

LEARNING TRUST

Iris Bohnet

Harvard University

Heike Harmgart

University College London and IFS

Steffen Huck (UCL)

University College London

Jean-Robert Tyran

University of Copenhagen

Abstract

We examine the effects of different forms of feedback information on the performance of markets that suffer from moral hazard problems due to sequential exchange. As orthodox theory would predict, we find that providing buyers with information about sellers' trading history boosts market performance. More surprisingly, this beneficial effect of incentives for reputation building is considerably enhanced if sellers, too, can observe other sellers' trading history. This suggests that two-sided market transparency is an important ingredient for the design of well-functioning markets that are prone to moral hazard. (JEL: C72, C91, L14)

1. Introduction

Reputation building in repeated trust games requires that trustors have some information about trustees' behaviour in the past. Consider a buyer–seller framework where sequential exchange induces a moral hazard problem. First, buyers make a decision about whether or not to send some money to a seller who has advertised a good. After having received the money, the seller then decides whether or not to deliver the promised good. In such a market a seller can build up a reputation for being honest if and only if buyers can at least partially observe the seller's trading history.

Thus, providing buyers with information about sellers' past should help to alleviate the moral hazard problem. We shall call such feedback provision to buyers one-sided market transparency. The first result that we establish in this paper is that it indeed helps to improve efficiency in laboratory markets that suffer from moral hazard. However, our key finding is that two-sided market

Acknowledgments: We are grateful to Arno Riedl for a very helpful discussion at the EEA meetings in Madrid. The third author acknowledges financial support from the Leverhulme Trust and the Economic and Social Research Council (UK) via ELSE.

E-mail addresses: Bohnet: iris_bohnet@harvard.edu; Harmgart: h.harmgart@ucl.ac.uk; Huck: s.huck@ucl.ac.uk; Tyran: Jean-Robert.Tyran@econ.ku.dk

Journal of the European Economic Association April–May 2005 3(2–3):322–329
© 2005 by the European Economic Association

transparency where both, buyers and sellers, have access to sellers' trading history improves market performance even further. From the vantage point of orthodox theory, this is a surprising result. Whether or not sellers can observe other sellers' past should be irrelevant. But as we conclusively show, it is not.

The key to understanding this result is very simple. There are some sellers who, when left to their own devices, simply do not understand the mechanics of reputation building. In markets with one-sided transparency only, they make use of any opportunity to rip off their customers despite the drastic consequences this implies for their reputation. Typically, it does not take long until such sellers establish a firm reputation as cheats and lose all business. This is different in markets with two-sided transparency because here sellers can learn from other sellers. In particular, sellers who initially do not understand the incentives for reputation building can now observe others who do. And they can see that those who do, get more business and are soon much better off than they are. Given a second chance, they can now imitate successful reputation building. This process of social learning gave this paper its title.

While the benefits of one-sided market transparency have already been documented in the literature,¹ the interaction of social learning and reputation incentives that makes two-sided transparency superior in our experiment has, to the best of our knowledge, not been demonstrated before.² Of course, the often cited example of eBay's feedback mechanism is one that implements two-sided transparency. On eBay, everybody, buyers and sellers, has access to information about sellers' histories. However, previous studies have—probably guided by orthodox reasoning—ignored the role of providing sellers with information about each other. Our results suggest that, in fact, two-sided transparency is an important ingredient for the design of well-functioning markets that are prone to moral hazard.

2. Experimental Design and Procedures

In our experiments, subjects play the binary-choice trust game shown in Figure 1. Payoffs are in pence and strategies and player roles are labeled exactly as in the experiment. Assuming that players maximise some monotone function of their monetary payoff and that this is common knowledge, the game has a unique Nash equilibrium, in which the first mover chooses “X”, i.e., not to trust, and the second mover chooses “left”, i.e., not to honour trust if being trusted. In the following we will refer to the first mover as the buyer and to the second mover as the seller.

1. See, for example, Keser (2002), Bolton, Katok, and Ockenfels (2004), or Bohnet and Huck (2004).

2. Studies that show how subjects can learn from other subjects to improve their decision making in other contexts include Offerman and Sonnemans (1998) and Slembeck and Tyran (2004).




FIGURE 1. The trust game.

Payoffs are deliberately chosen to be asymmetric in order to make the moral hazard problem as difficult as possible.³ Subjects play this game in all treatments for 30 periods. Keeping their roles they are randomly rematched at the start of each period. Each matching group consists of four sellers and four buyers.

The treatments differ in what subjects know about the past. In the baseline treatment, NOINFO, subjects have no information about the past. Whenever they are rematched, they are simply told “You have been rematched with a new participant” without knowing anything about this participant’s identity or history. In all other treatments sellers can be identified with labels (B1, B2, B3, and B4). In treatment REPUTATION, all buyers know all sellers’ pasts. In treatment IMITATION all sellers know each other’s past. And, finally, in treatment TWO-SIDED both, sellers and buyers, can observe all sellers’ pasts.⁴ This 2×2 design is summarised in Table 1.

The experiments were computerised⁵ and sellers’ histories were made available using a simple graphical tool. In the left part of the screen subjects could see

TABLE 1. The 2×2 design.

	Sellers know sellers’ histories	
	No	Yes
Buyers know	No	NOINFO
Sellers’ histories	Yes	REPUTATION
		IMITATION
		TWO-SIDED

3. With symmetric payoffs after honoured trust (Y, right) subjects find it much easier to achieve efficiency already in one-shot games, see, for example, Bacharach, Guerra, and Zizzo (2001) or Bolton, Katok, and Ockenfels (2004).

4. In each treatment, the information structure is publicly known. For example, in REPUTATION, both, buyers and sellers know that buyers can observe sellers’ pasts while sellers cannot.

5. We used Fischbacher’s (1999) z-tree.

TABLE 2. Average honour, trust, and efficiency rates in all four treatments.

	NOINFO	REPUTATION	IMITATION	TWO-SIDED
Honour rate	0.19 (0.11)	0.44 (0.26)	0.19 (0.16)	0.62 (0.23)
Trust rate	0.21 (0.15)	0.31 (0.17)	0.14 (0.11)	0.43 (0.15)
Efficiency rate	0.05 (0.05)	0.17 (0.14)	0.04 (0.04)	0.29 (0.18)

Note: Standard deviations in parentheses.

four columns consisting of 30 hash signs, each column representing one seller. Each row represented one period. Initially, all hash signs were white. Then, after each period, hash signs in the row representing this period changed their colour. They turned black if the seller had not to make a decision because his buyer did not trust him. They turned green if the seller honoured the buyer's trust. And they turned red if the seller exploited the buyer's trust. This colour coding is, of course, obvious and makes it rather easy to read the comparatively complex history information.

The experiments were conducted at the University of London. For each of the four treatments we conducted six separate sessions, each with eight subjects who had been recruited via e-mails to the college's entire student body. Altogether, 192 subjects participated in the experiments, which lasted on average less than an hour. Average earnings were £11.07 (including a £5 show-up fee).

3. Results

3.1. A Static View

Table 2 shows, for each treatment, average honour rates, i.e., the average frequency with which sellers honour buyers' trust, average trust rates, i.e., the average frequency with which buyers trust sellers, and, finally, average efficiency rates, i.e., the average frequency with which subjects play (Y, right) and reach the individually rational efficient outcome (which, from here on, we shall simply call the "efficient outcome" or refer to it as "efficient trade").⁶ Figures 2, 3, and 4 show the same information graphically and, in addition, honour, trust, and efficiency rates for, both, the best and the worst session in each treatment. A few observations are in order.

1. In treatments without incentives for reputation building (NOINFO and IMITATION) honour rates are very low (below 20%) and so are trust rates. Consequently, there is hardly any efficient trade.

6. Notice that the sum of payoffs is equal in both end nodes that can be reached after the first mover decision to trust. Again, this is a feature that makes it harder for subjects to cooperate. See also footnote 3.




FIGURE 2. Honour rates in all treatments.

2. Introducing incentives for reputation building in treatment REPUTATION more than doubles the average honour rate, which also boosts the trust rate. As a result the number of efficient trades is more than tripled. However, there is considerable variance between sessions and the overall outcome is far from perfect.
3. The reputation effects are considerably enhanced in treatment TWO-SIDED where both, buyers and sellers, have access to sellers' trading history. Again, there are considerable differences between sessions.

Conducting statistical tests⁷ reveals that the effects of introducing incentives for reputation building are highly significant.⁸ However, despite the consistently

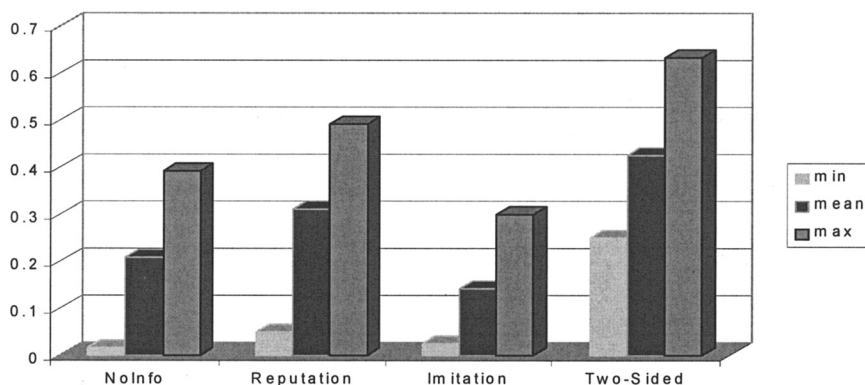


FIGURE 3. Trust rates in all treatments.

7. We always take one session as one independent observation and then perform pairwise MWU-tests with six against six observations.

8. Comparing NOINFO with REPUTATION we find that, both, honour and efficiency rates are significantly higher in the latter (one-sided $p = 0.023$ and $p = 0.074$, respectively). Comparing IMITATION

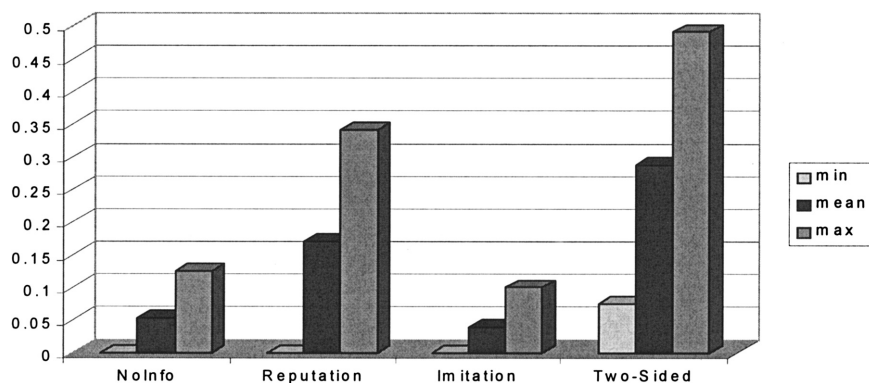


FIGURE 4. Efficiency rates in all treatments.

higher averages in treatment TWO-SIDED, tests fail to show any significant benefits of two-sided market transparency. Comparing treatments REPUTATION and TWO-SIDED, we find *no* significant differences—neither for honour rates, nor for trust and efficiency rates. Does this mean that there are indeed no effects of added market transparency and that, as orthodox theory predicts, it only matters whether or not buyers can observe sellers? In the next subsection we shall argue that this conclusion would be premature.

3.2. A Dynamic View

Two-sided market transparency has the advantage that sellers who do not understand the mechanics of reputation building can learn from other sellers who do. If such learning is important one would predict that two-sided transparency *crowds in* honour, trust, and efficiency over time.⁹ To test for such dynamic effects we will, therefore, analyse a very simple measure capturing the market dynamics. For each session we shall compute the difference between the average honour (trust/efficiency) rate over time and the initial honour (trust/efficiency) rate. In the case of trust and efficiency, these initial rates are simply computed for the first round. This approach does not work for initial honour rates since there are many sellers who do not have to make a decision in their first round. Hence to compute initial honour rates, we take for each seller the first instance where he or she had a decision to make.¹⁰ Table 3 shows the differences between average and initial rates for all four treatments.

and TWO-SIDED reveals that all three rates are higher when buyers know sellers' history ($p = 0.005$ for trust, $p = 0.008$ for honour, and $p = 0.004$ for efficiency).

9. Bohnet, Frey, and Huck (2001) provide a theoretical model for crowding in of trustworthiness.

10. Notice that these measures, while somewhat crude, are extremely clean. In particular, more sophisticated measures for the initial propensities to trust or honour would unavoidably be

TABLE 3. Crowding effects in all four treatments.

	NOINFO	REPUTATION	IMITATION	TWO-SIDED
Honour crowding effect	0.11 (0.17)	0.05 (0.31)	0.19 (0.16)	0.33 (0.49)
Trust crowding effect	-0.04 (0.11)	-0.06 (0.28)	-0.03 (0.15)	0.22 (0.26)
Efficiency crowding effect	0.01 (0.07)	0.00 (0.21)	0.04 (0.04)	0.20 (0.23)

Note: Crowding effect is computed as average rate minus initial rate. (Standard deviations in parentheses.) Notice that the honour crowding effect in treatment IMITATION is due to an initial honour rate of zero.

The table reveals rather dramatic effects of two-sided transparency. The average honour rate is 33 percentage points higher than the initial rate. This basically amounts to one-third of the seller population learning that building up a good reputation pays. Put differently, it amounts to every third seller turning from a cheat into a reliable trading partner.

In treatment REPUTATION there is also a slight increase in honour rates but, the effect is both, smaller and without consequences for overall market performance. Under two-sided market transparency increasing honour rates translate into increasing trust and, hence, increasing efficiency rates. In treatment TWO-SIDED overall efficiency is 20 percentage points higher than initial efficiency while in all other treatments efficiency does virtually not change over time.

Not surprisingly, these dynamic effects of two-sided market transparency are not only strong in size but also highly significant. A comparison of treatments REPUTATION and TWO-SIDED tests for the additional benefit of sellers observing other sellers in the presence of incentives for reputation building. Pairwise tests reveal that all three crowding effects are significantly higher in TWO-SIDED than in REPUTATION.¹¹ Thus, we see that two-sided market transparency has indeed an important beneficial effect for market performance that could not have been predicted by orthodox theory.

4. Conclusion

We examine the effects of different forms of feedback information on the performance of markets that suffer from moral hazard problems due to sequential exchange. We find that, as orthodox theory predicts, providing buyers with information about sellers' trading history boosts market performance. With such

confounded with learning or other dynamic effects. Also, by taking simply one dynamic measure we avoid making any assumptions about the functional form of the dynamics.

11. The p -values are $p = 0.055$ for the crowding in of trust and honour and $p = 0.075$ for the crowding in of efficiency (one-sided MWU-tests).

one-sided market transparency sellers have an incentive to build up a reputation as reliable trading partners and many sellers use this opportunity, which helps to alleviate moral hazard. This beneficial effect of incentives for reputation building is considerably enhanced if sellers, too, can observe other sellers' trading history. Apparently some sellers do not understand the mechanics of reputation building on their own. However, with two-sided market transparency these sellers can learn from those who manage to build a good reputation. Thus, there is a systematic learning process turning sellers who initially cheat into reliable trading partners. This dramatically increases market performance over time.

This result adds to the existing literature on feedback information, some of which has been motivated by the success of eBay's celebrated feedback mechanism. It suggests that eBay benefits from having sellers' feedback rating freely available to both market sides. While this was perhaps a very natural design choice for eBay where people act as both, buyers and sellers, this might be less obvious for specialised trading and procurement platforms where the two market sides are more separated. Here two-sided transparency might be a less obvious but—as our findings suggest—very recommendable choice.

References

- Bacharach, Michael, Gerardo Guerra, and Daniel Zizzo (2001). "Is Trust Self-Fulfilling? An Experimental Study." Working paper, University of Oxford.
- Bohnet, Iris, Bruno S. Frey, and Steffen Huck (2001). "More Order with Less Law: On Contract Enforcement, Trust, and Crowding." *American Political Science Review*, 95, 131–144.
- Bohnet, Iris and Steffen Huck (2004). "Repetition and Reputation: Implications for Trust and Trustworthiness When Institutions Change." *American Economic Review*, 94, 362–366.
- Bolton, Gary, Elena Katok, and Axel Ockenfels (2004). "How Effective Are Electronic Feedback Mechanisms? An Experimental Investigation." *Management Science*, 50, 1587–1602.
- Fischbacher, Urs (1999). "z-tree. Toolbox for Readymade Economic Experiments." IEW Working Paper 21, University of Zurich.
- Keser, Claudia (2002). "Trust and Reputation Building in e-Commerce." Working paper, Cirano Montreal.
- Offerman, Theo and Joep Sonnemans (1998). "Learning by Experience and By Imitating Successful Others." *Journal of Economic Behavior and Organization*, 34, 559–575.
- Slembeck, Tilman and Jean-Robert Tyran (2004). "Do Institutions Promote Rationality? An Experimental Study of the Three-Door Anomaly." *Journal of Economic Behavior and Organization*, 54, 337–350.