

IZA DP No. 8216

Identity Changes and the Efficiency of Reputation Systems

Matthias Wibrál

May 2014

Identity Changes and the Efficiency of Reputation Systems

Matthias Wibrals

*University of Bonn
and IZA*

Discussion Paper No. 8216
May 2014

IZA

P.O. Box 7240
53072 Bonn
Germany

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

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

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

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

ABSTRACT

Identity Changes and the Efficiency of Reputation Systems^{*}

Reputation systems aim to induce honest behavior in online trade by providing information about past conduct of users. Online reputation, however, is not directly connected to a person, but only to the virtual identity of that person. Users can therefore shed a negative reputation by creating a new account. We study the effects of such identity changes on the efficiency of reputation systems. We compare two markets in which we exogenously vary whether sellers can erase their rating profile and start over as new sellers. Buyer trust and seller trustworthiness decrease significantly when sellers can erase their ratings. With identity changes, trust is particularly low towards new sellers since buyers cannot discriminate between truly new sellers and opportunistic sellers who changed their identity. Nevertheless, we observe positive returns on buyer investment under the reputation system with identity changes, and our evidence suggests that trustworthiness is higher than in the complete absence of a reputation system.

JEL Classification: C91, D02, L14

Keywords: trust, reputation, identity changes

Corresponding author:

Matthias Wibral
Institute for Applied Microeconomics
University of Bonn
Adenauerallee 24–42
53113 Bonn
Germany
E-mail: wibral@uni-bonn.de

^{*} Financial support from the German Research Foundation (grant KR 2077/2-1 and SFB/TR 15) is gratefully acknowledged. The author would like to thank Johannes Abeler, Steffen Altmann, Christine Harbring and Armin Falk for many insightful discussions and Markus Antony, Holger Gerhardt, Alexander Koch, Sebastian Kube, Rosemarie Nagel, Axel Ockenfels, Gert Pönitzsch, Mirko Seithe, Dirk Sliwka, Florian Zimmermann and seminar and conference participants at Alicante, Amsterdam, Bonn, Cologne, Graz, Milan and Zurich for helpful comments.

1 Introduction

Online trade usually takes place between strangers, payment occurs before the good is shipped, and legal enforcement of an agreement may be prohibitively costly. This particular constellation generates moral hazard and adverse selection problems, especially on the seller side. Sellers have an incentive to ship a good of lower quality than promised or not to ship at all. Reputation systems are the most important tool in e-business to address these problems and to induce honest behavior among users. They store information about past conduct of a user provided by other users in a reputation profile and disseminate this information to the whole community. In principle, a reputation profile thus allows buyers to distinguish honest sellers from dishonest ones and to interact only with the former.

The successful use of reputation systems in traditional markets with similar moral hazard problems can be traced back at least to the beginnings of long-distance trade in the Middle Ages (Greif 1989, Milgrom et al. 1990). Online reputation in most settings, however, differs from reputation in traditional markets in a very important way. Online reputation is only connected to the virtual identity of a person, i.e., the user name, and not the person itself. After a bad rating, a dishonest seller can comparatively easily abandon his old virtual identity and create a new account under a new user name with no reputation attached to it. The implications of this distinctive feature of online reputation for the efficiency of reputation systems are not well understood.

At first sight, identity changes seem to severely weaken the disciplinary power of a reputation system. Dishonest behavior can still be punished with a negative rating but these ratings may lose their edge when they can easily be shed by creating a new identity. Real newcomers and sellers who have changed identity after a bad rating become indistinguishable. This argument suggests a high frequency of identity changes accompanied by opportunistic behavior and lower buyer trust. Theory on the other hand also provides some guidance on why reputation systems might still work effectively (Friedman and Resnick 2001, Ockenfels 2003). If buyers anticipate that dishonest sellers start over as new players they will not trust newcomers or only interact with them at very unfavorable conditions, e.g., low sale prices. Starting as a new seller may then become so costly that cheating and creating a new identity is not profitable anymore. In this case, we would observe a high level of seller trustworthiness and no identity changes. Buyer trust would be high in transactions with experienced sellers and low towards new sellers. This creates a negative externality for honest new sellers.

In this paper, we examine empirically how the option to change one's virtual identity affects the efficiency of markets and the performance of reputation systems in inducing trust and trustworthiness. We study two experimental markets in which buyers play the trust game (Berg et al. 1995) with varying sellers. Buyers can rate sellers after each transaction. Before deciding whether and how much to trust a seller, a buyer can see the rating profile of this seller. New players enter the market over time in both treatments. The only difference is that sellers in one market can change their identity (*change treatment*), i.e., erase their rating profile and start over as new players, while in the other market this is not possible (*no-change treatment*). In view of the discussion above we aim at answering the following questions: Do sellers use the opportunity to change their identity? If so, does this go along with an increase in opportunistic behavior? How is

buyer trust affected? In particular, are new members treated differently when identity changes are possible?

These questions are of high practical relevance given widespread reports that dishonest behavior remains an important problem for online interaction.¹ According to a survey conducted by the European Commission, 49% of internet users in the EU 27 countries are concerned about becoming a victim of online fraud, where goods are not delivered, counterfeit or not as advertised, and 19% have actually become a victim (European Commission 2012). In an earlier survey (European Commission 2004), 21% of the EU 15 population reported abstaining from buying online because they did not trust the internet. Recent studies on eBay's reputation system (Reichling 2004, Jin and Kato 2006, Dellarocas and Wood 2008, Klein et al. 2009, Bolton et al. 2013) also indicate that the low percentages of neutral and negative ratings (e.g., Resnick and Zeckhauser 2002) grossly underestimate the true extent of problematic transactions since users may be reluctant to provide negative feedback for fear of retaliatory ratings (Masclat and Peenard (2012)). Dellarocas and Wood (2008) estimate the true figure of mildly or very dissatisfied buyers at up to 21%.

It is difficult to determine whether these problems (also) arise because sellers can shed a negative reputation comparatively easily. So far, there is little evidence on the effects of identity changes, mainly because identity changes are inherently difficult to observe in field data. In addition, even if one could observe identity changes, comparing the same reputation system with and without the possibility to change one's identity would be extremely difficult in the field. We therefore use a controlled laboratory experiment to address our research questions.

The main findings can be summarized as follows. First, buyer trust and seller trustworthiness are significantly lower when sellers can change their identities. Trust is especially lower for new sellers. However, the reputation system in the change treatment maintains trustworthiness at a level that is high enough to make investing profitable for the buyers. The evidence is at least suggestive that trustworthiness is also higher than in the complete absence of a reputation system. Second, the basic principles of reputation systems function in both treatments. Lower trustworthiness translates into lower ratings, and buyers trust sellers with lower ratings less. Third, in the change treatment sellers with a sufficiently bad reputation circumvent lower buyer investments by switching identity and starting over as a new seller. Incentives to be trustworthy are therefore lower. Sellers seem to choose between two kinds of behavior—being trustworthy and maintaining a fixed identity, and behaving opportunistically and changing identity multiple times. The latter group drives the overall lower trustworthiness in the change treatment. Finally, opportunistic behavior does not pay off in the no-change treatment. In the change treatment, however, opportunistic players do earn more than their counterparts who share more equally with the buyer.

We believe that this work contributes on several fronts. The results for the change treatment demonstrate that the efficiency of reputation systems in inducing trust and trustworthiness is reduced if sellers have the opportunity to change their identity. While the negative effects of identity changes have long been discussed theoretically in the literature (e.g. Dellarocas and Wood 2003) this paper provides controlled evidence

¹See, for example, the reports prepared by the National Consumer League (www.fraud.org) or the Internet Crime Complaint Center (www.ic3.gov).

and establishes a *causal link* between the possibility of identity changes and a higher incidence of dishonest behavior. This is the main contribution of our paper.

Our results for the no-change treatment complement the findings of previous laboratory experiments on reputation systems without identity changes. Keser (2003), Bolton et al. (2004) and Masclet and Peenard (2012) find that a reputation system leads to high levels of trust and trustworthiness. We extend this finding to an environment in which new players enter the market over time: The mere fact that there are new players need not affect the efficiency of a reputation system as long as buyers can be sure that a new seller really is a person who has just entered the market.

In addition, the controlled environment of the laboratory allows us to take a closer look at the mechanisms through which identity changes influence market outcomes. Several of our specific findings, which we can causally attribute to potential identity changes, are consistent with observations from online platforms. For example, Ockenfels (2003) finds a very low percentage of sellers with more negative than positive ratings in a sample from the platform *half.com*. In the same sample, new sellers also ask for lower prices than experienced sellers. Resnick et al. (2006) conduct a randomized field experiment on eBay, in which matched pairs of vintage postcards are sold via a high reputation identity and accounts with little or no previous feedback. Buyers are willing to pay 8% more to the high reputation seller. While the design of their study nicely controls for many potential confounds that plagued earlier empirical studies, it is still unclear to which degree the effect is driven by personal experience of repeat customers of the high reputation identity. We can rule out this and other potential confounds in our setup. Finally, our results also show that a substantial part of the negative effects of a system which allows identity changes may arise from buyers simply refusing to participate in market interaction, thus effectively dropping out of the market. In this aspect, we complement survey evidence indicating that large efficiency losses occur because consumers do not trust online trade(rs).

The rest of the paper is structured as follows. Section 2 describes the experimental design. Section 3 reports and discusses our results. Section 4 concludes.

2 Experimental Design

2.1 Trust Game and Reputation System

We use the trust game, originally introduced as investment game by Berg et al. (1995), to mimic essential features of online interaction in our experiment.² The trust game is played by two players, which we refer to as buyer and seller in our setup. Both players have an endowment of 10. In the first stage, the buyer can send an amount between 0 and 10—the *investment*—to the seller. On the way to the seller this amount is tripled by the experimenter. In a second stage, the seller can then decide how much (if anything at all) of the tripled amount he wants to send back to the buyer.

²Two baseline paradigms have been used to study reputation systems. Keser (2003) and Masclet and Peenard (2012) who also use the standard trust game, while Bolton et al. (2004) use a binary variant (“buyer-seller game”), in which efficiency gains only arise when trust is honored. We use the standard trust game for our main analysis because it allows studying certain types of strategic seller behavior. For example, we can study whether sellers build up their reputation when it is relatively cheap to do so and then “milk” this reputation later on. However, we also test the robustness of our results using the buyer-seller game (see section 3.2.2).

The trust game captures the basic moral hazard problem inherent in a broad class of social interactions. There are efficiency gains to be realized but this requires that the buyer sends money without a guarantee of receiving anything in return from the seller. The buyer's investment can therefore be taken as a measure of his trust while the amount that the seller sends back (as a share of what he received) reflects the latter's trustworthiness. More precisely, our measure of trustworthiness is *return on (buyer) investment* which we define as $[(\text{Amount Received}/\text{Amount Sent}) - 1]$. Trustworthiness can, of course, only be measured when the seller was trusted in the first place, i.e., in those instances when buyer investment is greater than zero.

In our experiment subjects play the trust game described above, half of them in the role of buyers and the other half in the role of sellers. Every participant keeps his role throughout the entire experiment which lasts 20 rounds. This is common knowledge. Buyers and sellers are randomly rematched after every round. The matching mechanism ensures that players never play with the same person in two consecutive rounds, which is common knowledge as well.

Based on Keser (2003), we enhance the trust game with a reputation system in the following way. At the end of each round, after the seller has decided how much to return to the buyer, the latter *has to* rate this decision.³ He can give the seller a positive, neutral or negative rating. If there was no interaction, i.e., the buyer did not send anything to the seller, no rating can be made. The rating is stored in the seller's rating profile.

At the beginning of each round, the seller sees his rating profile. It lists four categories: the last rating and the total number of positive, neutral and negative ratings. In the *no-change treatment* every seller keeps his rating profile for the whole experiment and therefore is simply informed about his current profile. Then the buyer sees the seller's rating profile and decides how much money to send to the seller. In the *change treatment* the seller can decide *before* the buyer sees his rating profile whether he wants to erase the profile and start as a new player. This profile erasure is what we refer to as an identity change. If the seller keeps his rating profile it is then shown to the buyer. If the seller *erases* his profile, the buyer will only be given the information that he is matched with a *new* player before making his investment decision. Identity changes are free to the seller.

If buyers in real online markets face a new seller they are unable to tell whether this person really is a newcomer or a (dishonest) seller who has abandoned his old account and created a new one. To achieve this crucial feature in our laboratory experiment we let new subjects enter the experiment over time. Only six subjects (3 buyers, 3 sellers) start the experiment right from the beginning in every session. The others enter in the course of the experiment. The timing of the entry of new players is the same for all sessions in both treatments. One buyer and one seller each enter in rounds 4, 8 and 13, respectively.⁴ Players who entered the experiment at a later stage had to do a paid real effort task

³We thus abstract from the public good problem of voluntary feedback provision. Gazzale and Khopkar (2011) study a setup without identity changes, in which buyers can decide whether they want to leave feedback or not. Since there is no scope for opportunistic buyer behavior in our setup, sellers cannot rate buyers in our experiment.

⁴We chose this timing of entry to minimize the number of repeat encounters. In the post-experimental questionnaire, only 2 out of 192 subjects indicated that they thought they had recognized someone whom they had played with before.

which consisted of counting certain letters in a short text. We introduced the real effort task to keep subjects busy and to avoid income effects. Payments were calibrated such that earnings from the real effort experiment were not significantly different from those of the participants who played the trust game right from the beginning. Subjects know that new players enter the experiment but not when and how many. When faced with a new seller buyers in the change treatment thus cannot discern whether the seller really is a new player or has just erased his rating profile.

In addition to our two main treatments, we conducted a control treatment without a reputation system to establish baseline levels of trust and trustworthiness. Apart from the reputation system the control treatment is completely identical to the two other treatments in all aspects (matching, players per session, entry of new players, etc.). Subjects thus play a series of one-shot trust games with changing partners.

The experiment was programmed using z-Tree (Fischbacher 2007) and conducted at the BonnEconLab at the University of Bonn. Eight sessions of each treatment with 12 subjects each were played, yielding a total of 288 participants. At the beginning of each session participants were seated in separate cubicles. Before the experiment started the experimenter read out the instructions and answered all remaining questions privately.⁵ After the experiment, subjects answered a short questionnaire with free form questions regarding the motivation for their decisions. All subjects were undergraduates from different fields recruited via announcements posted on campus. The experiment lasted approximately one hour and subjects earned 12.35 Euro on average, including a show-up fee of 2 Euro.

2.2 Related literature and behavioral predictions

In the trust game presented above, the buyer does not send anything to the seller in the first place when agents only care about their own material payoffs. This follows from the standard unravelling argument. The introduction of a reputation mechanism does not change this result as long as the game is finitely repeated, independent from whether identity changes are possible or not. Based on this standard model we would therefore expect no differences between our treatments.

Previous experiments with finite repetitions of a trust game (Keser 2003, Bolton et al. 2004, Bohnet et al. 2005, Masclet and Peenard 2012) suggest that this is not a good description of actual behavior. In these experiments, a reputation system without identity changes substantially reduces seller moral hazard and improves efficiency. This finding can be rationalized by models which assume different types of players (Kreps et al. 1982). Brown et al. (2004) provide an example in which some players are inequity averse. If some sellers are intrinsically trustworthy in our setup it can be profitable for opportunistic sellers to camouflage as trustworthy at least until the final rounds of play.⁶ Another potential explanation for the success of reputation systems in previous experiments is that subjects are myopic in the sense that they fail to fully account for

⁵A translated version of the instructions can be found in the appendix.

⁶The behavioral reputation literature starting with Camerer and Weigelt (1988) tests the predictions of this type of reputation models by exogenously introducing uncertainty about sellers' preferences via the experimental design. The results of this literature also suggest that other-regarding preferences or an intrinsic concern for appropriate behavior play an important role in understanding reputation (Grosskopf and Sarin 2010).

the finite nature of the experiments.

The only models which explicitly consider the effects of identity changes (Friedman and Resnick 2001, Ockenfels 2003) use a setup with an infinite time-horizon. They stand in the tradition of earlier work (Kandori 1992, Okuno-Fujiwara and Postlewaite 1995) demonstrating that the seller moral hazard problem can be completely solved with a reputation system when identity changes are not possible, even in large communities, in which agents only interact infrequently and enter and exit over time. The simple intuition is that sellers, who behave opportunistically, can be punished later on by other members of the community if past conduct is common knowledge. With identity changes, direct punishment of opportunistic players is not feasible anymore since players with a bad reputation can—in anticipation of punishment—simply reenter the game as new players with no reputation at all. However, even though the direct link “opportunistic behavior–bad reputation–future punishment” is severed, this need not lead to a complete breakdown of cooperation. Both models of identity changes (Friedman and Resnick 2001, Ockenfels 2003) show that an equilibrium characterized by high levels of cooperation can be sustained.

The idea behind this result is straightforward. Since there is no way to distinguish between truly new players and opportunistic players, who have reentered the game, new players will not be trusted and receive unfavorable treatment from experienced players. This is referred to as making new players “pay their dues”. In an online auction setup, for example, new sellers could receive a lower sales price for their goods. If this unfavorable treatment of new players is sufficiently severe, i.e., the dues for newcomers are high enough, opportunistic behavior and subsequent identity changes do not pay off anymore. These “pay your dues” strategies lead to a welfare maximizing (second-best) equilibrium, in which no one behaves opportunistically and changes identities. Efficiency losses compared to a situation without identity changes can arise because truly new players “undeservedly” also receive unfavorable treatment from experienced players. The insights in Friedman and Resnick (2001) and Ockenfels (2003) are derived for an infinite time horizon, but the key mechanisms may also be informative for our finite setup, in particular, if some agents are myopic or intrinsically cooperative.

Summarizing the theoretical arguments and previous experimental evidence, the following hypotheses can be formulated. The two basic principles, upon which reputation systems are based, should work in both reputation treatments: Opportunistic behavior will translate into lower ratings and buyers will condition their investment decision on the rating profile in the sense that lower ratings will lead to lower buyer investments. As a consequence, we expect high levels of trust and trustworthiness in the no-change treatment. The degree to which these carry over to the change treatment crucially depends on the treatment of new players. Based on the models, we hypothesize that buyer investment towards new sellers (and therefore investment overall) will be lower in the change treatment compared to the no-change treatment. If this discount is sufficiently big, the lower buyer investment could make identity changes unprofitable and prevent sellers from behaving opportunistically and changing their reputation.

3 Results

In this section we first present the basic treatment effects. We then explore whether the mechanisms described in the previous section can explain the treatment differences

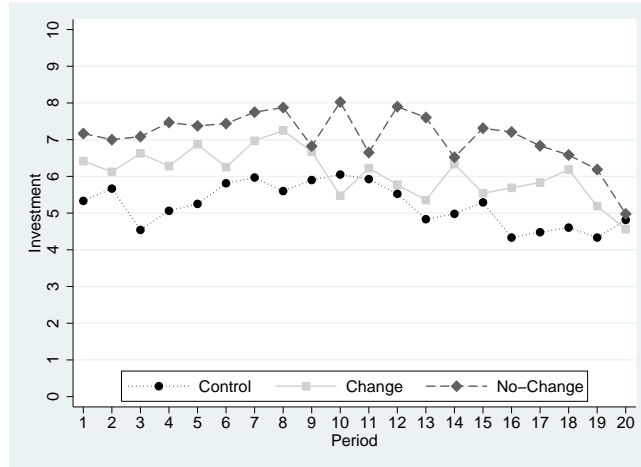


Figure 1: Investment over time

observed between the no-change and the change treatment.

3.1 Overall Treatment Effects

A reputation system should ultimately promote seller trustworthiness and, as a consequence, buyer trust. As described above, we use buyer *investment* to measure trust and *return on (buyer) investment* as a measure of sellers' trustworthiness. Figure 1 shows average buyer investment over time for the two reputation treatments and the control treatment without a reputation system, including those instances in which the buyer did not invest anything at all. Figure 2 displays return on buyer investment over time for the different treatments.

We first examine the effect of identity changes by comparing our two main treatments. Buyers trust more when sellers cannot change their identity. While buyers in the no-change treatment invest 7.04 on average, they send only 6.00 to the sellers in the change treatment. The difference is (weakly) significant using a non-parametric U-test ($p = 0.09$) and when comparing the change to the no-change treatment using a random-effects estimation of buyer investment on treatment dummies (Table 1, columns 1 and 2).⁷ Regarding the relative frequency of the different investment levels the biggest treatment difference occurs for investments of zero, i.e., those instances in which buyers withdraw from the market altogether and which are difficult to observe with field data. Buyers in the no-change invest zero in 64 cases or 8% of all rounds.⁸ Strikingly, buyers in the change treatment do not invest anything almost twice as often, namely in 117 or 15% of all rounds (U-test $p=0.06$).

A similar pattern can be found for trustworthiness. Return on investment is highest in the no-change treatment where buyers on average make a return of 61% on their investment. For every period including the last one, average return on investment is

⁷All reported p-values are two-sided. Unless otherwise noted all non-parametric tests use session averages as independent observations, i.e., $N = 8$ for each treatment. Table 4 in the appendix provides detailed information for each session.

⁸In total 98 rounds were played in each session (3 players in each role played 20 rounds, 1 player each played 17, 13, or 8 rounds respectively). For each role there are thus $(60 + 17 + 13 + 8) \times 8 = 784$ rounds per treatment.

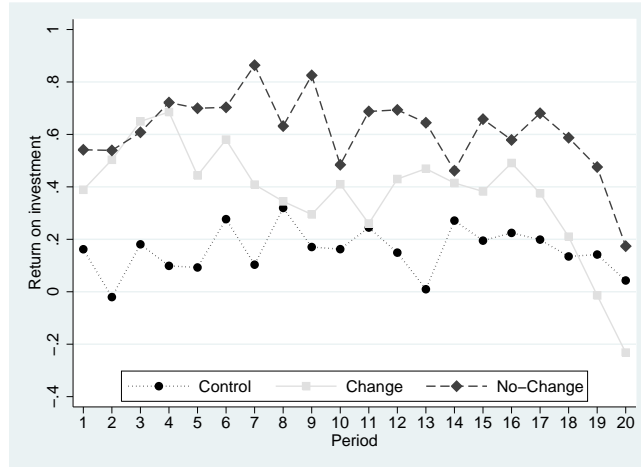


Figure 2: Return on investment over time

positive. The reputation system is less successful in inducing trustworthy behavior in the market in which participants can erase their reputation. Return on investment is significantly lower in the change treatment (U-test, $p = 0.06$; random-effects estimation $p = 0.02$, columns 3 and 4 in Table 1). For example, the most extreme case in which a buyer makes an investment but the seller does not return anything occurs 63 times (8.8% of all investments made) in the no-change treatment but 102 times (15.3%) in the change treatment. However, buyers in the change treatment on average still make a substantial profit of 36% on their investment.⁹

Result 1: *Trust and trustworthiness are significantly lower when sellers can change their identities.*

In the control treatment without reputation system buyers invest 5.16, and return on investment is 16%.¹⁰ A main finding of previous experiments on reputation systems without identity changes (Bohnet et al. 2005, Bolton et al. 2004 and Keser 2003) is that a reputation system increases trust and trustworthiness compared to a situation

⁹Note that, in principle, lower *average* trustworthiness in the change treatment could simply be driven by lower buyer investments if sellers reciprocate lower investments with lower returns on investment. Two analyses indicate that the treatment difference in our setup is not just due to the different distribution of investments, but persists conditional on investment. First, we calculate a counterfactual average return on investment by using the investment distribution of the change treatment and the returns on investment observed for each investment level from 1 to 10 in the no-change treatment. The counterfactual return on investment is 56% and thus substantially higher than the 36% actually observed. Second, we regress return on investment in the two reputation treatments on a treatment dummy controlling for investment in an random-effects estimation with standard errors adjusted for clustering at session level. While the coefficient on investment is significant and positive, the treatment dummy remains economically and statistically significant in this regression. Detailed results of both analyses are available upon request.

¹⁰While some previous studies observe a negative or zero return on investment in trust games in a stranger environment (e.g., Keser 2003), the results of our control treatment are in line with those of a recent meta-study of 162 trust games regarding investment and return on investment (Johnson and Mislin 2011), which finds an average investment of 5.00 and a small, but significant positive return on investment of 0.11.

without reputation system. Our no-change treatment replicates this finding for an environment in which new players enter over time (U-test control vs. no-change: trust $p=0.01$; trustworthiness $p < 0.01$).

Table 1: Random-effects panel regressions of investment ((1) and (2)) and return on investment ((3) and (4)) on treatment dummies (“Control” and “Change”). “Last” is a dummy equal to 1 in period 20. *** and ** indicate significance on the 1-percent and 5-percent level respectively. Reported standard errors (in parentheses) are adjusted for clustering at session level.

	Investment		Return on investment	
	(1)	(2)	(3)	(4)
Control	-1.940** (0.76)	-2.068*** (0.73)	-0.461*** (0.09)	-0.478*** (0.09)
Change	-0.959* (0.53)	-0.998** (0.51)	-0.253** (0.11)	-0.246** (0.11)
Last		-2.115*** (0.54)		-0.457*** (0.14)
Control × Last		1.902*** (0.66)		0.335* (0.18)
Change × Last		0.581 (0.66)		-0.156 (0.19)
Constant	6.951*** (0.36)	7.094*** (0.34)	0.606*** (0.04)	0.629*** (0.05)
R^2	0.046	0.057	0.056	0.073
N	2352	2352	2019	2019

In view of result 1 the question arises whether a reputation system under which identity changes are possible still represents an improvement over a situation without any reputation system. The evidence on this is mixed. While average investment is higher in the change treatment compared to the control in all but two periods, the difference is not significant (U-test $p = 0.43$; Wald test for equality of coefficients in Table 1: column 1, $p = 0.20$; column 2, $p = 0.16$). This finding is unexpected in view of the models discussed above. Compared to the control treatment, return on investment in the change treatment is more than twice as high (36% vs. 16%). The difference is not statistically significant using a U-test ($p = 0.14$), but is (weakly) significant in the random-effects regressions (Wald test for equality of coefficients in Table 1: column 3, $p = 0.09$; column 4, $p = 0.05$). The evidence thus is at least suggestive that the reputation system with identity changes increases trustworthiness compared to an environment without any reputation system.

3.2 The impact of identity changes

The results presented above suggest that the ability of reputation systems to foster trust and trustworthiness seems to be negatively affected if sellers have the opportunity to change their identity. In the following, we analyze how reputation information

influences behavior and also discuss whether the treatment differences between the treatments with a reputation system can be traced back to the mechanisms discussed in section 2.

3.2.1 Do the basic principles of reputation systems work?

A first prerequisite for the functioning of a reputation system is that buyer ratings reflect seller behavior. In particular, opportunistic behavior should translate into a bad reputation. We use ordered probit regressions of return on investment on the rating given by the buyer for this return to investigate this question. We code a good rating as 1, a neutral rating as 0, and a negative rating as -1 . Table 2 reports the regression results. The coefficient for return on investment is positive and significant for both treatments, also if we control for the level of investment and a variable indicating how long a buyer has been in the experiment (“personal period”). In both treatments, sellers thus receive lower ratings for lower returns.

To gain an impression of potential differences in rating standards across treatments we calculate the return on investment thresholds which explain the highest percentage of ratings for each treatment assuming that buyers rate a return on investment below a threshold b_{neg} as negative, between b_{neg} and b_{neu} as neutral, and above b_{neu} as positive. The optimal b_{neg} is 0.19 in the change and 0.59 in the no-change treatment. The optimal threshold between a neutral and a good rating, b_{neu} , is 0.65 in the change and 0.88 in the no-change treatment. These thresholds predict around 75% of choices in both treatments.¹¹ Overall, rating standards seem to be a bit lower in the change treatment. Nevertheless, we observe a higher average rating in the no-change compared to the change treatment. While in the no-change treatment positive (negative) ratings account for 47% (30%) of all ratings, this figure is 41% (37%) in the change treatment. Using the coding of ratings described above, the average rating in the no-change treatment is 0.17, and 0.04 in the change treatment.

A second prerequisite for the success of a reputation system is that buyers condition their actions on the reputation information they receive, i.e., reward a good reputation with higher and punish a bad reputation with lower investments. Table 3 shows how the reputation profile a buyer sees at the beginning of a period influences his investment decision in the two treatments. The table reports the results of random-effects estimations with standard errors adjusted for clustering at session level. The results indicate that buyers take a seller’s reputation into account when deciding how much to invest. The coefficients for the number of good and bad ratings are significant and have the expected sign. Bad ratings seem to have a stronger impact. On average, one additional good rating increases investment by 0.17 in the no-change treatment whereas the coefficient for an additional negative rating is -0.52 (0.21 and -0.75 in the change treatment, column 1). A Wald-test confirms that the absolute values of the coefficients are significantly different (no-change: $p < 0.01$; change: $p < 0.01$). A stronger influence of negative ratings is in line with findings for laboratory data and field data from eBay.¹²

¹¹Since buyers did not make a decision for every potential return on investment only intervals can be determined for the b_{neg} and b_{neu} , which explain the highest percentage of choices. All $b_{neg} \in [0.58; 0.6]$ and all $b_{neu} \in [0.88; 0.88]$ are optimal in the no-change treatment (change: $b_{neg} \in [0.17; 0.2]$; $b_{neu} \in [0.63; 0.66]$). The values reported represent the mid points of these intervals.

¹²See Bajari and Hortaçsu (2004), Resnick et al. (2006) or Bolton et al. (2013) for references on the

Table 2: Buyer rating depending on seller return on investment, ordered probit estimates. “Personal Period” indicates how many periods a buyer has been in the experiment. *** and ** indicate significance on the 1-percent and 5-percent level respectively. Reported standard errors (in parentheses) are adjusted for clustering at session level.

	Dependent variable: Rating					
	no-change treatment			change treatment		
	(1)	(2)	(3)	(4)	(5)	(6)
Return on investment	2.271*** (0.409)	2.282*** (0.344)	2.363*** (0.371)	1.513*** (0.186)	1.510*** (0.190)	1.514*** (0.189)
Investment	-	0.099** (0.050)	0.105** (0.050)	-	0.053** (0.026)	0.0524* (0.027)
Personal Period	-	-	-0.028*** (0.006)	-	-	0.007 (0.014)
Log-likelihood	-510.19	-499.26	-495.05	-462.79	-458.62	-458.35
N	720	720	720	667	667	667

Neutral ratings seem to be perceived as a positive signal. The coefficient for neutral ratings is positive in both treatments, but only significant for the change treatment (column 1, Wald test, $p < 0.01$). Overall, the coefficients of the interaction terms indicate that ratings tend to have a slightly stronger (i.e., more positive for positive and neutral ratings, and more negative for negative ratings) influence in the change treatment. However, the difference is only significant for neutral ratings and this result is not robust to including a time trend (column 2). Note that the rating profile of a seller also listed last period’s rating separately. Including dummies for a positive, neutral or negative rating in the last period and the respective interaction terms does not qualitatively change the results reported above (column 3). Regarding the influence of the last rating we find that only a positive last rating has a significant additional impact on investment. However, this impact is substantial, as a positive last rating increases investment by almost 1.3.

Result 2: *The basic principles of reputation systems function in both treatments. Buyers give lower ratings for lower returns on investment. They also base their investment decision on the reputation profile of the seller. A worse reputation leads to lower buyer investments.*

Finally, we study whether a buyer’s personal experience matters for his investment decision in addition to the reputation information he receives. To this end we include “personal period” and a variable indicating how often a buyer made a loss on his investment in the past (“no. loss”) in the regression (column 4). The results indicate that own negative experiences in previous rounds lead to a significant and sizable reduction in buyer investment, especially in the no-change treatment. Opportunistic sellers thus

growing number of field studies on the influence of reputation information on outcomes in online trade.

create a negative externality. A similar “own history effect” has been found in Bolton et al. (2004).

Table 3: Buyer investment depending on seller reputation. “Change” is a dummy variable equal to 1 for the change treatment. “Good” is a variable containing the number of positive ratings in the rating profile of the seller a buyer was matched with, “Neutral”, and “Bad” are defined analogously. Variables with the prefix “Last” are dummy variables for a good, neutral, or bad rating in the previous period. “No. Loss” contains the number of times a buyer has made a loss so far. ***, **, and * indicate significance on the 1-, 5- and 10% level respectively. Reported standard errors (in parentheses) are adjusted for clustering at session level.

	Dependent variable: Investment			
	(1)	(2)	(3)	(4)
Change	-1.844*** (0.512)	-1.285*** (0.489)	-1.943*** (0.588)	-1.327** (0.513)
Good	0.165*** (0.027)	0.242*** (0.035)	0.109*** (0.027)	0.252*** (0.044)
Good × Change	0.049 (0.061)	0.023 (0.060)	-0.018 (0.047)	-0.005 (0.072)
Neutral	0.088 (0.067)	0.176** (0.075)	0.052 (0.075)	0.171** (0.085)
Neutral × Change	0.191* (0.113)	0.119 (0.114)	0.119 (0.091)	0.128 (0.120)
Bad	-0.525*** (0.065)	-0.406*** (0.066)	-0.419*** (0.069)	-0.396*** (0.075)
Bad × Change	-0.226 (0.161)	-0.248 (0.170)	-0.043 (0.132)	-0.262 (0.170)
Period	-	-0.115*** (0.030)	-	-0.102* (0.053)
Last_Good	-	-	1.282*** (0.315)	-
Last_Good × Change	-	-	0.889 (0.572)	-
Last_Neutral	-	-	0.705 (0.542)	-
Last_Neutral × Change	-	-	0.601 (0.638)	-
Last_Bad	-	-	-0.415 (0.069)	-
Last_Bad × Change	-	-	0.490 (0.462)	-
No. Loss	-	-	-	-0.831*** (0.211)
No. Loss × Change	-	-	-	0.579** (0.251)
Personal Period	-	-	-	0.046 (0.045)
Constant	7.355*** (0.280)	8.021*** (0.317)	6.826*** (0.393)	7.982*** (0.321)
N	1568	1568	1568	1520
R ²	0.187	0.206	0.223	0.222

3.2.2 Identity changes, the quality of new sellers, and consequences for buyer trust

We hypothesized above that investment towards new sellers would be lower in the change treatment. Note that the treatment dummy in the regression in column 1 of Table 3 represents the difference in investment towards sellers without any ratings between the no-change and the change treatment. As predicted players without a reputation receive a substantially lower investment in the change treatment.¹³

Result 3: *Investment towards new players is substantially lower in the change treatment.*

In the models of reputation systems with identity changes described in section 2.2, new players receive the same treatment as a player with the worst possible reputation. Changing identity is not a profitable strategy and no identity changes occur in equilibrium. We therefore look at how new sellers are treated compared to those with a bad reputation in the change treatment. We calculate the aggregate reputation of a seller by subtracting the number of bad ratings from the number of good ratings. In the change treatment, sellers with an aggregate reputation of -1 , i.e., with exactly one more negative than positive rating, received 3.29, those who had an equal number of positive and negative ratings received 4.89, and those with an aggregate reputation of 1 received 6.79. As a comparison, buyers invested 4.88 when they were facing a new player. According to this, admittedly coarse, analysis strategic sellers should thus switch their identity when their aggregate reputation reaches -1 . While identity changes are not readily observable in the field our data allow an analysis of whether and when sellers actually change their identity. On average, sellers erase their reputation profile every 3.5 rounds, i.e., in 28.4% of the rounds.¹⁴ An analysis of *when* sellers erase their reputation indicates that they understand that there is scope for a strategic use of identity changes. 63% of identity changes occur when the aggregate reputation is -1 , another 25% when it is 0.¹⁵

There is also tentative evidence for strategic behavior in terms of reputation dynamics. In view of the difference in investment towards new players and those with an aggregate reputation of 1 shown above, opportunistic players have an incentive to strategically build up a reputation in order to ‘milk’ this reputation later on. This holds for those instances, in which they can buy a good rating cheaply, i.e., when they receive a small investment. Two observations suggest that at least some sellers in the change treatment employ this strategy. First, there is a negative (though not significant) corre-

¹³Players without reputation include those players who entered as a new seller in the previous round but did not receive any investment. If we only consider players with the label “new” the difference in investment is even larger (see below).

¹⁴For the analysis of identity change behavior we exclude the first round of each seller since identity could not be changed in this round. 736 rounds are thus included in the analysis.

¹⁵The conditional probability of changing identity is also highest for an aggregate reputation of -1 . 86% of the sellers with an aggregate reputation of -1 change their identity. Note also that the substantial positive impact of a positive last rating on investment found above could make a strategy profitable where sellers split equally in one period to receive a positive rating, behave opportunistically in the next, and change their identity after having received a negative rating. This could explain why 25% of sellers change their identity at an aggregate reputation of zero.

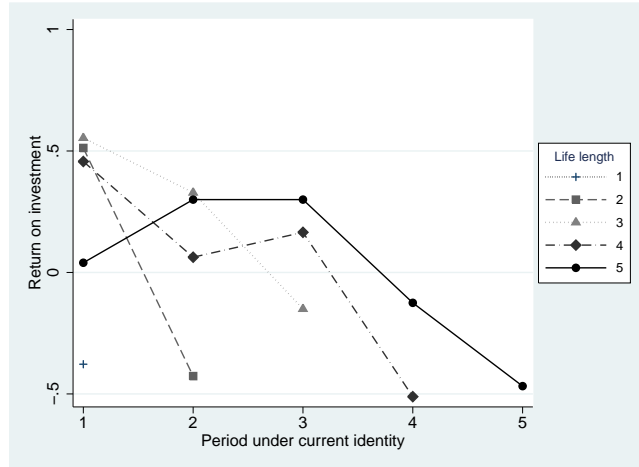


Figure 3: *Return on investment for identities of different length*

lation between investment and return on investment for new sellers (Spearman rank order correlation, $\rho = -0.11$, $p = 0.15$, $N = 181$), whereas this correlation is positive and significant overall ($\rho = 0.11$, $p = 0.01$, $N = 644$). Second, we study return on investment dynamics over the life span of a seller *identity*. For example, a seller who participates in the experiment right from the beginning and changes identity every fifth period has four identities of five periods each. Figure 3 shows average return on investment per period for identities of different length (1-5 periods). Return on investment in the first period of an identity is negative for identities which only last one period but starts out above zero for those who stay with their current identity for more than one period. In addition, we see positive returns on investment for the early periods of multi-period identities and a negative return for the last at about the level observed for one-period-identities. These findings are in line with anecdotal evidence from eBay.¹⁶

Finally, we confirm that sellers who behave opportunistically and then change their identity are driving the lower trustworthiness in the change treatment. Return on investment and the frequency of identity changes are negatively correlated on an individual seller basis (Spearman rank order correlation $\rho = -0.67$; $p < 0.01$; $N = 48$).¹⁷ Sellers can roughly be divided into two groups—those who change their identity at most once and exhibit comparatively high returns on investment, and those who change their identity more frequently and return little to the buyer.

Result 4: *Sellers in the change treatment use the opportunity to strategically change identities. Frequent identity changes are associated with low returns on investment.*

The previous analyses suggest that the composition of the pool of new players differs drastically between the treatments. In contrast to the no-change treatment, sellers with

¹⁶Jin and Kato (2006), for example, observe several sellers who build up their reputation first and then fail to deliver on a large number of parallel auctions in the eBay submarket for unrated baseball cards. See Liu (2011) for a model which produces similar reputation dynamics.

¹⁷The correlation is also significant if we take each session as an independent observation (Spearman rank order correlation $\rho = -0.90$; $p < 0.01$; $N = 8$).

the label “new” in the change treatment are a mix of truly new sellers (10%) and sellers who erased their rating profile (90%). Average return on investment by real new sellers is 40%.¹⁸ In contrast, the remaining 90% of the sellers labeled as “new”, i.e., those who changed their identity, display an average return on investment of –10%. Overall, we observe a significant difference in the behavior of sellers labeled as “new”, and those with a reputation profile (two-sided Wilcoxon test $p = 0.02$). While buyers on average make a loss when they interact with new sellers in the change treatment, interactions with sellers with a reputation profile yield a return on investment of 52%.

Taking the returns on investment presented in the previous paragraph as given it would be optimal for buyers in the change treatment not to invest anything when facing a new seller. However, while buyers in the change treatment do indeed invest less towards new sellers compared to the no-change treatment (4.88 vs 7.38; two-sided U-test, $p = 0.01$), the investment is still substantial.¹⁹ As shown above, sellers can therefore avoid the negative consequences of opportunistic behavior to some extent when identity changes are possible: new sellers on average receive a higher investment than those with a negative aggregate reputation and about the same as sellers with a neutral aggregate reputation. Compared to the no-change treatment, incentives to behave trustworthy are therefore lower. An important question in this respect is whether a reputation system with identity changes can still prevent that opportunistic sellers make a higher profit than trustworthy sellers. We therefore turn to the analysis of payoffs in the next section.

3.2.3 Payoffs

Buyers generate a substantial surplus in both treatments by investing more than the Nash prediction of 0, but the surplus generated is not distributed equally. In line with the trust game literature, buyers on average earn less than sellers in both treatments. Both, buyers and sellers, are nevertheless better off compared to their initial endowment. In the no-change treatment buyers earn 14.49 per round while sellers earn 19.58 (Wilcoxon test $p = 0.01$). The difference is even larger for the change treatment where the respective figures are 12.33 for buyers and 19.67 for sellers (Wilcoxon test $p = 0.01$). There is no significant difference between seller profits in the two treatments (U-test $p = 0.92$), but buyers earn significantly less when sellers can change identity (U-test $p = 0.04$). The efficiency losses arising from lower investments in the change treatment are thus entirely borne by the buyers.²⁰

For judging the performance of a reputation system it is important to consider the payoff consequences of different types of behavior. Ideally, a reputation system prevents that opportunistic agents make a higher profit than honest ones. In the models described above this desirable property is achieved in equilibrium. Our final analysis is therefore a simple test of whether bad behavior is profitable. We classify a seller as “bad” (“good”) if he has collected more bad than good ratings (more good than bad) in the

¹⁸Given that all players were new players at the start of the experiment and the label could thus not carry any information we exclude the first period from the analyses in this paragraph. The qualitative results are robust to including these observations.

¹⁹One reason why buyers might not treat new sellers in the worst possible way could be that they are guided by “preferences for appropriate choices” (Grosskopf and Sarin 2010).

²⁰This result also holds for the control treatment where buyers earn 11.07 and sellers earn 19.25.

course of the experiment.²¹ In the no-change treatment, good sellers earn 19.22 and bad sellers earn 19.75 on average, but the difference is not significant (Wilcoxon test $p = 0.21$). In the change treatment, good sellers earn 18.39 on average, bad sellers make 20.74. In contrast to the no-change treatment, this difference is significant (Wilcoxon test $p = 0.02$). The reputation system in the change treatment thus does not succeed in making good behavior at least as profitable as bad behavior. Since opportunistic players camouflage as new players, all new players receive lower investments. But the reduction for new sellers is not large enough to offset the lower returns by opportunistic sellers.²²

Result 5: *Opportunistic sellers do not earn more than sellers with high returns in the no-change treatment. However, opportunistic behavior does pay off in the change treatment.*

3.2.4 Robustness Check

Two different designs have been used to study reputation systems, the standard trust game (Keser 2003, and Masclet and Peenard 2012) and a binary variant called the buyer-seller game (Bolton et al. 2004). The main difference between the two games is that efficiency gains in the standard trust game arise once the buyer trusts while in the buyer-seller game they only arise when trust is honored. In addition, Bolton et al. (2004) abstract from a reputation system based on subjective ratings and instead provide the history of past actions of a seller. To check whether our results also hold in this environment we ran two additional buyer-seller control treatments ($N = 48$ for each treatment). Subjects either participated in the no-change or the change treatment, but, instead of the trust game, played the buyer-seller game from Bolton et al. (2004). Payoffs and move structure in this game are as follows. Buyers can choose buy/not buy, and sellers can choose ship/not ship if the buyer has decided to buy. If the buyer does not buy, both players earn 35. If the buyer buys and the seller ships, both earn 50. If the buyer buys and the seller does not ship, the buyer earns 0 and the seller 70. The feedback mechanism showed the history of a seller's actions.

To keep the setup in our buyer-seller control treatments as similar as possible to our original treatments, everything else was kept equal, i.e., number of participants per session, entry and timing of new players, number of rounds, and possibility to erase one's history in the change treatment. For the same reason buyers and sellers did not switch roles, in contrast to Bolton et al. (2004).

Figure 4 and Figure 5 in the appendix show the results for buyer and seller decisions. The main results from the trust game treatments are replicated in the buyer-seller game setup—trust and trustworthiness are lower in the change treatment. The average buy

²¹There are 20 (23) bad and 27 (24) good sellers in the no-change (change) treatment. For our analysis, we ignore the remaining seller in each treatment who ends the session with an equal number of good and bad ratings. Including them either with the good or the bad sellers does not change the results.

²²In view of this result one might ask why not more subjects choose to behave opportunistically and change their identity afterwards. Analyzing behavior in the last round suggests that a substantial fraction of sellers does not act purely out of strategic concerns. 11 out of the 34 sellers who receive an investment in the last period provide a return on investment for the buyer of 50% or more, and thus seem to have an intrinsic preference for trustworthy behavior.

rate is 81% in the no-change and 67% in the change treatment. The average shipping rate conditional on a buy decision is 88% in the no-change and 69% in the change treatment. The treatment difference for both measures is significant in a random-effects probit regression of the respective measure on a treatment dummy (Buy: $p = 0.01$; Ship: $p = 0.02$) or using a non-parametric U-test (Buy: $p = 0.06$; Ship: $p = 0.02$). Regarding the treatment of new sellers we also observe a qualitatively similar pattern in the buyer-seller control treatments. Buyers facing a new seller decide to buy in 96% of the cases in the no-change treatment, but only in 52% of the cases in the change treatment. Overall, the results from the control treatment clearly indicate that our main results hold for both the trust game and the buyer-seller game.

4 Conclusions

In this paper, we studied the influence of identity changes on the efficiency of reputation systems by comparing two markets. In both markets, new buyers and sellers enter over time, but in one market sellers cannot erase their reputation while in the other market they can. We find reduced seller trustworthiness and buyer trust in the latter market. A substantial fraction of sellers uses the opportunity to behave opportunistically and change their identity afterwards. Buyers cannot distinguish between these players and those who are really new and react with lower investments towards players with the label “new player” as predicted by (infinite-horizon) models of reputation systems with identity changes. However, this reduction in investment is not large enough to make opportunistic behavior with frequent identity changes an unprofitable strategy in the change treatment. Nevertheless, return on buyer investment is positive under the reputation system with identity changes and somewhat higher than in an environment without reputation system at all.

An important implication of our findings for online market design is that any measure which reduces the uncertainty connected with the label “new player” and thus increases the ability to distinguish between good and bad sellers can increase efficiency.²³ An obvious suggestion is to introduce fixed identifiers, another less invasive measure is to give sellers the opportunity to have their identity verified by a third party. However, users may value a certain degree of anonymity (Friedman and Resnick 2001). The results in this paper thus need not be interpreted as arguments against allowing identity changes in general. They rather suggest that doing so comes at a cost which has to be weighed against the potential benefits of increased anonymity. These are likely to depend on the environment in which the reputation system is to be used.

Our tentative finding of an improvement in trustworthiness between the control and the change treatment hints at the potential reputation systems offer even in environments in which identities can be changed. An interesting question for future research is whether introducing competition in the form of partner choice increases the efficiency of reputation systems in these environments. On the seller side competition is likely to create stronger incentives for trustworthy behavior (Bolton et al. 2008, Cassar and Rigdon 2011). Competition on the buyer side (Keck and Karelaia 2011) could lead to a greater buyer willingness to take risks with sellers whose reputation is not entirely posi-

²³Friedman and Resnick (2001) draw the same conclusion from their model and provide a discussion of different measures to achieve this.

tive. This seems to hold in particular for new buyers (Jin and Kato 2006). Which of these effects dominates is ultimately an empirical question. Our simple and parsimonious framework can successively be enriched to study partner choice and other measures such as verified IDs in future research.

References

- Bajari, Patrick and Ali Hortaçsu (2004): “Economic Insights from Internet Auctions.” *Journal of Economic Literature*, 42:457–486.
- Berg, Joyce E., John Dickhaut and Kevin McCabe (1995): “Trust, Reciprocity, and Social History.” *Games and Economic Behavior*, 10:122–142.
- Bohnet, Iris, Heike Harmgart, Steffen Huck and Jean-Robert Tyran (2005): “Learning Trust.” *Journal of the European Economic Association*, 3:322–329.
- Bolton, Gary E., Ben Greiner and Axel Ockenfels (2013): “Engineering Trust: Reciprocity in the Production of Reputation Information.” *Management Science*, 59:265–285.
- Bolton, Gary E., Elena Katok and Axel Ockenfels (2004): “How Effective are Online Reputation Mechanisms? An Experimental Investigation.” *Management Science*, 50:1587–1602.
- Bolton, Gary E., Claudia Loebbecke and Axel Ockenfels (2008): “Does Competition Promote Trust and Trustworthiness in Online Trading? An Experimental Study.” *Journal of Management Information Systems*, 25:145–170.
- Brown, Martin, Armin Falk and Ernst Fehr (2004): “Relational Contracts and the Nature of Market Interactions.” *Econometrica*, 72:747–780.
- Camerer, Colin and Keith Weigelt (1988): “Experimental Tests of a Sequential Equilibrium Reputation Model.” *Econometrica*, 56:1–36.
- Cassar, Alessandra and Mary Rigdon (2011): “Trust and Trustworthiness in Networked Exchange.” *Games and Economic Behavior*, 71:282–303.
- Dellarocas, Chrysanthos and Charles A. Wood (2003): “The Digitization of Word of Mouth: Promise and Challenges of Online Feedback Mechanisms.” *Management Science*, 49:1407–1424.
- Dellarocas, Chrysanthos and Charles A. Wood (2008): “The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias.” *Management Science*, 54:460–476.
- European Commission (2004): *Issues Relating to Business and Consumer E-commerce*. Special Eurobarometer 60.0/Wave 201, European Opinion Research Group, Brussels.
- European Commission (2012): *Cyber Security*. Special Eurobarometer 390/Wave EB 77.2, TNS Opinion & Social, Brussels.
- Fischbacher, Urs (2007): “z-Tree: Zurich Toolbox for Ready-made Economic Experiments.” *Experimental Economics*, 10:171–178.
- Friedman, Eric J. and Paul Resnick (2001): “The Social Cost of Cheap Pseudonyms.” *Journal of Economics & Management Strategy*, 10:173–199.

- Gazzale, Robert S. and Tapan Khopkar (2011): "Remain silent and ye shall suffer: seller exploitation of reticent buyers in an experimental reputation system." *Experimental Economics*, 14:273–285.
- Greif, Avner (1989): "Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders." *Journal of Economic History*, 49:857–882.
- Grosskopf, Brit and Rajiv Sarin (2010): "Is Reputation Good or Bad? An Experiment." *American Economic Review*, 100:2187–2204.
- Jin, Ginger Zhe and Andrew Kato (2006): "Price, quality, and reputation: evidence from an online field experiment." *The RAND Journal of Economics*, 37:983–1005.
- Johnson, Noel D. and Alexandra A. Mislin (2011): "Trust games: A meta-analysis." *Journal of Economic Psychology*, 32:865–889.
- Kandori, Michihiro (1992): "Social Norms and Community Enforcement." *Review of Economic Studies*, 59:63–80.
- Keck, Steffen and Natalia Karelaia (2011): "Does competition foster trust? The role of tournament incentives." *Experimental Economics*, 14:273–285.
- Keser, Claudia (2003): "Experimental games for the design of reputation management systems." *IBM Systems Journal*, 43:498–506.
- Klein, Tobias J., Christian Lambertz, Giancarlo Spagnolo and Konrad O. Stahl (2009): "The actual structure of eBay's feedback mechanism and early evidence on the effects of recent changes." *International Journal of Electronic Business*, 7:301–320.
- Kreps, David, Paul Milgrom, John Roberts and Robert Wilson (1982): "Rational cooperation in the finitely repeated prisoners' dilemma." *Journal of Economic Theory*, 27:245–52.
- Liu, Qingmin (2011): "Information Acquisition and Reputation Dynamics." *Review of Economic Studies*, 78:1400–1425.
- Masclot, David and Thierry Peenard (2012): "Do reputation feedback systems really improve trust among anonymous traders? An experimental study." *Applied Economics*, 44:4553–4573.
- Milgrom, Paul R., Douglass C. North and Barry R. Weingast (1990): "The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs." *Economics & Politics*, 2:1–23.
- Ockenfels, Axel (2003): "Reputationsmechanismen auf Internet-Marktplattformen: Theorie und Empirie." *Zeitschrift für Betriebswirtschaft*, 73:295–315.
- Okuno-Fujiwara, Masahiro and Andrew Postlewaite (1995): "Social Norms and Random Matching Games." *Games and Economic Behavior*, 9:79–109.
- Reichling, Florian (2004): "Effects of Reputation Mechanisms on Fraud Prevention in eBay Auctions." *Stanford University Working Paper*.

Resnick, Paul and Richard Zeckhauser (2002): "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System." *The Economics of the Internet and E-Commerce*, 11:127–157.

Resnick, Paul, Richard Zeckhauser, John Swanson and Kate Lockwood (2006): "The value of reputation on eBay: A controlled experiment." *Experimental Economics*, 9:79–101.

Appendix A

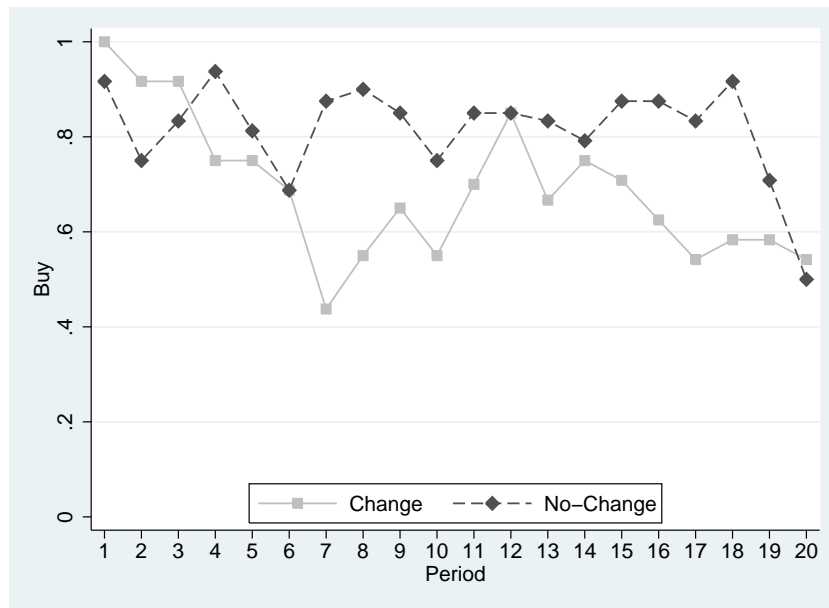


Figure 4: Relative frequency of buy decisions in the buyer-seller treatments

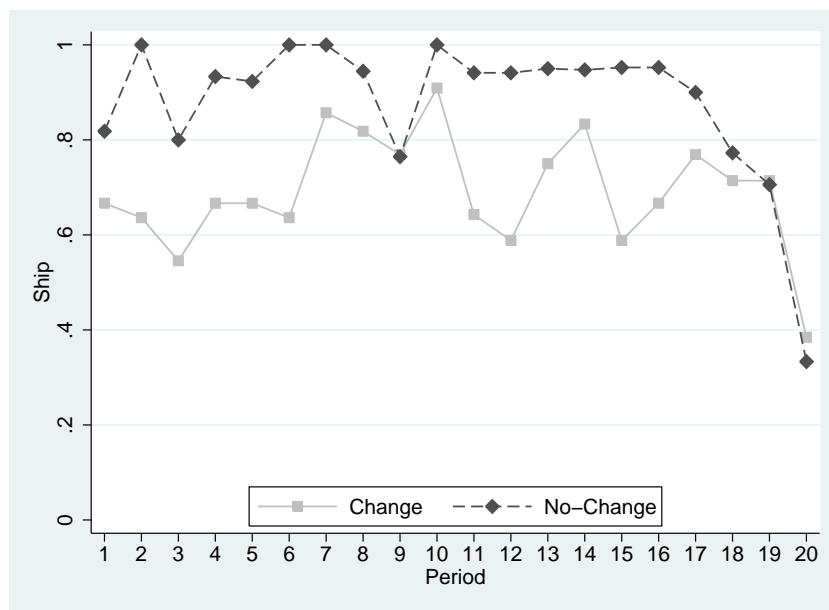


Figure 5: Frequency of ship decisions in the buyer-seller treatments conditional on buying

Treatment	Session	Inv	Inv_Half1	Inv_Half2	ROI	No_Inv	Inv_10	Good	Neutr	Bad
Control	1	8.22	7.45	8.76	.45	2	1	-	-	-
Control	2	1.24	1.9	.79	-.19	60	7	-	-	-
Control	3	5.93	7.1	5.12	.35	15	5	-	-	-
Control	4	5.21	5.72	4.86	.28	14	7	-	-	-
Control	5	6.12	6.1	6.14	.15	1	2	-	-	-
Control	6	4.64	5.03	4.38	.1	13	4	-	-	-
Control	7	4.86	6.05	4.03	.05	19	0	-	-	-
Control	8	5.07	5.2	4.98	-.19	28	4	-	-	-
No-Change	9	7.27	8.25	6.59	.47	9	0	31	21	37
No-Change	10	8.27	8.45	8.14	.58	6	0	45	27	20
No-Change	11	8.47	8.25	8.62	.74	1	0	54	23	20
No-Change	12	5.59	6.47	4.98	.49	22	4	34	13	29
No-Change	13	6.71	7.88	5.91	.55	8	5	36	26	28
No-Change	14	7.46	6.6	8.05	.81	4	1	48	21	25
No-Change	15	6.13	6.8	5.67	.5	12	1	30	22	34
No-Change	16	6.4	6.8	6.12	.7	2	7	60	12	24
Change	17	4.98	5.3	4.76	.08	23	0	26	12	37
Change	18	5.82	5.9	5.76	.11	13	3	26	17	42
Change	19	4.82	5.5	4.34	.55	15	7	33	23	27
Change	20	5.07	6.07	4.38	.26	23	4	18	19	38
Change	21	7.51	7.7	7.38	.81	3	1	63	20	12
Change	22	7.16	7.25	7.1	.46	10	0	39	18	31
Change	23	5.28	5.53	5.1	-.06	23	2	23	17	35
Change	24	7.38	8.75	6.43	.55	7	2	48	16	27

Table 4: Descriptive statistics by session. Inv_Half1 and Inv_Half2 refer to investment in periods 1-10 and 11-20, respectively. No_Inv and Inv_10 state how often an investment of 0 or 10, respectively was made. Good, Neutr and Bad describe how often the respective rating was given in a certain session.

Appendix B: Instructions for the change treatment

In what follows, we present a translation of the instructions for players in the change treatment. The instructions for the no-change treatment differed only in the paragraph regarding the possibility to erase one's rating profile.

Instructions

There are two types of players in this experiment, **player A** and **player B**. You will be assigned the role of either player A or player B, and you will keep your assigned role throughout the entire experiment.

The experiment lasts **20 rounds**. At the beginning of each round the number in the upper part of the screen indicates the current round. In each round a player A is matched with a player B. This matching process is random. However, you will **never** be matched **with the same player two rounds in a row**.

Structure of each round:

At the beginning of each round each player is endowed with 10 Taler. Player A then decides how many Taler of his endowment he wants to send to player B. Player A can send any (integer) amount between 0 and 10 Taler.

The experimenter triples this amount, so that player B receives the tripled amount of what player A decided to send.

Player B then decides how many Taler he wants to send back to player A. He can send any (integer) amount between 0 Taler and the tripled amount transferred by player A. The amount of Taler that player B sends back to player A is **not tripled** by the experimenter.

Rating:

After both players made their decisions, player A can **rate** player B's decision as "**good**", "**neutral**" or "**bad**". The rating is stored in the rating profile. There is no rating stage if player A sends zero Taler to player B. In that case the phrase "No rating, because no exchange took place" is displayed on the screen.

Before player A decides on how much to send to player B, the **rating profile** of the player B he is matched with is displayed on the screen. The rating profile includes the rating in the previous round as well as the total number of good, neutral and bad ratings that player B has received so far.

Deleting the rating profile:

At the beginning of each round player B can **delete his existing rating profile**, and start over as a **new player**.

New players:

During the course of the experiment new players enter the experiment. Players which enter the experiment in later rounds find paper and pen at their desks and have to work on a different task before starting as a player in the computerized experiment. These players will also be paid for working on the non-computerized task at the end of the experiment. They will find a more detailed description of the task at their desks. These players will start with their decisions in the computerized experiment as soon as the screen in their cubicle displays an entry mask.

The label "New player" is used for players who deleted their rating profile, as well as for players who have just entered the experiment.

Questionnaire:

After the last round of the experiment a short questionnaire is displayed on the screen. Please answer the questionnaire as precisely as possible. When you are done with the questionnaire please wait in your cubicle until we ask you to collect your payment.

Payment: After the experiment the amount of Taler you earned is exchanged into Euros at an **exchange rate of 3,5 Cent/Taler**.