Will Any Gossip Do? Gossip Does Not Need To Be Perfectly Accurate To Promote Trust*

Miguel A. Fonseca[†], Kim Peters[‡]

December 3, 2017

Abstract

The fact that gossip can be inaccurate, intentionally or otherwise, has led to questions over its ability to build cooperation in large societies. We explore the impact of gossip accuracy on trust and trustworthiness in a population playing decentralized, two-player trust games. We observed non-trivial levels of spontaneous inaccuracy in gossip, and there was evidence that this was largely due to gossipers' desire to punish untrustworthy players. Although this endogenous inaccuracy did not adversely affect levels of trust and trustworthiness, introducing high levels of exogenous inaccuracy did. Importantly though, we observed greater trust and trustworthiness when highly inaccurate gossip was present than when communication was impossible. This suggests that even inaccurate gossip induces a degree of reputational concern in gossip targets and some willingness among gossip recipients to discriminate between partners on the basis of the gossip they received. Thus, gossip need not be perfectly accurate to effectively induce cooperation.

JEL – classification numbers: C93, D03, H41
Keywords: Gossip, Trust, Communication, Experiment.

*Author names are in alphabetical order. We thank the Editor, Associate Editor, and two anonymous reviewers, as well as Bill von Hippel, and participants at seminars at University of Exeter, University of Porto, University of Minho, University of Queensland, and the 2016 ESA meeting at Jerusalem for useful comments and discussions. The authors thank the University of Queensland, as well as ESRC grant ES/N00762X/1 for financial support. We thank Tim Miller for programming the experimental software.

†University of Exeter; Tel: +44 1392 262584; Email: *m.a.fonseca@exeter.ac.uk*.
‡University of Queensland; Tel: +61 73346 9157; Email: *k.peters@uq.edu.au*.

1 Introduction

"Trust is an important lubricant of a social system. It is extremely efficient; it saves a lot of trouble to have a fair degree of reliance on other people's word." Arrow (1974, p. 23)

One of the oldest problems facing organizations and societies has been how to induce good behaviour from their members or citizens, particularly when economic exchange occurs in large and/or geographically dispersed populations. In these circumstances, repeated interactions are rare and it is very costly for any one individual to observe and keep track of the behaviour of others over time. One way one can realign incentives between different parties to economic exchange and induce cooperative behaviour is by disseminating reputational information (Greif, 1993). There is evidence that when such information is available, people will make use of it, discriminating in favor of those with positive reputations. Indeed, the availability of reputational information has been the cornerstone of the success of online marketplaces such as eBay (Dellarocas, 2003; Lucking-Reiley et al., 2007).

However, the development of decentralized mechanisms that disseminate reputation information occurred well before the advent of the Internet. It has been suggested that one such mechanism is gossip, the class of speech that transmits information about the behaviors and attributes of third parties. By some accounts, gossip is said to account for the majority of human conversation; this raises the possibility that people's everyday conversations may be one of the most important contexts for the transmission of reputational information (Dunbar et al. 1997). The sheer ubiquity of gossip suggests that it may impact on economic exchange even in the presence of more formal reputational mechanisms, like those in online marketplaces. Indeed, word-of-mouth is still widely recognised as a powerful driver of new business (Trusov et al., 2009). However, there is (at least) one potential limitation to gossip serving this economic and social function: its inaccuracy. If the reputational information contained in gossip is inaccurate (through error or deliberate distortion), then recipients are more likely to make mistakes like choosing to trust an untrustworthy or unproductive partner. In these situations, people should be less willing to discriminate on the basis of gossip that they receive. Although there is no behavioural evidence for this claim, simulation work has shown that as gossip becomes less accurate, discriminating players are less successful

(Nakamaru and Kawata, 2004; Roberts, 2008; Rauwork et al., 2015).

In this paper, we explore the impact of gossip and its (in)accuracy on trust and trustworthiness in a population of individuals playing decentralized, two-player trust games. We suggest that gossip inaccuracy erodes the cooperative potential of gossip if it interferes with people's tendencies to discriminate on the basis of gossip and their partners' tendencies to anticipate this discrimination (Basu et al., 2009; Boero et al., 2009; Sperber and Baumard, 2009; Sylwester and Roberts, 2013). If gossip provides accurate information about a potential partner's past trustworthiness and competence, recipients are better able to choose to play with good players and avoid playing with bad players. If people receive inaccurate information, they are hampered in their ability to discriminate in this way. At the same time, gossip creates a demand for reputability, as long as targets are aware of the transmission of gossip, and believe that recipients are likely to act on it. If the accuracy of gossip is in question, the targets of gossip may infer that the recipients of the piece of gossip are unlikely to discriminate on the basis of the information it contains. This reduction in reputational concern is likely to reduce the motivation of the targets of gossip to engage in cooperative acts.

In each round of our experiment, we randomly match an Investor with an Agent; the Investor can send any proportion of her endowment to that Agent. Agents can be High-Productivity or Low-Productivity types: any amount sent to the former type is multiplied by six, while any amount sent to the latter is multiplied by three. Investors do not know what type they are facing when deciding how much to send. Agents can then return any amount they wish to the Investor. Before the round ends, each Investor sends a message summarizing the events of the current round: (i) the amount the Agent received; (ii) her type; (iii) the amount the Agent returned to the Investor. At the start of the following round, and prior to making their sending decision, Investors receive a message pertaining to the Agent they are matched with.

In a first set of treatments, we explore whether inaccuracy in messages sent between Investors spontaneously occurs in populations playing our game, and we study what effect such inaccuracy has on levels of trust and trustworthiness. To achieve this, we generated different experimental treatments by varying who Investors sent their messages *to* and who they received them *from*. In the Gossip treatment, Investors sent their messages to the Investor who would play with their Agent in the following round. In the Truth treatment, Investors sent their messages to the experimenter, and received factual summaries of Agent behavior from the experimental software. In the No Message treatment, Investors neither sent nor received messages. By comparing behavior in Gossip to behaviour in No Message, we are able to understand to what extent Investors are willing to discriminate on the basis of the information they receive. Comparing Gossip to Truth allows us to understand if Investors are more likely to lie when communicating to one another than when communicating to the experimenter. This comparison also shows to what extent spontaneous message inaccuracy undermines the ability of Investors to discriminate between Agents and the resulting willingness of Agents to build positive reputations.

In an additional treatment, we study how exogenous message inaccuracy affects trust and trustworthiness. This Inaccurate treatment was a variation on the Gossip treatment. As before, Investors sent their messages to the Investor who would next play with their Agent. However, in this new treatment there was a 50% chance in a given round that an Investor would unknowingly receive a message intended for another Investor. Both Investors and Agents knew of this manipulation, although this was not common knowledge (i.e., Investors were not told that Agents were also aware of the potential for misdirected gossip and vice versa). In this treatment, Investors have weaker incentives to discriminate on the basis of the information they receive; this translates into a weaker incentive for Agents to be trustworthy. This treatment allows us to understand how high levels of inaccuracy affect Investors' reputation-based discrimination and Agents' reputational concern.

The rest of the paper is organized as follows. Section 2 contextualises our paper in the literature on reputation and gossip. Section outlines the research questions and summarizes the experimental design and procedures. Section 4 presents the results and Section 5 offers some concluding remarks.

2 Related Literature

In populations where there is uncertainty about the type of player one is matched with, players may want to discriminate against those who have past histories of uncooperative behavior, and discriminate in favor of those who have reliably cooperated in the past (Wilson, 1985). This means that if reputation-based discrimination is prevalent, strategic reputation building can be a sensible strategy. In these populations, people may find that the immediate costs of cooperation are outweighed by the rewards that their positive reputation delivers over time (Kreps and Wilson, 1982; Milgrom and Roberts, 1982; Fudenberg and Levine, 1992). Indeed, people do appear to be highly sensitive to the possibility that their behaviours may have reputational consequences, as they cooperate more when they believe that they may be observed (Haley and Fessler, 2005; Charness et al., 2011; Barclay, 2012; Yoeli et al., 2013).

When actions are not perfectly observable, it is still possible for people to acquire reputation information through their social networks by means of gossip (Dunbar, 1996, 2004; Foster, 2004; Sommerfeld et al., 2007). Cognitive and evolutionary psychologists have long argued that gossip may be an important mechanism through which people can resolve social dilemmas (Wilson et al., 2000; Dunbar, 2004; Willer et al., 2010). Enquist and Liemar (1993) show that cooperation can be an evolutionary stable strategy if group members can exchange gossip that conveys reputational information. Indeed, the mere possibility that others may gossip about them leads to individuals giving more in a one-shot dictator game (Piazza and Bering, 2008), or contributing more in a one-shot public good game (Beersma and van Kleef, 2011). In an organizational setting, gossip can allow people to maintain and regulate their social networks. For instance, employees can engage in gossip in order to learn about, and influence, the reputations of other organisational members (Emler, 1990, 2001). At an interpersonal level, sharing gossip may also shore up trust between gossipers — especially where sharing this gossip is risky. For instance, employees who share negative gossip about high power individuals within the organisation are opening themselves to the possibility that their audience will report the gossip to the individuals in question (McAndrew and Milenkovic, 2002). Evidence that making oneself vulnerable in this way may facilitate trust and cooperation between gossipers is provided by Kopányi-Peuker et al. (2017) who showed that agents who voluntarily entered into a contract that permitted them to punish one another if they defected achieved higher rates of cooperation than agents who did not.

Recent work has shown that allowing the exchange of gossip in social dilemma games can bolster levels of cooperation. For instance, Abraham et al. (2016) found that if players in a population playing decentralized trust games were able to exchange objective information about their transactions at no cost, they were able to sustain higher trust and trustworthiness. When there was a cost to exchanging this information, however, it was less widely used, and less effective. In the context of public goods games, Wu, Balliet and van Lange (2015) found that allowing players to exchange gossip increased levels of cooperation, and indeed was more effective than punishment at doing so. Finally, Feinberg et al. (2014) showed that the capacity for gossip to shore up cooperation in a public goods game was especially pronounced when players were able to discriminate on the basis of the information that it contained (see also Feinberg et al., 2012). Together, then, this work shows that when people exchange gossip, and use it in order to discriminate among potential partners, these partners are more likely to behave cooperatively.

The accuracy of gossip is central to its function as a catalyst of cooperation. As gossip becomes less accurate, people are less likely to use it to discriminate between partners; in turn, their partners' willingness to engage in strategic reputation building could also diminish. In particular, if people believe that inaccurate gossip reduces reputation-based discrimination then they may be less concerned about the reputational implications of their behaviours and less willing to cooperate (Semmann et al., 2004). To date, there is little work that has examined the effect of inaccuracy in information transmission on cooperation. Most of the work that has done so has employed agent-based simulation models (Nakamura and Kawata, 2004; Rauwolf et al., 2015; Roberts, 2008). Importantly, these studies have shown that increasing the inaccuracy of gossip leads to a decline in cooperation.

The general approach of these studies has been to examine the fitness of cooperators who discriminate on the basis of the gossip they receive relative to agents who do not. While these studies found that discriminating gossipers are highly successful when gossip is perfectly accurate (ensuring high levels of cooperation), this was no longer the case when inaccuracy is introduced into gossip. In particular, inaccurate gossip increased the success of defectors playing a prisoner's dilemma game, especially when this gossip spread widely and rapidly (Nakamura and Kawata, 2004). In a similar way, inaccurate gossip in an indirect reciprocity game increased the relative success of defectors (Rauwolf et al., 2015) or players using a strategy of direct reciprocity (Roberts, 2008).

These simulations also provide evidence that the negative effects of inaccuracy can be seen when rates of inaccuracy are quite low. In particular, Rauwolf et al. (2015) found that cooperation rates declined dramatically if more than 20 percent of messages misreported an agent's previous cooperation as defection, as vice versa. Roberts (2008) found similar patterns when more than 30 percent of messages were replaced with a random value (in a population equally composed of cooperators and defectors, this would mean that 15 percent of messages were wrong). There is also some behavioural support for these findings, as Pfeiffer et al. (2012) found that allowing people who were playing an infinitely repeated prisoner's dilemma game to sell and buy positive reputations between rounds (so reducing the association between past behaviour and reputation) was associated with substantially lower levels of cooperation.

While these studies suggest that even small increases in inaccuracy should lead to reductions in trustworthiness and trust, Nowak et al. (2000) provide a basis for expecting that even inaccurate information may be better than no information at all. In particular, these authors report that allowing reputational information to flow through a network of agents playing the ultimatum game makes it much more likely that players will settle on the fair strategy of making and accepting high offers than the rational strategy of making and accepting low offers. The beneficial impact of reputational information was evident even when when a small fraction of players (i.e., 10-20 percent) in the population learned about the offer that a given agent had previously accepted.Although this study does not examine inaccurate information, it does suggest that (at least in some circumstances), inaccurate gossip may be better than no information at all.

To the best of our knowledge, there is almost no empirical work to date looking at the impact of gossip inaccuracy on behaviour. Furthermore, no paper to date has disentangled the effect that inaccurate gossip has on the behaviour of those who receive the gossip and those being gossiped about. The closest related work to the present paper is that by Fehr and Sutter (2016), who independently of our work, look at the effect of observability by a third party in a similar vein to that proposed by Feinberg et al. (2012), as well as the possibility of that third party engaging in gossip as a non-monetary form of punishment in populations playing trust games. Fehr and Sutter (2016) find that mere observability by a third party increases trust and trustworthiness. The ability of that third party to send a free form message about the events of the game increases trustworthiness further, even when the message is noisy.

Our study contributes to this literature by exploring how inaccuracy can impact discrimination and reputational concern. These are intrinsically interlinked processes that are both implicated in the effectiveness of gossip (and reputation systems more generally). That is, discriminating in favour of cooperative agents increases the likelihood of future cooperation (at least when similar discrimination is anticipated). At the same time, however, in most everyday circumstances, these processes are likely to be loosely coupled, as agents may assume that levels of discrimination are higher (or lower) than they actually are, and consequently exhibit more (or less) trustworthiness than is warranted. By examining discrimination and reputation building in the presence of endogenous and exogenous inaccuracy, our study will inform an understanding of the way in which these processes interact to determine the efficacy of gossip.

3 Experimental Design and Hypotheses

In this pre-registered experiment (https://tinyurl.com/hc27ajw), we implemented a modified version of the trust game (Berg et al., 1995) in networks of 12 people. Participants played the game on networked computers that were separated by partitions to ensure anonymity. At the start of each session, equal numbers of participants were randomly assigned to be Investors (N=6) or Agents (High-Value N=3; Low-Value N=3), a role they retained for the duration of the experiment. We used the neutral terminology 'Sender' and 'Receiver' in the instructions and experimental software when referring to 'Investors' and 'Agents'. Please see the Appendix for copies of the instructions.

The different types of Agents capture the idea that there is a distribution of productivity in the economy, which manifests itself in the surplus that is generated when an Investor enters in an economic relationship with an Agent. By manipulating Agent types we were able to examine the different kinds of reputational information exchanged by people in everyday life. In particular, while previous work has focused exclusively on reputational information about a person's morality (i.e., trustworthiness or cooperative behaviour), real social networks also convey reputational information about a person's competence (Peters and Kashima, 2015).

In each of the 15 experimental rounds, an Investor was randomly paired with an Agent with the exception that no two individuals could be paired twice in a row (over 80 per cent of Investors faced High-Value Agents on at least 5, but fewer than 11 rounds). At the start of each round, Investors received an endowment of 10 Experimental Currency Units (ECU), and could send any amount (including zero) to their Agent. To proxy the role ability plays in generating economic benefits, any amount sent to a High-Value Agent was multiplied by 6, while any amount sent to a Low-Value Agent was multiplied by 3. On receiving their ECU, Agents could return any amount, including zero, to their Investor. Each player's behaviour, the Agent's type and their respective payoffs for that round were then summarized onscreen, bringing the round to a close.

The experiment consisted of four treatments. In our baseline treatment, No Message, there was no opportunity for Investors to communicate with each other. In all other treatments, we allowed reputational information to flow within networks. In these reputational treatments, at the end of each round Investors were asked to construct a message that told the message recipient what they wanted him or her to know about the number of ECU that they had sent to their Agent, the number their Agent received, the number their Agent returned and their Agent's type. Messages about ECU sent, received and returned were numerical in nature: subjects had to type the relevant number in a text box. Messages about the Agent's type were selected from a set menu with two options (High- or Low-Value). To preclude memory effects, the actual behaviour in that round was summarized onscreen for reference. From the second round onwards, Investors in the reputation treatments received a message that described their new Agent's most recent behaviour.

We did not assign any label or identifying information to Investors and/or Agents. It was therefore impossible for Investors to track the behaviour of different Agents over time. Likewise, Agents could not track the past behavior of different Investors. This also made it impossible for Investors to track the predictive power of messages from different Investors over time. We opted for this environment in order to have the cleanest possible measure of what the impact of communication between Investors is on cooperation in a population. It is impossible for reputational information to be acquired in any form other than through gossip.

[insert Table 1 here]

We generated different experimental treatments by varying who Investors sent their messages *to* and who they received them *from*, outlined in Table 1. We ran an initial set of three treatments, and subsequently collected an additional treatment in which gossip was exogenously made inaccurate. The initial set consisted of a treatment with information transmission between Investors, Gossip; a direct observation treatment, Truth; and a no reputation treatment, No Message. In the Gossip treatment, Investors sent their messages to the Investor who would play with their Agent in the following round. In the Truth treatment, Investors sent their messages to the experimenter, and received factual summaries of Agent behaviour from the experimental software. In the No Message treatment, Investors neither sent nor received messages.

In a one-shot interaction, self-interested Agents will keep all the ECU they receive. Investors who anticipate this response should not send any ECU to their Agents. If a population of randomly matched players repeatedly plays this game, and if reputational information is available, discrimination and reputation building may come into play, making it easier for Investors to trust their Agents. If gossip supports these reputational processes, it should increase social welfare, because the number of ECU that Investors send to Agents determines joint profits.

Hypothesis 1: Levels of trust and trustworthiness will be higher in Truth and Gossip than in No Message.

It is noteworthy that in our game, there is no financial incentive for an Investor to send truthful information about an Agent to the next Investor who will play with that Agent. While there is no financial incentive associated with Investors sending inaccurate messages either, they are free to do so if they want to. We anticipate that Investors will be more likely to send inaccurate messages in Gossip, where these messages are passed on to other players, than in Truth, where they are passed on to the Experimenters. If Investors are inaccurate in the messages that they send to one another, or believe that such inaccuracy is possible, then this should reduce their willingness to discriminate on the basis of gossip. This should, in turn, reduce the reputational concern and trustworthiness of Agents. Investors as a result should be less trusting.

Hypothesis 2: Investors will send more inaccurate messages in Gossip than in Truth.

Hypothesis 3: Levels of trust and trustworthiness in will be higher in Truth than in Gossip.

We complemented this initial set of treatments with a fourth treatment, Inaccurate. As in Gossip, Investors sent their messages to the Investor who would next play with their Agent. However, in this new treatment there was a 50% chance in a given round that an Investor would unknowingly receive a message intended for another Investor. Both Investors and Agents knew of this manipulation although this was not common knowledge (i.e., Investors were not told that Agents were aware of the potential for misdirected gossip and vice versa).

The built-in inaccuracy in the Inaccurate treatment should compound any spontaneous inaccuracy inherent to communication in the Gossip treatment, and have a dramatic effect on the reliability of the information being transmitted between Investors. This should have a knock-on effect on trust and trustworthiness.

Hypothesis 4: Levels of trust and trustworthiness will be higher in Gossip than in Inaccurate.

We recruited 288 participants in March and October 2015 from a pool of volunteer participants through ORSEE (Greiner, 2015). We conducted a total of 24 sessions, six for each of the four treatments. A total of 72 participants took part in each treatment. Each participant took part in only one treatment, and no participant had taken part in any trust experiments previously run in the same lab. Upon arrival to the laboratory, participants sat in an individual booth, which had partitions to minimize visual contact between neighboring participants. The instructions requested that participants maintained silence at all times. All interactions were done through computer terminals, using the z-Tree software (Fischbacher, 1997). The experiment was run in the FEELE laboratory at the University of Exeter. The average duration of the experiment was one hour. The experimental payment was calculated on the basis of the show-up fee plus their payoff from one randomly selected experimental round. The average payment was approximately £11.90 (US$18.90).

4 Results

4.1 Message (In)accuracy

There was a considerable degree of spontaneous inaccuracy in the messages that were sent by Investors in the reputation treatments: 53 per cent of Investors sent one or more inaccurate messages, such that about 20 per cent of all messages had at least one inaccurate component. Notably, this approaches the levels of inaccuracy that have been shown to undermine the cooperative potential of gossip in simulation studies (Roberts, 2008; Rauwolf et al., 2015). This inaccuracy was more often seen in reports of an Agent's trustworthiness (ECU returned: 15 per cent of messages) than his or her type (5 per cent) or Investors' self-reported trust (ECU sent: 8 per cent of messages). This inaccuracy was also significantly more prevalent when Investors were writing to one another (Gossip and Inaccurate treatments: 11 to 12 per cent of the three message components) than when they were writing to the experimenter (Truth treatment: 4 per cent of the three message components). The fact that rates of inaccuracy were about three times higher when communicating with other Investors (vs. the Experimenter), suggests that experimental demand effects and/or random imprecision are unlikely to be their only —or indeed the most important— determining factors.

Evidence that the observed inaccuracy was driven by strategic considerations is provided by an examination of Investors' responses to an open question in the post-session questionnaire that asked how they had arrived at their experimental decisions. In their responses, 49 per cent of Investors in the reputation treatments made some reference to the accuracy of messages. In the Truth treatment, 70 per cent of such Investors stated that they always told the truth and another 23 per cent said that they lied at least occasionally. In the Gossip treatment, the pattern was reversed, with 32 per cent of such Investors saying that they always told the truth and 63 per cent admitting having lied. In the Inaccurate treatment, there was a different pattern again, with 85 per cent of these Investors commenting instead on the inaccuracy of the messages they received.

[insert Figure 1 here]

Importantly, Investors who admitted lying always specified how this helped them to achieve one or more social ends. In particular, 53 per cent said they exaggerated negatives to punish untrustworthy Agents, 26 per cent said they exaggerated positives to reward trustworthy Agents, 16 per cent said they misrepresented Agent behaviour to compete with other Investors and 21 per cent said they put a positive spin on all behaviour to encourage other Investors to send more — perhaps indicating efficiency concerns (Engelmann and Strobel, 2004). While the contradictory form of these different lies will tend to obscure general patterns of inaccuracy in the data, there is some evidence for the punishment motive in the data. Figure 1 plots the relative frequency of lies as a function of the trustworthiness of the Agent (i.e. the amount returned by the Agent as a proportion of what she received) in the Gossip and Inaccurate treatments. Negative lies (i.e. messages portraying the Agent as less trustworthy than reality) account for 31% of all messages when the Agent returned no more than 10% of what she received. Thereafter, the relative frequency of negative lies is between 10-15%, a significant decrease (all comparisons $\chi^2(1) \geq 4.89, p \leq 0.027$). In contrast, the proportion of positive lies (i.e. messages that portray the Agent as more trustworthy than reality) do not vary with the Agent's actual trustworthiness.

Observation 1: There was a substantial amount of inaccuracy in messages sent by Investors; the primary motivation for message inaccuracy was the punishment of misbehaving Agents.

In short, these findings are consistent with claims that inaccuracy is intrinsic to gossip; they also suggest that this is due, at least in part, to gossipers' perceptions that lying can serve a range of social functions. We will now examine the implications of this inaccuracy for the cooperative functions of gossip.

4.2 Spontaneous Inaccuracy and Cooperation

[insert Figure 2 here]

To see whether spontaneous inaccuracy eroded the ability of reputational information to shore up cooperation, we compared Investor trust behaviour in the Gossip, Truth and No

Message treatments, pooling the data in five-round blocks (all our findings replicate if we analyse behaviour round-by-round; see Table C3 in the Appendix). Trust levels are presented in Figure 2.[1] Although there were no differences in Investors' trust levels across the three treatments in the first third of the experiment (all comparisons $F(1, 29) \leq 2.33, p \geq 0.140$), differences emerged over time. By the last third of the experiment, levels of trust were significantly higher in the Truth and Gossip treatments than they were in the No Message treatment (all comparisons $F(1, 35) \geq 9.24, p \leq 0.006$). Unexpectedly, there was no evidence that the ability of reputational information to boost trust depended on whether it was conveyed through gossip or direct observation (all comparisons $F(1, 35) \leq 0.07, p \geq 0.800$). This suggests that reputational information can be less than perfectly accurate and still function very well.

Observation 2: The ability to transmit reputational information about Agents leads to higher levels of trust, even when messages may be inaccurate.

To understand why inaccuracy in the Gossip treatment did not interfere with its ability to shore up cooperation, we looked at the way in which it affected key reputational processes. We first examined Investors' tendencies to discriminate between Agents on the basis of the messages that they received in the Truth and Gossip treatments. The average amount of trust shown by Investors as a function of gossip about the Agent's trustworthiness and Type is shown in Figures 3 and 4. As expected, Investors discriminate between Agents as a function of the Agent's trustworthiness (sending more points to reputedly more trustworthy Agents), and the tendency to discriminate on this basis appears to be higher in the Truth than the Gossip treatment. Investors in Truth and Gossip also discriminated positively towards reputedly High Type Agents; interestingly, this effect appears to be more pronounced in Gossip.

[insert Figures 3 and 4 here]

Table 2 presents the output of the three-level random effects maximum likelihood

[1] We present summary statistics results in graphical form for ease of exposition. The regression results on which Figures 2, 3, 4, 5 and 6 are based can be found in Appendix Tables C4, C8, C7, C4 and C6 respectively.

estimation regressing ECU sent by Investor j in round t in session k on the content of the message received by that Investor (Agent type and proportion ECU returned, which we calculated by dividing ECU returned by ECU received), the interaction between Agent type and proportion ECU returned, as well as a linear time trend and its interactions with the message variables. All these variables were interacted with treatment dummies.

[insert Table 2 here]

The first two columns in Table 2 show that in the Truth and Gossip treatments Investors trusted Agents more if the message said that the Agent had previously been more trustworthy (i.e., returned a higher proportion of received ECU). Importantly though, this tendency to discriminate on the basis of previous trustworthiness was about two and a half times higher in the Truth treatment than it was in the Gossip treatment ($\chi^2(1) = 5.06, p = 0.025$). This is consistent with the possibility that Investors were less willing to discriminate between Agents on the basis of potentially inaccurate reputational information. At the same time, the significant positive interaction between Agent trustworthiness and experimental round in the Gossip treatment ($\chi^2(1) = 6.93, p = 0.009$) suggests that over time Investors learned to rely on these messages.

Interestingly, the analysis represented in Table 2 reveals that Investors in the Gossip treatment were sensitive to their Agent's Type, as they sent significantly more points to High-Type Agents. However, there is no statistically significant difference in the differential treatment of different types of Agents when we compared Truth to Gossip ($\chi^2(1) = 0.19, p = 0.666$); the lack of statistical significance is presumably due to the large variance in the estimate for the Truth condition. There was no evidence that Investors in the Gossip treatment learned that Type was more likely to be accurate, as the interaction with experimental round was not significant. In addition, there was no evidence of a similar increase in sensitivity in the Inaccurate treatment.

Observation 3: Message inaccuracy leads to Investors discriminating less on the basis of gossip they receive about Agent trustworthiness.

To understand whether the unconditional trust shown by Investors in the Gossip treatment was a reflection of the trustworthiness of their Agents, we compared the proportion of ECU returned by Agents in the three conditions (Figure 5). By this measure, Agents in the Truth and Gossip treatments were about twice as trustworthy as those in the No Message treatment in all three five-round blocks (all comparisons $F(1, 35) \geq 18.22, p \leq 0.001$). Therefore, the transmission of reputational information appeared to motivate strategic reputation building. Importantly though, there was no evidence that Agents in the Gossip treatment were any less trustworthy than those in the Truth treatment at any stage in the experiment (all comparisons $F(1, 35) \leq 1.06, p \geq 0.314$). This means that Agents in the Gossip treatment might not have been able to anticipate their Investors' weaker discrimination at the start of the experiment, nor to take advantage of it as the experiment progressed (at least over the time frame examined here).

Observation 4: Trustworthiness rates are higher in Gossip and Truth than No Message, which indicates that transmission of reputational information appeared to motivate strategic reputation building.

[insert Figure 5 here]

Our data suggest that inaccurate gossip does a better job of shoring up trust in behavioural settings than simulation work implies. We are also able to explain why: Investors were sensitive to the manner in which reputational information was conveyed, as their rate of trust was less responsive to information about past trustworthiness in Gossip than in Truth. However, Agents were as trustworthy in Gossip as in Truth; this may have led Investors to be more trusting unconditionally in the former than in the latter.

To see whether Agents can better calibrate their trustworthiness to Investors' actual discrimination when levels of observed inaccuracy were higher, in the next section we examine behaviour in the Inaccurate treatment, where there was a 50 per cent chance of a message being misdirected.

4.3 Misdirected Messages and Cooperation

The messages that Investors received in the Inaccurate treatment were highly inaccurate, as they only accounted for, on average, 7 to 12 per cent of variance in the behaviors that they were supposed to describe. For instance, the behaviour-to-message correlation for ECU sent was $r = 0.35$ in the Inaccurate treatment; the same association was $r = 0.80$ in the Gossip treatment.[2] Further analysis (see Table **??** in Appendix) showed that the information that Investors received was significantly less accurate in the Inaccurate treatment than in the Gossip treatment across all message components (all comparisons, $\chi^2(1) \geq 69.50, p < .001$).

To see whether this increased inaccuracy reduced the effectiveness of gossip, we compared Investor trust in the Inaccurate treatment with the Gossip and No Message treatments (Figure 2). There was no evidence of treatment differences in the first third of the experiment (all $F(1, 29) \leq 1.55, p \geq 0.226$). However, from this point onwards, levels of trust in the Inaccurate treatment dropped below those in the Gossip treatment (all comparisons $F(1, 29) \geq 5.20, p \leq 0.032$) and could not be distinguished from those in the No Message treatment (all comparisons $F(1, 29) \leq 2.71, p \geq 0.113$). Together, these findings lead to our next observation.

Observation 5: Exogenously increasing the inaccuracy of gossip had a deleterious effect on trust.

To better understand these findings, we return to the key reputation processes. Figure 4 shows that Investors seem to be less willing to discriminate between Agents as a function of the Agent's trustworthiness in the Inaccurate treatment, Figure 3 shows that Investors in Inaccurate do not appear to discriminate between Agents on the basis of their reputed Type. The last column of Table 2 shows that Investors in the Inaccurate treatment discriminated in favour of Agents who were reputed to be more trustworthy. At intercept, Investors in the Inaccurate treatment were no less willing to discriminate on the basis of Agent trustworthiness than were those in the Gossip treatment ($\chi^2(1) \leq 0.04, p \geq 0.841$). Importantly however, unlike Investors in the Gossip treatment, there was no evidence that Investors in the Inaccurate treatment increased their reliance on these messages over time ($\chi^2(1) = 5.01, p = 0.025$).

Observation 6: Exogenously increasing the inaccuracy of gossip leads to Investors (i)

[2]For a full set of correlations between behavior and messages, see Table **??** in the Appendix

reducing their levels of trust and (ii) maintaining their reduced reliance on gossip information over time, relative to the Truth and Gossip conditions.

Next, we compared the trustworthiness of Agents in the Inaccurate, Gossip and No Message treatments (Figure 5). Agents in the Inaccurate treatment appeared to anticipate diminished discrimination: in the first third of the experiment they returned significantly less than Agents in the Gossip treatment ($F(1, 29) = 28.18, p < 0.001$). This relatively low level of trustworthiness was retained through the remainder of the experiment ($F(1, 29) \geq 8.23, p \leq 0.009$). At the same time, however, Agents in the Inaccurate treatment were consistently more trustworthy than those in the No Message treatment (all comparisons $F(1, 29) \geq 5.38, p \leq 0.030$).

Observation 7: Exogenously increasing inaccuracy in gossip diminishes strategic reputation building and trustworthiness.

4.4 Investor Responsiveness to Messages

We next look at how Investors reacted to the messages they received at the start of each round, and the resulting payoff consequences. Tables B1, B2, and B3 in the Appendix display the average profit that each available action (i.e. ECU sent) yielded an Investor conditional on the message they received from another Investor concerning the trustworthiness of the Agent they were about to play with. In order to make the analysis tractable, we constructed a composite message that condensed the three pieces of information being transmitted (ECU sent, Agent type and ECU returned) into one metric: the percentage amount returned by the Agent, which we define as the ratio of ECU returned to ECU received (i.e. post-multiplication, and thus controlling for Agent type). We conditioned the profitability of different actions on five sets of messages: 0; (0, 0.2]; (0.2, 0.4]; (0.4; 0.6]; (0.6, 1].[3] We took this approach in order to overcome two obstacles: the dimensionality of message space, and sample size. Constructing the composite message and pooling messages in five sets allowed us to have enough observations across outcomes to make meaningful statistical inference about what action was most profitable to Investors for a given message. We focus our analysis on

[3]We excluded three observations where Investors reported return rates greater than 1, which are impossible.

instances with at least six observations, and we test for significant differences in profits using non-parametric statistics.[4]

We start by focusing on the cases where the message either reported the Agent not having returned anything (i.e. $\{0\}$) or less than 20% of the amount he or she received. In all three treatments, the average profit does not noticeably change with ECU sent, and there are no instances in any of the three treatments where Investors made higher profits by sending positive amounts than by sending 0 ECU. When the message sent to the Investor reported that the Agent returned more than 20% of the amount received, there is a clear difference between the Innacurate treatment and the Truth and Gossip treatments. In the latter treatments, average profit dramatically increases with ECU sent, and the most profitable action is 9 ECU or 10 ECU (between 80% and 100% more profitable than sending 0 ECU). In contrast, in the Inaccurate treatment, the increase in average profit is much more moderate; the most profitable action yielded 50% higher payoffs than sending nothing.

As a consequence, the distribution of ECU sent conditional on the message type in Inaccurate is much more dispersed than in Truth and Gossip. In latter two treatments, there is a clear modal amount sent conditional on message: 0 for low reported trustworthiness and 10 for high reported trustworthiness. In Inaccurate, this is only the case when the message reports a % return in the 0.4-1 range. This is corroborated by Figure 4, which plots average ECU sent conditional on the message about the percentage amount returned by Agents. The average amount sent conditional on a message of no-trustworthiness is actually higher in Inaccurate than in Gossip ($F(1,16) = 5.60, p = 0.030$) or Truth ($F(1,16) = 11.64, p = 0.003$), which reflects the higher dispersion in ECU sent conditional on a message of $\{0\}$ (see tables B1, B2, and B3). For messages reporting low trustworthiness, $(0, 0.2]$ and $(0.2, 0.4]$, we find no difference in average ECU sent when comparing Inaccurate to Gossip ($F(1,16) = 0.53, p = 0.478$; $F(1,16) = 0.79, p = 0.387$) or Truth ($F(1,16) = 0.34, p = 0.570$; $F(1,16) = 0.19, p = 0.672$). However, for messages conveying high trustworthiness $(0.4, 1]$, the average ECU sent in Inaccurate was significantly smaller than that in Gossip ($F(1,16) = 3.54, p = 0.0.077$) and Truth ($F(1,16) = 10.24, p = 0.005$).

[4] We nevertheless had to pool observations across sessions, as well as ignoring the time dimension in the experiment.

Observation 8: In both Gossip and Truth treatments, responding to (low) high reported trustworthiness with (low) high trust yields significantly higher payoffs than any other strategy. This association is substantially less evident in the Inaccurate treatment.

4.5 Efficiency

We conclude by looking at the impact gossip had on efficiency, defined as the joint payoffs for players in the population. The left panel in Figure 6 displays the average total payoff in each treatment, as well as the share of Investors and Agents. Allowing for information transmission leads to a significant increase in efficiency, irrespective of the origin or reliability of the message, relative to the case where no information is transmitted (Inaccurate = No Message, $F(1, 23) = 3.18, p = 0.088$; all other comparisons with No Message, $F(1, 23) \geq 8.08, p < 0.001$). Average total payoffs are lower in Inaccurate than Gossip ($F(1, 23) = 8.66, p = 0.007$) but not significantly different than in the Truth condition ($F(1, 23) = 2.48, p = 0.129$).

The right panel in Figure 6 breaks down the differences in average payoffs by role (i.e. Investor, Low-Type Agent and High-Type Agent). It is immediately apparent that the bulk of the gains in efficiency in Gossip and Truth relative to No Message are borne by Investors (both comparisons, $F(1, 23) \geq 22.98, p \leq 0.001$), as well as High Value Agents (both comparisons, $F(1, 23) \geq 4.39, p \leq 0.047$), who turn out to be the beneficiaries from the Investors' ability to discriminate on the basis of reputation. In contrast, Low-Value Agents do not gain from the existence of gossip between Investors relative to the No Message baseline ($F(1, 29) = 2.71, p = 0.114$). Low-Value Agents do benefit from exogenous noise: they earn significantly more in the Inaccurate treatment than in Gossip ($F(1, 29) = 4.46, p = 0.046$).

[insert Figure 6 here]

Observation 9: Gossip leads to increased efficiency, primarily by increasing average Investor and High-Value Agent payoffs. Exogenous inaccuracy in messages benefits Low-Value Agents.

5 Conclusion

It has been suggested that inaccuracy poses a major threat to the capacity of gossip to boost cooperation in society. Our results suggest that while inaccuracy is intrinsic to gossip —most Investors misrepresented their Agent's behaviour at least once— and perceived to serve a range of strategic purposes, it may be less of a threat to its functionality than previous simulation work suggests. Even though spontaneous inaccuracy in messages was widespread, meaning that gossip provided a weak reflection of Agent behaviour, rates of trust and trustworthiness were higher than when there was no transmission of reputational information at all. As a result, even inaccurate gossip had a positive impact on welfare: total earnings were higher when gossip was allowed than when it was not. Investors were the main beneficiaries, with an increase in earnings of up to 50% relative to the No Message condition.

Our finding that inaccurate gossip is less effective than accurate gossip, but more effective than no gossip is consistent with that reported by Fehr and Sutter (2016). The fact that this consistency is observed in the context of non-trivial methodological differences—including levels of inaccuracy, gossiper role and message forms—points to the likely robustness of the findings. Specifically, across the two papers, it seems that inaccurate gossip is more effective than no gossip whether it is misdirected on 20% or 50% of occasions; the capacity of gossip to increase efficiency does not appear to depend on whether the gossiper merely observed the described behaviour or was materially affected by it; and the impact of gossip does not appear to depend on whether the gossip is strictly numerical or allows free-form verbal messages.

Our analysis further suggests that the efficiency gains that are associated with inaccurate gossip can be attributed to the fact that it enables its recipients to discriminate on the basis of the social information they receive and creates a need for reputability among its targets. Importantly, this is the case even when extreme levels of inaccuracy reduce levels of discrimination; this suggests that the capacity for gossip to shore up cooperation is likely to depend as much on the reputational concern it generates as on people's tendencies to act on it. As long as people underestimate the way in which inaccuracy diminishes the discriminatory behavior of gossipers, then the mere existence of gossip (however accurate) should suffice.

An additional plausible explanation for the effectiveness of gossip in the Inaccurate treatment is that while gossip may not have been informative about the individual with whom Investors were about to play, it still provided valuable information about the trustworthiness of the population of Agents as a whole. Although this reflects a particularity of the current experimental design, it does have real life analogies. For instance, in a situation of mistaken identity, an individual could receive information about an actual behaviour but attribute it to the wrong person. While this could lead to discrimination mistakes, it could also improve this individual's understanding of the local descriptive norms (e.g., people can be highly (un)trustworthy), allowing them to adapt their behavior accordingly.

It is also important to note that in the present experimental setup, Agents could be excluded from future transactions at a relatively low cost to Investors. Specifically, Investors could avoid interacting with an Agent by sending them zero ECU (thus keeping their endowment) if they wished to. It is possible that inaccurate gossip may be less effective if Investors cannot avoid interacting with Agents (e.g., Feinberg, Willer, and Schultz, 2014).

Our study suggests that online platforms may benefit from imperfect reputation systems, at least as long as this system serves to increase levels of reputational concern (and consequently boost levels of cooperation). There is, however, an important difference between the inaccuracy that we introduced in the lab and that accompanying most online reputation systems. While the gossip target did not observe specific instances of inaccuracy in our study (and hence was unable to challenge them), in many online marketplaces feedback is visible to both parties and can be countered. It is an open question as to the impact that these differences would have for the generalizability of the findings that we describe in our study; we believe that this question is worthy of future research attention.

References

[1] Abraham, M., Grimm, V., Neess, C. and Seebauer, M., 2016. Reputation formation in economic transactions. Journal of Economic Behavior and Organization 121, 1-14.

[2] Arrow, K.J., 1974. The Limits of Organization. New York: W. W. Norton.

[3] Barclay, P., 2012. Harnessing the power of reputation: strengths and limits for promoting cooperative behaviors. Evolutionary Psychology 10, 868-883.

[4] De Backer, C.J.S., and Gurven, M., 2006. Whispering down the lane: the economics of vicarious information transfer. Adaptive Behavior 14, 249-264.

[5] Beersma, B., and Van Kleef, G.A., 2011. How the grapevine keeps you in line: gossip increases contributions to the group. Social Psychology and Personality Science 2, 642-649.

[6] Bolton, G., Greiner, B., and Ockenfels, A., 2012. Engineering trust: Reciprocity in the production of reputation information. Management Science 59(2), 265-285.

[7] Charness G, Du N., and Yang C-L., 2011. Trust and trustworthiness reputations in an investment game. Games and Economic Behavior 72, 361-375.

[8] Dellarocas, C., 2003. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. Management Science 49(10), 1407-1424.

[9] Dellarocas, C., and Wood, C.A., 2007. The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias. Management Science 54(3), 460-476.

[10] Dunbar, R.I.M., Marriott, A., and Duncan, N.D.C., 1997. Human conversational behavior. Human Nature 8(3), 231-246.

[11] Emler, N., 1992. The truth about gossip. Social Psychology Newsletter 27, 23-37.

[12] Emler, N., 2001. Gossiping. The New Handbook of Language and Social Psychology 317-338.

[13] Engelmann, D., and Strobel, M., 2004. Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. American Economic Review 94(4), 857-869.

[14] Fehr, D., and Sutter, M., 2016. Gossip and the efficiency of interactions. IZA DP No. 9704.

[15] Feinberg M., Cheng, J.T., and Willer R., 2012. Gossip as an effective and low cost form of punishment. Behavioral and Brain Sciences 35(1), 25.

[16] Feinberg M., Willer R., Stellar J., and Keltner, D., 2012. The virtues of gossip: reputational information sharing as prosocial behavior. Journal of Personality and Social Psychology 102, 1015-1030.

[17] Fischbacher, U., 2007. z-Tree - Zurich toolbox for Readymade Economic Experiments. Experimental Economics 10(2), 171-178.

[18] Greif, A., 1993. Contract enforceability and economic institutions in early trade: The Maghribi traders' coalition. American Economic Review 83(3), 525-548.

[19] Greiner, B., 2015 Subject pool recruitment procedures: Organizing experiments with ORSEE. Journal of the Economic Science Association 1(1), 114-125.

[20] Haley K.J., and Fessler D.M.T., 2005. Nobody's watching? Subtle cues affect generosity in an anonymous economic game. Evolution and Human Behavior 26, 245-256.

[21] Kopányi-Peuker, A., Offerman, T., and Sloof, R., 2017. Fostering cooperation through the enhancement of own vulnerability. Games and Economic Behavior 101, 273-290.

[22] Lucking-Reiley, D., Bryan, D., Prasad, N. and Reeves, D., 2007. Pennies from Ebay: The determinants of price in online auctions. Journal of Industrial Economics 55(2), 223-233.

[23] McAndrew, F. T., and Milenkovic, M. A., 2002. Of tabloids and family secrets: The evolutionary psychology of gossip. Journal of Applied Social Psychology, 32(5), 1064-1082.

[24] Nosko, C., and Tadelis, S., 2015. The limits of reputation in online markets: An empirical analysis and field experiment. Mimeo.

[25] Nowak, M.A., Page, K.M., and Sigmund, K., 2000. Fairness versus reason in the ultimatum game. Science 289(5485), 173-1775.

[26] Peters K., and Kashima, Y., 2015. Bad habit or social good? How perceptions of gossiper morality are related to gossip content. European Journal of Social Psychology 45:784-798.

[27] Pfeiffer T., Tran L., Krumme C., and Rand, D.G., 2012. The value of reputation. Journal of the Royal Society Interface 9: 2791-2797.

[28] Piazza, J., Bering, J.M., 2008. Concerns about reputation via gossip promote generous allocations in an economic game. Evolution and Human Behavior 29, 172-178.

[29] Semmann, D., Krambeck, H-J., and Milinski, M., 2004. Strategic investment in reputation. Behavioral Ecology and Sociobiology 56, 248-252.

[30] Sommerfeld, R. D., 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 of the U.S.A., 104(44), 17435-17440.

[31] Sommerfeld, R.D., Krambeck, H-J., and Milinski, M., 2008. Multiple gossip statements and their effect on reputation and trustworthiness. Proceedings of the Royal Society B 275, 2529-2536.

[32] Trusov, M., Bucklin, R.E., and Pauwels, K., 2009. Effects of word-of-mouth versus traditional marketing: Findings from an internet social Nnetworking site. Journal of Marketing 73, 90-102.

[33] Wu, J., Balliet, D., and van Lange, P.A.M., 2014. When does gossip promote generosity? Indirect reciprocity under the shadow of the future. Social Psychological and Personality Science 6(8), 923-930.

[34] Yoeli, E., Hoffman, M., Rand, D.G., and Nowak, M., 2013. Powering up with indirect reciprocity in a large-scale field experiment. Proceedings of the National Academy of the U.S.A. 110, 10424-10429.

Tables

Treatment	No Message	Truth	Gossip	Inaccurate
# sessions, # subjects	6, 72	6, 72	6, 72	6, 72

Table 1: Experimental Design

DV: ECU Sent	Truth	Gossip	Inaccurate
Message High-Value Type	0.531	1.743***	-0.225
	(0.716)	(0.666)	(0.606)
Message Proportion Returned	8.150***	2.769*	3.179***
	(1.726)	(1.656)	(1.190)
Message High-Value Type ×	1.544	1.675	2.311**
Proportion Returned	(1.451)	(1.230)	(0.934)
Experimental Round	0.020	0.014	0.023
	(0.073)	(0.069)	(0.057)
Message High-Value Type ×	0.002	-0.034	-0.024
Experimental Round	(0.058)	(0.058)	(0.057)
Message Proportion Returned ×	0.317*	0.448***	-0.012
Experimental Round	(0.177)	(0.170)	(0.115)
Constant	2.246***	3.097***	4.512***
	(0.751)	(0.733)	(0.635)
Random Effects Parameters			
Session parameter: (< 0.001)			
Investor, intercept parameter: (1.649)			
Investor, residual: (2.431)			
N (obs, Investors)= (1,394, 108); Wald $\chi^2(20) = 677.00, p < 0.0001$.			

*** : $p < 0.01$; ** : $p < 0.05$; * : $p < 0.10$. Coefficients refer to interactions of IVs in the first column with treatment dummy IVs in first row.

Table 2: Regression estimates of message determinants of Investor trust.

Figures

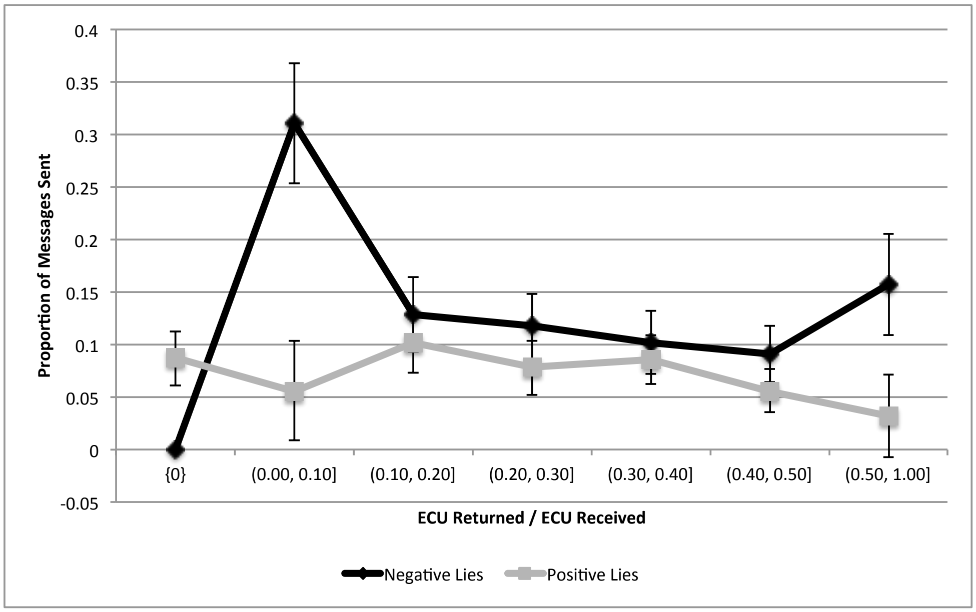

Figure 1: Proportion of negative and positive lies about trustworthiness as a function of actual trustworthiness. Error lines represent robust standard errors clustered at the network level around the mean.

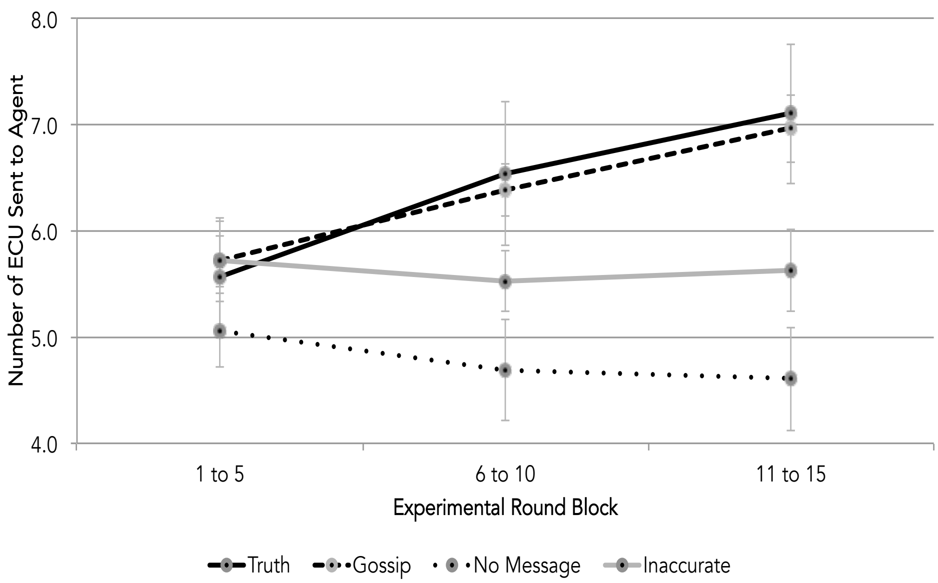

Figure 2: Trust as measured by ECU sent to Agents. Error lines represent robust standard errors clustered at the network level around the mean.

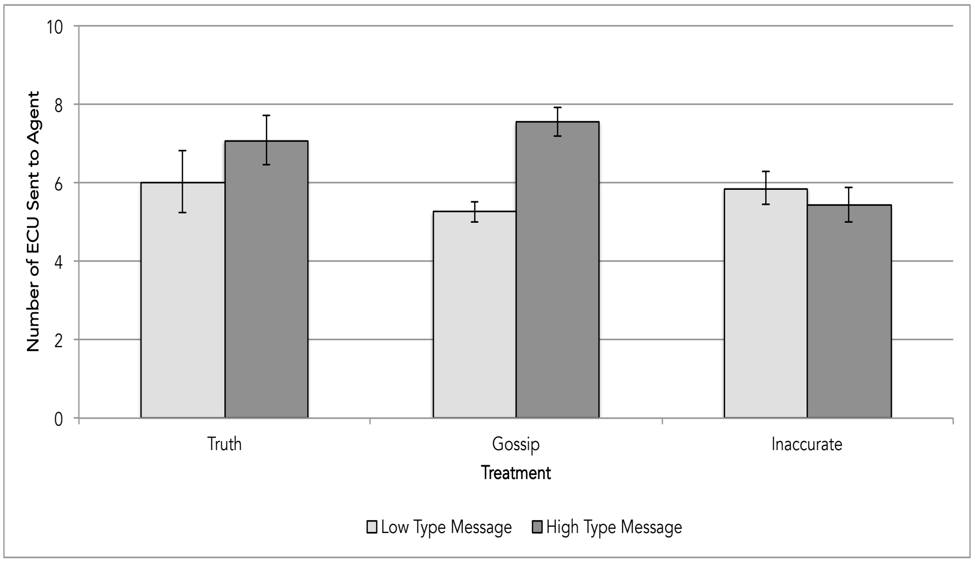

Figure 3: Trust as measured by ECU sent to Agents as a function of reported Agent type. Error lines represent robust standard errors clustered at the network level around the mean.

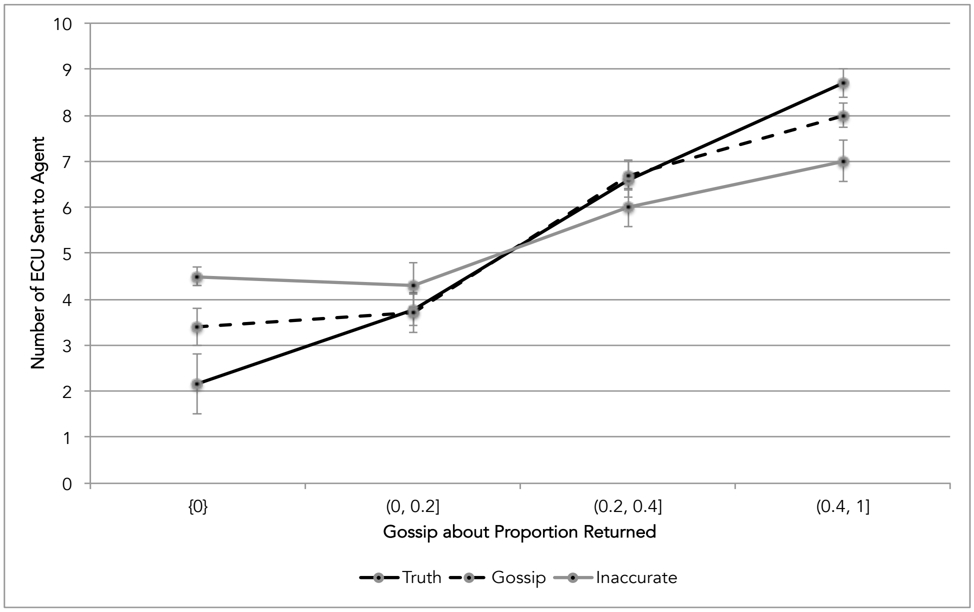

Figure 4: Trust as measured by ECU sent to Agents as a function of reported Agent trustworthiness. Error lines represent robust standard errors clustered at the network level around the mean.

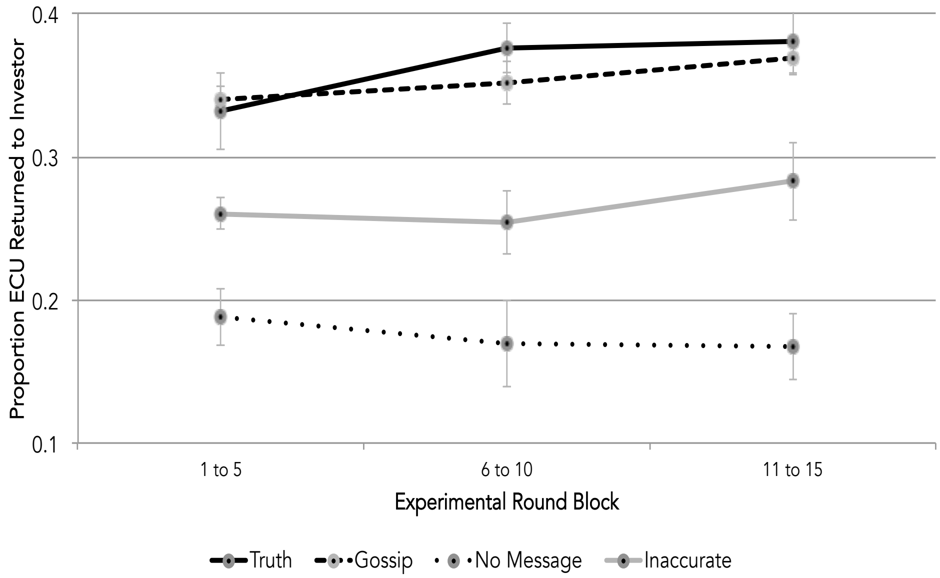

Figure 5: Trustworthiness as measured by proportion of ECU received that was returned to Investors. Error lines represent robust standard errors clustered at the network level around the mean.

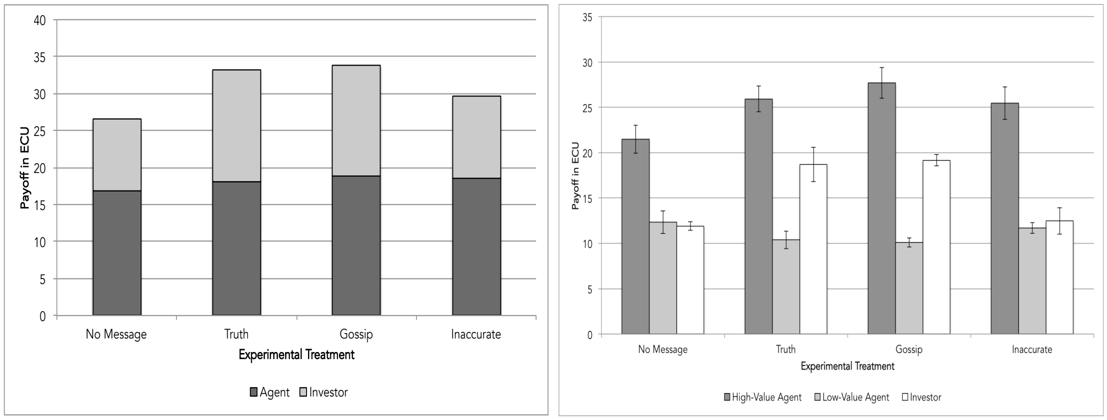

Figure 6: **Left Panel**: Average joint payoffs as a function of treatment. Dark (light) section of each bar denotes average Agent (Investor) payoff, and their share of welfare. **Right Panel**: Average payoffs as a function of role and treatment. Error bars denote robust standard errors clustered at the network level.

Appendix A: Instruction sets

In this section of the Appendix we reproduce the instructions presented to participants. Within a given treatment, instructions for different roles only differed in the first page for all treatments.

Instruction Set (No Message – Sender)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Sender**. You will retain this role throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.

Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

Instruction Set (No Message – Receiver)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Receiver**. There are two kinds of Receivers: **High-Type** and **Low-Type Receivers**. Your type will be allocated randomly by the computer. You will find out what type you are when the experiment starts. You will retain this role and type throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.
Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

Instruction Set (Truth – Sender)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Sender**. You will retain this role throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.

Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

End of Round Sender Decision.

At the end of each round, Senders are asked to state to the experimenter what happened in the previous round: that is, the amount they sent, the type of Receiver they were matched with and the amount sent back by that Receiver.

Round 2 onwards.

From the start of Round 2 onwards, we will also inform each Sender about what type of Receiver they have been paired with for that round (i.e. High or Low). We will also inform each Sender about the receiver's decision in the previous round (i.e., the amount that they received, and the amount they chose to send back).

Instruction Set (Truth – Receiver)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Receiver**. There are two kinds of Receivers: **High-Type** and **Low-Type Receivers**. Your type will be allocated randomly by the computer. You will find out what type you are when the experiment starts. You will retain this role and type throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.
Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.
Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

End of Round Sender Decision.

At the end of each round, Senders are asked to state to <u>the experimenter</u> what happened in the previous round: that is, the amount they sent, the type of Receiver they were matched with and the amount sent back by that Receiver.

Round 2 onwards.

From the start of Round 2 onwards, we will also inform each Sender about what type of Receiver they have been paired with for that round (i.e. High or Low). We will also inform each Sender about the receiver's decision in the previous round (i.e., the amount that they received, and the amount they chose to send back).

Instruction Set (Gossip – Sender)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Sender**. You will retain this role throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.

Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

End of Round Sender Decision.

At the end of each round, Senders are asked to state to <u>the next Sender who will play with the Receiver they played with in that round</u> what happened in the previous round: that is, the amount they sent, the type of Receiver they were matched with and the amount sent back by that Receiver.

Round 2 onwards.

From the start of Round 2 onwards, Senders will receive a message from the Sender who played with their current match in the previous round, which was produced at the end of the round. The message will contain information about the type of Receiver they have been paired with for that round (i.e. High or Low), and the Receiver's decision in the previous round (i.e., the amount that they received, and the amount they chose to send back).

Instruction Set (Gossip – Receiver)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Receiver**. There are two kinds of Receivers: **High-Type** and **Low-Type Receivers**. Your type will be allocated randomly by the computer. You will find out what type you are when the experiment starts. You will retain this role and type throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.

Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].
- The receiver's payoff for the round = amount received – amount returned.

End of Round Sender Decision.

At the end of each round, Senders are asked to state to the next Sender who will play with the Receiver they played with in that round what happened in the previous round: that is, the amount they sent, the type of Receiver they were matched with and the amount sent back by that Receiver.

Round 2 onwards.

From the start of Round 2 onwards, Senders will receive a message from the Sender who played with their current match in the previous round, which was produced at the end of the round. The message will contain information about the type of Receiver they have been paired with for that round (i.e. High or Low), and the Receiver's decision in the previous round (i.e., the amount that they received, and the amount they chose to send back).

Instruction Set (Inaccurate - Sender)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Sender**. You will retain this role throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.

Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

End of Round Sender Decision.

At the end of each round, Senders are asked to state to the next Sender who will play with the Receiver they played with in that round what happened in the previous round: that is, the amount they sent, the type of Receiver they were matched with and the amount sent back by that Receiver.

Round 2 onwards.

From the start of Round 2 onwards, Senders will receive a message stating the type of Receiver they have been paired with for that round (i.e., High or Low), and the Receiver's decision in the previous round (i.e., the amount that they received, and the amount they chose to send back).

Before the message is received, the computer will flip a coin: if heads comes up, the message a Sender receives will be exactly that which was sent by the previous Sender; if tails come up, the message will instead be about another Receiver's previous behaviour. You will not know whether the computer flipped heads or tails.

Instruction Set (Inaccurate – Receiver)

Welcome to our experiment. Please read these instructions carefully, as your payment from taking part in this experiment will depend on the decisions you and other people in the room take. It is important therefore that you have a good understanding of the rules of the experiment.

Payoffs in the experiment are denominated in Experimental Currency Units (ECU). 3 ECU are worth £1. At the end of the experiment, we will select one round at random and convert your payoff in that round to pounds and pay you in cash. You will additionally receive your £5 show up fee.

In the experiment there are two different players: Senders and Receivers. Your role is that of a **Receiver**. There are two kinds of Receivers: **High-Type** and **Low-Type Receivers**. Your type will be allocated randomly by the computer. You will find out what type you are when the experiment starts. You will retain this role and type throughout the session.

You will play a number of rounds in the experiment. At the beginning of every round, we will randomly pair a Sender and a Receiver, who will play together for that round. This matching will be done in such a way that you will never be paired with the same person in two consecutive rounds. Each round consists of a sender decision and a receiver decision.

Sender Decision.

Each **Sender** starts each round with 10 ECU and must decide how much of that amount to send to the Receiver with whom he/she is playing. This could be any whole number between 0 and 10.

There are two kinds of **Receivers**: High-Type and Low-Type Receivers.

Each ECU a Sender sends to a **High-Type Receiver** is multiplied by **6**.

- For instance if a Sender sends 4 ECU, the Receiver will get 24 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 48 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Each ECU a Sender sends to a **Low-Type Receiver** is multiplied by **3**.

- For instance if a Sender sends 4 ECU, the Receiver will get 12 ECU.
- Alternatively, if a Sender sends 8 ECU, the Receiver will get 24 ECU.
- If however, a Sender sends 0 ECU, the Receiver will get 0 ECU.

Receiver Decision.

Once the Receiver learns how much money he or she has received from the sender, he or she has to decide how much money to return to the Sender. The Receiver can return any amount he or she wants, from 0 to the maximum amount possible (that is the total amount received from the sender). The only restriction is that this amount is in **whole** numbers (i.e., 1, 2, 3 etc). Any money that the receiver does not return he or she gets to keep.

Payoff for Round.

Once the Sender receives whatever ECUs the Receiver chose to return, the round is over.

- The Sender's payoff for the round = [10 – amount sent to receiver] + [amount returned by receiver].

- The receiver's payoff for the round = amount received – amount returned.

End of Round Sender Decision.

At the end of each round, Senders are asked to state to <u>the next Sender who will play with the Receiver they played with in that round</u> what happened in the previous round: that is, the amount they sent, the type of Receiver they were matched with and the amount sent back by that Receiver.

Round 2 onwards.

From the start of Round 2 onwards, Senders will receive a message stating the type of Receiver they have been paired with for that round (i.e., High or Low), and the Receiver's decision in the previous round (i.e., the amount that they received, and the amount they chose to send back).

Before the message is received, the computer will flip a coin: if heads comes up, the message a Sender receives will be exactly that which was sent by the previous Sender; if tails come up, the message will instead be about another Receiver's previous behaviour. The Sender will not know whether the computer flipped heads or tails.

Appendix B: Supplementary Tables

	Message about % returned by Agent							
	{0}		(0, 0.2]		(0.2, 0.4]		(0.4, 1]	
ECU Sent	Sender Payoff	N	Sender Payoff	N	Sender Payoff	N	Sender Payoff	N
0	10.00*	25	10.00*	14	10.00	11	10.00	1
	(0.00)		(0.00)		(0.00)		(0.00)	
1	9.00	2	9.50	2	10.00	1		0
	(0.00)		(0.71)		(0.00)			
2	8.67	3	9.40	5	9.67	3	10.00	1
	(1.15)		(1.67)		(1.53)		(0.00)	
3	-	0	9.00*	8	9.33	9	10.00	5
			(2.27)		(1.66)		(2.24)	
4	8.25	4	10.33	3	10.83	12	12.00	10
	(4.50)		(4.16)		(3.64)		(3.74)	
5	-	0	12.40	5	12.91	11	15.60	5
			(5.37)		(3.67)		(2.51)	
6	6	1	10.25	4	14.14	7	14.07	14
	(0.00)		(0.50)		(4.26)		(3.97)	
7	-	0	15.00	3	13.42	12	15.43	14
			(7.21)		(5.25)		(4.55)	
8	6.00	1	14.33	3	12.31	16	16.77	22
	(0.00)		(5.13)		(5.46)		(6.56)	
9	-	0	5.00	1	21.00*	8	19.14	21
			(0.00)		(5.73)		(6.39)	
10	8.00	5	11.00	5	20.27*	26	21.32*	143
	(7.58)		(6.52)		(7.79)		(7.89)	

Standard deviations in parentheses.

*: (joint) maximum profit conditional on message for cases with at least 6 obs.

Table B1: Average Profit conditional on ECU sent and message concerning % of amount returned by Receiver in previous round, Truth treatment.

	Message about % returned by Agent							
	$\{0\}$		$(0, 0.2]$		$(0.2, 0.4]$		$(0.4, 1]$	
ECU Sent	Sender Payoff	N	Sender Payoff	N	Sender Payoff	N	Sender Payoff	N
0	10.00*	16	10.00	24	10.00	12	10.00	8
	(0.00)		(0.00)		(0.00)		(0.00)	
1	10.67	3	10.50	2	9.75	4	9.00	1
	(1.53)		(2.12)		(1.50)		(0.00)	
2	10.00	2	9.25	4	9.75	8	8.00	1
	(1.41)		(0.50)		(1.50)		(0.00)	
3	13.00	3	8.80	5	9.91	11	11.00	2
	(0.00)		(1.10)		(2.51)		(1.41)	
4	10.50	2	11.13*	8	11.17	12	10.77	13
	(2.12)		(6.31)		(2.76)		(1.74)	
5		0	10.00	9	11.67	18	11.92	12
			(2.60)		(2.20)		(3.53)	
6	8.33	3	13.00	3	12.00	13	12.94	16
	(2.08)		(4.58)		(3.43)		(4.09)	
7		0	15.00	3	16.67*	6	15.69	13
			(6.08)		(2.42)		(4.59)	
8	13.00	1	11.50	2	16.07	14	19.11	19
	(0.00)		(3.54)		(3.81)		(5.88)	
9	27.00	1		0	16.44*	9	20.13	15
	(0.00)				(8.79)		(5.85)	
10	9.71*	7	13.60*	10	18.83*	53	22.07*	98
	(5.56)		(7.29)		(6.70)		(9.30)	

Standard deviations in parentheses.

*: (joint) maximum profit conditional on message for cases with at least 6 obs.

Table B2: Average Profit conditional on ECU sent and message concerning % of amount returned by Receiver in previous round, Gossip treatment.

	Message about % returned by Agent							
	$\{0\}$		$(0, 0.2]$		$(0.2, 0.4]$		$(0.4, 1]$	
ECU Sent	Sender Payoff	N	Sender Payoff	N	Sender Payoff	N	Sender Payoff	N
0	10*	20	10.00*	7	10.00	5	10.00*	6
	(0.00)		(0.00)		(0.00)		(0.00)	
1	9.77*	22	10.00	5	10.00	5	10.00	1
	(1.11)		(1.22)		(1.22)		(0.00)	
2	9.55*	11	8.86*	7	9.25	4	8.00	2
	(1.97)		(1.46)		(0.96)		(0.00)	
3	8.64*	11	11.00	4	9.50*	10	9.88*	8
	(2.69)		(4.69)		(2.59)		(3.27)	
4	11.16*	19	10.30*	10	12.00*	22	11.31*	16
	(4.21)		(3.47)		(3.61)		(2.55)	
5	8.25	4	15.20	5	11.00*	12	10.94*	18
	(4.72)		(3.03)		(5.83)		(4.53)	
6	11.14*	7	9.86*	7	12.93*	15	11.85*	13
	(7.97)		(3.24)		(3.56)		(5.54)	
7	12.92*	12	9.50	4	13.56*	18	12.76*	25
	(8.56)		(4.51)		(7.16)		(6.68)	
8	14.25	4	12.50	2	15.67*	6	12.89*	9
	(12.47)		(6.36)		(17.81)		(7.83)	
9	4.00	1			13.00	3	9.40	5
	(0.00)				(6.56)		(5.50)	
10	6.62	29	15.00	5	12.07*	29	12.04*	54
	(8.15)		(8.75)		(9.26)		(10.46)	

Standard deviations in parentheses.

* denotes (joint) maximum profit conditional on message.

Table B3: Average Profit conditional on ECU sent and message concerning % of amount returned by Receiver in previous round, Inaccurate treatment.

Supplementary Information

Additional Econometric Analysis

In this section, we provide supporting evidence that could not be included in the text due to space restrictions, but to which we refer in the text.

Table C1 displays the correlation coefficients between (i) gossip about points sent by Investors and actual points sent; (ii) gossip about Agents' type and their actual type; (iii) gossip about proportion returned by Agents and actual proportion returned by Agents. It displays the behaviour/message correlation coefficients for the contents of the messages sent, as well as the messages received. These two metrics differ in the Truth treatment because, here, sent messages were replaced with the summary of the Agents' actual behaviour; they differ in the Inaccurate treatment because of the 50% probability that messages were misdirected.

Table C2 displays estimates of inaccuracy in Investor messages, comparing Inaccurate and Gossip treatments.

Table C3 summarises regression estimates of determinants of single round Investor trust as a function of reputation treatment. This replicates findings presented in the main text by showing that levels of Investor trust do not differ at intercept across treatments, but that trust levels increase across rounds in the Truth and Gossip treatments only.

Table C4 provides the estimates on which Figures 1 and 4 are based. Table C5 provides the estimates on which Figure 2 is based. Table C6 provides the estimates on which both panels in Figure 5 are based. Table C7 provides the estimates on which Figures 3 is based. Table C8 provides regression model estimates of trust as a function of message about Receiver type.

Messages Sent	Truth	Gossip	Inaccurate
Agent Type	0.96	0.92	0.85
ECU Sent	0.96	0.80	0.83
ECU Returned	0.93	0.80	0.83
Messages Received	**Truth**	**Gossip**	**Inaccurate**
Agent Type	1.00	0.92	0.27
ECU Sent	1.00	0.80	0.35
ECU Returned	1.00	0.80	0.26

Table C1: Correlation between messages sent, messages received and actual behaviour. Notes: entries on cells are Spearman correlation coefficients. The top half of the table displays the correlations between an Agent's behaviour and messages sent about it; the bottom half displays the correlations between messages received and the behaviour of the Agent it is supposed to describe.

DV:	Inaccuracy in Receiver Type	Inaccuracy in Message ECU Sent	Inaccuracy in Message ECU Returned
Gossip	-0.304***	-0.411***	-0.326***
	(0.022)	(0.025)	(0.039)
Constant	0.341***	0.506***	0.513***
	(0.016)	(0.018)	(0.028)
Random Effects Parameters			
Session parameter:	< 0.001 (< 0.001)	0.001 (< 0.001)	0.001 (< 0.001)
Investor, intercept parameter:	0.001 (< 0.001)	0.001 (< 0.001)	0.001 (< 0.001)
Investor, residual:	0.361 (0.008)	0.410 (0.009)	0.446 (0.010)
N (obs, Investors)	(1,080, 72)	(1,080, 72)	(1,080, 72)
Wald $\chi^2(3)$	$191.34, p < 0.001$	$272.04, p < 0.001$	$69.50, p < 0.001$

Standard errors in parentheses. $^{**}: p < 0.01$.

Table C2: Regression model estimates of inaccuracy in Investor messages, comparing Inaccurate and Gossip treatments. Notes: Model is a three-level random effects maximum likelihood regression model that nested rounds in participants in sessions; intercepts were allowed to vary.

DV: ECU sent		
Truth	0.006	(0.623)
Gossip	0.154	(0.623)
Inaccurate	0.555	(0.623)
Round	-0.041	(0.028)
Truth× Round	0.201**	(0.040)
Gossip× Round	0.177**	(0.040)
Inaccurate × Round	0.036	(0.040)
Constant	5.119**	(0.441)
Random Effects Parameters		
Network:	0.407	(0.324)
Investor (intercept):	1.904	(0.141)
Investor (residual):	2.837	(0.045)
N (obs, Investors):	(2,160, 144)	
Wald $\chi^2(7) = 69.22, p < 0.0001$		

$^{**}: p < 0.01.$

Table C3: Regression model estimates of Investor trust as a function of treatment. Notes: Model was a three-level random effects maximum likelihood regression model that nested rounds in participants in sessions; intercepts were allowed to vary. No message treatment provides the reference category.

DV	% Returned By Agents		ECU Sent	
Truth × (Periods 1-5)	0.332***	(0.027)	5.567***	(0.531)
Truth × (Periods 6-10)	0.376***	(0.017)	6.539***	(0.675)
Truth × (Periods 11-15)	0.381***	(0.022)	7.106***	(0.663)
Gossip × (Periods 1-5)	0.340***	(0.010)	5.717***	(0.242)
Gossip × (Periods 6-10)	0.352***	(0.016)	6.389***	(0.248)
Gossip × (Periods 11-15)	0.369***	(0.012)	6.967***	(0.320)
Control × (Periods 1-5)	0.188***	(0.021)	5.067***	(0.350)
Control × (Periods 6-10)	0.169***	(0.029)	4.694***	(0.481)
Control × (Periods 11-15)	0.167***	(0.023)	4.606***	(0.487)
Inaccurate × (Periods 1-5)	0.260***	(0.011)	5.728***	(0.399)
Inaccurate × (Periods 6-10)	0.254***	(0.022)	5.528***	(0.285)
Inaccurate × (Periods 11-15)	0.283***	(0.027)	5.628***	(0.385)
N	1,883†		2,160	
R^2	0.74		0.74	

Session-level clustered standard errors in parentheses.

*** : $p < 0.01$; †: 277 observations excluded due to 0 ECU sent.

Table C4: OLS estimates of percentage returned by Agents in 5-round blocks in all treatments.

DV:	Negative Lie		Positive Lie	
PctRet $\in (0, 10]$	-	-	-0.031	(0.050)
PctRet $\in (10, 20]$	-0.181***	(0.060)	0.015	(0.032)
PctRet $\in (20, 30]$	-0.193***	(0.058)	-0.008	(0.032)
PctRet $\in (30, 40]$	-0.207***	(0.057)	-0.001	(0.027)
PctRet $\in (40, 50]$	-0.219***	(0.055)	-0.030	(0.025)
PctRet $\in (50, 1]$	-0.153***	(0.069)	-0.055	(0.042)
Constant	0.310***	(0.057)	0.087***	(0.026)
Random Effects Parameters				
Network:	<0.001	(<0.001)	<0.001	(<0.001)
Investor (intercept):	0.193	(0.019)	0.134	(0.013)
Investor (residual):	0.255	(0.007)	0.225	(0.005)
N (obs, Investors):	(830, 72)		(975, 72)	
	Wald $\chi^2(2) = 18.24, p = 0.0027$		Wald $\chi^2(6) = 5.19, p = 0.5201$	

Table C5: Estimates of Relative Frequency of Investor Lying as a function of Agent trustworthiness. Notes: Model was a three-level random effects maximum likelihood regression model that nested rounds in participants in sessions; intercepts were allowed to vary. Observations are restricted to cases where Investor sent a positive amount. PctRet $\in \{0\}$ was the reference category for Positive Lies regression; PctRet $\in (0, 10]$ was the reference category in the Negative lies regression, as negative lies were not possible if no amount was returned by Agent. $^{***}: p < 0.01$.

DV: Payoff	All	Investors	Agents	High-Type Agent	Low-Type Agent
Truth	16.560***	15.035***	18.156***	25.926***	10.385***
	(0.976)	(1.019)	(1.136)	(1.438)	(0.971)
Gossip	16.882***	14.881***	18.883***	27.678***	10.089***
	(0.407)	(0.315)	(0.841)	(1.693)	(0.491)
Control	13.297***	9.693***	16.902***	21.478***	12.326***
	(0.627)	(0.451)	(1.379)	(1.560)	(1.268)
Inaccurate	14.814***	11.061***	18.567***	25.441***	11.693***
	(0.573)	(0.613)	(1.116)	(1.801)	(0.579)
N	4,320	2,160	2,160	1,080	1,080
R^2	0.65	0.79	0.63	0.73	0.69

Session-level clustered standard errors in parentheses. *** : $p < 0.01$;

Table C6: OLS estimates of average payoffs in all treatments.

	Message about % returned by Agent			
DV: ECU Sent	$\{0\}$	$(0, 0.2]$	$(0.2, 0.4]$	$(0.4, 1]$
Truth	2.146***	3.792***	6.233***	8.725***
	(0.656)	(0.358)	(0.394)	(0.314)
Gossip	3.395***	3.671	6.463***	7.955***
	(0.415)	(0.429)	(0.334)	(0.270)
Inaccurate	4.500***	4.143***	5.984***	6.962***
	(0.215)	(0.489)	(0.421)	(0.452)
N	219	179	405	591
R^2	0.55	0.58	0.80	0.91

Session-clustered standard errors in parentheses. *** : $p < 0.01$.

Table C7: Average ECU sent conditional on message concerning % of amount returned by Receiver in previous round.

DV: ECU sent	Truth	Gossip	Inaccurate
Constant	6.034***	5.372***	5.694***
	(0.417)	(0.417)	(0.418)
High-Value Type	1.035***	2.105***	-0.118
	(0.265)	(0.266)	(0.268)
Random Effects Parameters			
Network:	0.601	(0.270)	
Investor (intercept):	1.689	(0.153)	
Investor (residual):	2.930	(0.055)	
N (obs, Investors):	(1,512, 108)		
Wald $\chi^2(1) = 81.30, p < 0.0001$			

$^{***}: p < 0.01$.

Table C8: Regression model estimates of trust as a function of message about Receiver type. Notes: Model is a three-level random effects maximum likelihood regression model that nested rounds in participants in sessions; intercepts were allowed to vary. Coefficients are interactions between IVs in first column with dummy IVs in the first row.