment. *European Economic Review*, 53(3): 309-326.
- Bohnet, I., and Huck, S. (2004). Repetition and reputation: Implications for trust and trustworthiness when institutions change. *American Economic Review, Papers and Proceedings*, 94(2): 362-366.
- Bohnet, I., Harmgart, H., Huck, S. and Tyran, J.-R. (2005). Learning trust. *Journal of the European Economic Association*, 3(2-3): 322-329.
- Bolton, G., Katok, E., and Ockenfels, A. (2004). How effective are electronic reputation mechanisms? An experimental investigation. *Management Science*, 50(11): 1587-1602.
- Bracht, J., and Feltovich, N. (2009). Whatever you say, your reputation precedes you: observation and cheap talk in the trust game. *Journal of Public Economics*, 93(9-10): 1036-1044.

- Brown, M., and C. Zehnder (2007). Credit reporting, relationship banking, and loan repayment. *Journal of Money, Credit, and Banking*, 39(8): 1884-1918.
- Burt, R., and Knez, M. (1995). Kinds of third-party effects on trust. *Rationality and Society*, 7(3): 255-92.
- Cappelen A., Nielsen, U., Tungodden, B., Tyran, J-R., Wengström, E. (2015). Fairness is intuitive. *Experimental Economics*, forthcoming.
- Charness, G., and Dufwenberg, M. (2006). Promises and partnership. *Econometrica*, 74(6): 1579-1601.
- Charness, G., and Dufwenberg, M. (2010). Bare promises: An experiment. *Economics Letters*, 107(2): 281-283.
- Charness, G., Cobo-Reyes, R., and Jimenez, N. (2008). An investment game with third-party intervention. *Journal of Economic Behavior and Organization*, 68(1): 18-28.
- Charness, G., Du, N., and Yang, C-L., (2011). Trust and trustworthiness reputation in an investment game. *Games and Economic Behavior*, 72(2): 361-375.
- Charness, G., and Yang, C-L. (2014). Starting small toward voluntary formation of efficient large groups in public goods provision. *Journal of Economic Behavior and Organization*, 102: 119-132.
- Coleman, J. (1990). *Foundations of social theory*. Cambridge, MA: Harvard University Press.
- Coricelli, G., Fehr, D., and Fellner, G. (2004). Partner selection in public goods experiments. *Journal of Conflict Resolution*, 48(3): 356-378.
- Crawford, V. (1998). A survey of experiments on communication via cheap talk. *Journal of Economic Theory*, 78(2): 286-298.
- Duffy, J., and Feltovich, N. (2002). Do actions speak louder than words? An experimental comparison of observation and cheap talk. *Games and Economic Behavior*, 39(1): 1-27.
- Duffy, J., and Feltovich, N. (2006). Words, deeds and lies: Strategic behavior in games with multiple signals. *Review of Economic Studies*, 73(3): 669-688.
- Duffy, J., Xie, H., and Lee, Y-J. (2013). Social norms, information and trust among strangers: theory and evidence. *Economic Theory*, 52(2): 669-708.
- Dunbar, R. (1994). *Grooming, gossip, and the evolution of language*. London: Faber & Faber.
- Dunbar, R. (2004). Gossip in evolutionary perspective. *Review of General Psychology*, 8(2): 100-110.

- Engelmann, D., and Fischbacher, U. (2009). Indirect reciprocity and strategic reputation building in an experimental helping game. *Games and Economic Behavior*, 67(2): 399-407.
- Fehr, E., and Fischbacher, U. (2003). The nature of human altruism, *Nature* 425, 785-791.
- Fehr, E., and Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior*, 25(2): 63-87.
- Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics*, 10(2): 171-178.
- Foster, E. (2004). Research on gossip: Taxonomy, methods, and future directions. *Review of General Psychology*, 8(2): 78-99.
- Greiner, B. (2015). Subject pool recruitment procedures: Organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1): 114-125.
- Hamilton, B., Nickerson, J., and Hideo O. (2003). Team incentives and worker heterogeneity: An Empirical analysis of the impact of team on production and participation. *Journal of Political Economy*, 111(3): 465-497.
- Huck, S., Lünser, G., and Tyran, J.-R. (2012). Competition fosters trust. *Games and Economic Behavior*, 76(1): 195-209.
- Isaac, M., and Walker, J. (1988). Communication and free-riding behavior: The voluntary contributions mechanism. *Economic Inquiry*, 26(4): 585-608.
- Kamei, K., and Putterman, L. (2015). Play it again: Partner choice, reputation building and learning from finitely-repeated dilemma games. *Economic Journal*, forthcoming.
- Knez, M., and Simester, D. (2001). Firm-wide incentives and mutual monitoring at continental airlines. *Journal of Labor Economics*, 19(4): 743-772.
- La Porta, R., Lopez de Silanes, F., Shleifer, A., and Vishny, R. (1997). Trust in large organizations. *American Economic Review*, 87(2): 333-338.
- Leibbrandt, A., and Lopez Perez, R. (2012). An exploration of third and second party punishment in ten simple games. *Journal of Economic Behavior & Organization*, 84(3): 753-866.
- Lergetporer, P., Angerer, S., Glätzle-Rützler, D., and Sutter, M. (2014). Third-party punishment increases cooperation in children through (misaligned) expectations and conditional cooperation. *Proceedings of the National Academy of Sciences*, 111(19): 6916-6921.
- Mussweiler, T., and Ockenfels, A. (2013). Similarity increases altruistic punishment in humans. *Proceedings of the National Academy of Sciences*, 110(48): 19318-19323.

- Nikiforakis N., and Mitchell, H. (2014). Mixing the carrots with the sticks: Third party punishment and reward. *Experimental Economics*, 17 (1): 1-23.
- Noon, M., and Delbridge, R. (1993). News from behind my hand: Gossip in organizations. *Organization Studies*, 14(1): 23-36.
- Oprea, R., Charness, G. and Friedman, D. (2014). Continuous time and communication in a public-goods experiment. *Journal of Economic Behavior and Organization*, 108: 212-223.
- Page, T., Putterman, L., and Unel, B. (2005). Voluntary association in public goods experiments: reciprocity, mimicry, and efficiency. *Economic Journal*, 115(506): 1032-1053.
- Rand, D., Greene, J. and Nowak, M. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416): 427-430.
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: A study of response times, *Economic Journal*, 117(1424): 1243-1259.
- Seinen, I., and Schram, A. (2006). Social status and group norms: Indirect reciprocity in a repeated helping experiment. *European Economic Review*, 50(3): 581-602.
- Sommerfeld, R., Krambeck, H-J., Semmann, D., and Milinski, M (2007). Gossip as an alternative for direct observation in games of indirect reciprocity. *Proceedings of the National Academy of Sciences*, 104(44): 17435-17440.
- Sutter, M., and Kocher M. (2007). Trust and trustworthiness across different age groups. *Games and Economic Behavior*, 59(2): 364-382.
- Tinghög, G., Andersson, D., Bonn, C., Böttiger, H., Josephson, C., Lundgren, G., Västfjäll, D., Kirchler, M., and Johannesson, M. (2013). Intuition and cooperation reconsidered. *Nature*, 498(7452): E1-E2.
- Wilson, R., and Sell, J. (1997). Liar, liar, ..., cheap talk and reputation in repeated public goods settings. *Journal of Conflict Resolution*, 41(5): 695-717.

Tables and Figures

Table 1. Summary of Treatments

Treatment	# of sessions (# of indep. obs.)	# of subjects	Third party's possible actions
<i>Baseline</i>	5 (8)	96	No third party present
<i>Gossip</i>	8 (8)	144	Gossip
<i>Observation</i>	8 (8)	144	no Gossip, only observation
<i>Noisy Gossip</i>	8 (8)	144	Gossip, noisy transmission

Table 2. Summary of Results for all Treatments

	Investor		Trustee	Gossiper		N
	Amount sent	Expected return	Amount returned	WTP	% message transmitted	
<i>Baseline</i>	4.55 (2.90)	7.34 (5.82)	5.03 (5.38)	--	--	480
<i>Gossip</i>	6.41 (2.17)	10.61 (4.90)	10.17 (5.24)	2.23 (2.30)	0.26 (0.44)	480
<i>Observation</i>	5.60 (2.86)	8.62 (5.89)	6.95 (6.09)	--	--	480
<i>Noisy Gossip</i>	5.47 (2.94)	9.03 (5.70)	6.65 (5.97)	2.06 (2.29)	0.25 (0.43)	480

Notes: Standard deviations in parentheses. WTP: Gossiper's average willingness to pay for transmitting their message.

Table 3. Trustor and trustee behavior in *Baseline* and *Gossip* (GLS regressions)

Dependent Variable	Investment	Return
Expected Return	0.376*** (0.029)	
Round	-0.011 (0.026)	-0.131*** (0.040)
<i>Gossip</i>	0.638** (0.303)	2.813*** (0.654)
Investment		1.245*** (0.096)
Constant	1.848*** (0.285)	0.088 (0.641)
N	960	960
R ²	0.62	0.48

Notes: Random-effects GLS regression with standard errors clustered at the session level. *** (**) [*] indicates significance at the 1% (5%) [10%] level.




Figure 1: Evolution of Trustor and Trustee Behavior in *Baseline* and *Gossip*.




Figure 2: Trustees' return ratio across all rounds in all treatments.




Figure 3: Frequency of gossip and willingness to pay (WTP) in treatment *Gossip*.




Figure 4: Evolution of Trustor and Trustee Behavior in all Treatments.

Appendix

Instructions – *Gossip* Treatment

Welcome to the experiment.

Note that from now on it is not allowed to talk to other participants and to exchange any information. If you have a question, please raise your hand and we will come to your place and answer your question in private. It is important that you do not ask your question in public. If you do not comply with the rules of the experiment, we have to discontinue the experiment.

General Overview

There will be three different player roles in this experiment (player A, B and C). Depending on your player role you will make different decisions. The roles will be determined randomly in the beginning and are fixed throughout the experiment.

The experiment consists of 10 rounds. At the beginning of a round, we will randomly form groups of three. Each group consists of one player A, one player B and one player C. You will neither learn the identity of the other players in your group nor the identity of players in other groups. That is, you remain anonymous throughout the experiment and afterwards. No other participant can associate you with your role and your decisions in the experiment.

Your endowment and all payoffs in the experiment will be denoted as ECU (Experimental Currency Units) and converted into Euro at the end of the Experiment. The conversion rate is:

$$10 \text{ ECU} = 1 \text{ Euro}$$

Task Description

In the following we will give you an overview of the tasks of the three players in this experiment.

In each round each player will be endowed with a certain amount of ECU. Player A has the possibility to send any share of their endowment to player B. Each sent ECU will be tripled by the experimenter. Player B will then decide how much of the available ECU to return to player A. Player C will observe the decision of player B. More specifically, Player C will learn which share of the received amount (in percent) player B returned to player A. Player C can then decide whether to comment the behavior of player B and to send a message to the player A who is next paired with the observed player B. This message will be available before this player A will make their decision in the next round.

In the following, we will describe each player role for you in more detail.

Player A

You are randomly matched with a player B and a player C. In each round you will receive an endowment of 8 ECU.

Your task is to decide how much of your endowment to send to player B. You can send any amount between 0 and 8 ECU. The amount you send will be tripled by the experimenter and credited to the account of Player B (in addition to their endowment). Player B then decides how much of the available amount to return to you.

Example: If you send 4 ECU to B, player B will receive 12 ECU. If you send your whole endowment of 8 ECU, player B will receive 24 ECU.

In addition you have to indicate how much you expect back from player B (based on your decision). Note that this information will not be transferred to player B.

If some player C decides to send you a message at some point during the experiment, you will get this information before you make your decision of how much to send to player B. The message may contain some information on the past behavior of player B you are paired with.

Your payoff in a given round is:

$$\text{Payoff} = \text{Endowment} - \text{Amount sent} + \text{returned Amount of Player B}$$

Player B

You are randomly matched with a player A and a player C. In each round, you will receive an endowment of 8 ECU. Player A has the option to send you any amount between 0 and 8 ECU. This amount will be tripled by the experimenter. That is, you can receive up to 24 ECU in addition to your endowment of 8 ECU in each round. You will see all necessary information on your screen.

You will then decide how much of your total available ECU to return to player A. You can return any amount to player A.

Example: Suppose Player A sent you 6 ECU. This amount is tripled and you receive 18 ECU in addition to your endowment. That is, your total available amount is 26 ECU. You can now return any amount between 0 and 26 ECU to player A.

Player C will learn which share of the received amount (18 ECU in the example above) you returned to player A. Player C can use this information to send a message to the next player A you are paired with. That is, in the next round, player A will possibly learn about your behavior in this round.

Your payoff in a given round is:

$$\text{Payoff} = \text{Endowment} + \text{received Amount} - \text{Return to Player A}$$

Player C

You are randomly matched with a player A and a player B. In each round, you will get information about the behavior of player B. You will learn how much player B returned to player A, expressed as the share of the amount that player B received from player A (return ratio).

Example 1: Player A sent 6 ECU. This amount was tripled such that player B received 18 ECU. Player B returned 6 ECU. That means the return ratio is 33.3 percent. Calculation: Return (6) divided by the received amount (18) multiplied by 100 ($6/18 \times 100 = 33.3\%$).

Example 2: Player A sent 6 ECU. This amount was tripled such that player B received 18 ECU. Player B returned 0 ECU. That means the return ratio is 0 percent.

Example 3: Player A sent 0 ECU. This amount was tripled such that player B received 0 ECU. Player B returned 0 ECU. That means the return ratio is 0 percent. However, in this case you will also learn that player A has sent 0 ECU.

Based on this information, you can write a message, which will be transmitted to the player A, who is paired next with the player B you have observed in this round. This player A will receive your message before making a decision. Whether and how your message will be transmitted is explained below.

If you decide to write a message, you have 300 characters to formulate your message. **Note that you are not allowed to use offensive language and to reveal your identity or others identity, for example by revealing your computer number, sex, appearance, etc.**

When will your message be transmitted?

After you wrote your message, you will be asked to indicate your willingness to pay to transmit your message to player A. You can indicate any amount between 0 and 8 ECU (and up to 2 decimal points). In each round you will get an endowment of 16 ECU.

In each round the computer will draw a random number, which will be compared to your willingness to pay. If your willingness to pay is equal to or higher than the random number drawn by the computer, your message will be transmitted and you will have to pay the random number. If your willingness to pay is smaller than the random number, your message will not be transmitted and you will pay nothing.

Note that if your message is transmitted to player A, you will not pay your stated willingness to pay but the random number drawn by the computer, which is always *lower or equal* to your stated number. Therefore, it is in your best interest to state your true willingness to pay.

If you want to increase the probability that your message is transmitted, you should state a high willingness to pay. If it is not so important for you that a player A gets your message, you should state a smaller amount. This will reduce the probability that your message is transmitted.

Example: You write a message and your willingness to pay is 5.20 ECU. Suppose the random number is 4.01. Since your willingness to pay is greater than the random number, your message

will be transmitted to player A and you have to pay the amount of the random number, i.e., 4.01 ECU. That is, your payoff in this round is 11.99 ECU

Your payoff in a given round is:

Willingness to pay \geq random number \rightarrow *Message transmitted*

Payoff = Endowment – random number

Willingness to pay $<$ random number \rightarrow *No Message transmitted*

Payoff = Endowment

Final Earnings

At the end of the experiment, your payoffs in ECU from all rounds are added up and converted into Euro according to the conversion rate stated in the beginning of these instructions. Your final earnings will be rounded to nearest 0.5 Euro.

Are there any remaining questions? If yes, please raise your hands. We will answer your questions in private.

We would like to thank you for participating in this experiment.

Instructions – Coders

Thank you for participating in this study. Through your participation you contribute in significant ways to the scientific analysis of economic decision making. You will receive 15 Euro for your participation today.

If you have questions, please raise your hand. We will come to your place and answer your question in private. Please do not talk to others during this session.

General Overview

Your task will be to evaluate statements and comments from a previous experiment. To give you a better sense of the context in which these statements and comments were made, we will first describe the experiment and then explain your task.

Experiment

There were three different player roles in the experiment (A, B and C), which were randomly assigned in the beginning. The experiment consisted of 10 rounds and in each round one player A, one player B and one player C were randomly matched into a group of three. At the beginning of each round, player A and B received an endowment of 8 points and player C received an endowment of 16 points. The players faced the following task: Player A had to decide how much of their endowment to send to player B. Each sent point was tripled by the experimenter, such that player B received three times the sent amount of player A (in addition to their endowment). Player B had then to decide how much of the available amount to return to player A. The returned amount was not tripled.

Example: Player A sent 8 points. These were triple such that player B received 24 points. Thus the total available amount was 32 points for player B. Player B could therefore return any amount between 0 and 32 points to player A.

Player C observed the decision of player B. That is, player C learned which share of the received amount (return ratio) player B returned to player A.

Examples:

Suppose player B returned 8 points in the example from above. In this case player A would have earned 8 points and player B 24 points. Player C would have seen the following information in this case: 33.3%.

Suppose that Player B returned 12 points. In this case player A would have earned 8 points and player B 20 points. Player C would have learned that the return ratio was 50%.

Suppose that Player B returned 16 points. In this case both players would have earned 16 points. Player C would have learned that the return ratio was 66.6%.

Suppose that Player B returned 0 points. IN this case the return ratio would be 0% (Player C would learn if a player A sent nothing to player B).

Player C could then comment on player B's behavior and write a message to the player A, who was paired next with the observed player B.

Task

Your task is now to evaluate the comments of player Cs. We present you with a list of comments of player Cs along with the corresponding return ratio that player Cs saw when making their comment. Below you will see a table with four different pre-defined categories each consisting of several attributes. For each comment we will show you, you have to assign one or more attributes from each category from below.

Category I	Category II	Category III	Category IV
positive	detailed	evaluates behavior	fair
negative	vague	does not evaluate behavior	generous
neutral	wrong	--	praise
--		--	advice
			selfish, greedy
			warning
			neutral
			wrong, non-sense
			comparing
			not applicable

Please assign only one attribute from each of category I to III, whereas you can assign multiple attributes from category IV.

Example: Suppose you see that the return ratio was 66.7% and that player C wrote the following comment: „B has returned 66.7%, which is a fair distribution.” In this case the following attributes may apply:

Category I: positiv

Category II: detailed (because the message includes explicit information on the return ratio)

Category III: evaluates behavior

Category IV: fair, generous

Important: You should evaluate the different comments from an objective view point and you should try to be consistent. Please take your time when evaluating the messages.

Thank you for your participation.