

Competition Fosters Trust

STEFFEN HUCK, GABRIELE K. LÜNSER AND JEAN-ROBERT TYRAN

October 2007

Abstract

We study the effects of reputation and competition in a stylized market for experience goods. If interaction is anonymous, such markets perform poorly: sellers are not trustworthy, and buyers do not trust sellers. If sellers are identifiable and can, hence, build a reputation, efficiency quadruples but is still at only a third of the first best. Adding more information by granting buyers access to all sellers' complete history has, somewhat surprisingly, no effect. On the other hand, we find that competition, coupled with some minimal information, eliminates the trust problem almost completely.

Keywords: Experience Goods; Competition; Reputation; Trust; Moral hazard; Information conditions

JEL Codes: C72, C92, D40, L14

Authors

Steffen Huck

Department of Economics & ELSE
University College London
Gower Street
London WC1E 6BT
United Kingdom
phone: +44-207-679-5895
e-mail: s.huck@ucl.ac.uk

Gabriele K. Lünser

Department of Economics & ELSE
University College London
Gower Street
London WC1E 6BT
United Kingdom
phone: +44-207-679-5852
e-mail: g.luenser@ucl.ac.uk

Jean-Robert Tyran

Department of Economics
University of Copenhagen
Studiestræde 6, DK-1455 Copenhagen
and
CEPR, Goswell Road, London
phone: +45-353-23-027
e-mail: jean-robert.tyran@econ.ku.dk

Huck acknowledges financial support from the Economic and Social Research Council (UK) via ELSE and a grant on "Competition and Trust", Lünser acknowledges financial support from the University of Erfurt, Tyran acknowledges financial support from the Department of Economics at the University of Copenhagen.

1. Introduction

Lack of trust and moral hazard are widely seen as serious impediments to economic efficiency and growth. In this paper, we study two remedies against these obstacles: reputation and competition. Using experimental methods we isolate the partial effects of both and obtain surprisingly unambiguous results. Identification of sellers, as a prerequisite for reputation building, has strong positive effects but fails to overcome the moral hazard problem. Adding more information by granting buyers access to all sellers' complete history has, somewhat surprisingly, no effect. With full public information efficiency remains at just around 30% of the first best. On the other hand, we find that competition, coupled with minimal information, eliminates the trust problem completely. As long as buyers can *identify* sellers, competition is *sufficient* to achieve almost full efficiency.

We implement a repeated binary trust game that resembles a market for an experience good with a fixed price where the seller can only choose quality.¹ We vary the amount of information that buyers can access about sellers' past quality choices and study treatments with and without competition. Without competition, buyers are in each period randomly assigned to sellers. With competition, buyers *choose* in each period the seller from whom they want to buy. It turns out that the introduction of competition is greatly effective. With choice of sellers market efficiency rises from a meager 30% to over 80%.

Regarding the market's informational structure, we start with treatments where buyers only know the *identity* of different sellers and remember their own experiences with sellers from whom they have bought in the past.² We refer to these treatments as treatments with "private information". On the other hand, we study treatments with "full information" where buyers know *all* past quality choices of *all* sellers. With private information, sellers can build some limited reputations vis-à-vis their own past customers. With full information, sellers can build much stronger reputations vis-à-vis *all* potential customers. However, somewhat

¹ Similar games have previously been used to investigate issues of trust by GÜTH, OCKENFELS, and WENDEL (1997), BOHNET, FREY and HUCK (2001), ANDERHUB, ENGELMANN, and GÜTH (2002), BOLTON, KATOK, and OCKENFELS (2004), BOHNET and HUCK (2004), and BOHNET, HARMGART, HUCK, and TYRAN (2005). The binary structure of the game lends itself to a much simpler interpretation than the continuous "investment game" popularized by BERG, DICKHAUT, and MCCABE (1995). For how trust measures gained from experiments correlate with survey data, see GLAESER, LAIBSON, SCHEINKMAN, and SOUTTER (2000).

² Knowledge of sellers' identities requires essentially that sellers have the same name in all periods. Notice also that if sellers' names were changing, choice and random matching would effectively be the same.

surprisingly, we find that introducing full information has no effect — neither with nor without competition.

While we find very strong effects of competition and very little effect of improved possibilities for reputation building, it is worthwhile to note that in our environment (where prices are fixed) some reputation is necessary for competition. In the absence of any information about sellers' histories, a buyer's choice of seller would necessarily be random and competition would, hence, be meaningless. Indeed, in our environment competition takes place *via* reputations. In some sense, this makes our second finding — that the strength or richness of reputations has no effect on market outcomes — more remarkable.

The effect of different forms of feedback information, i.e., the information subjects have about the past, has been studied in a small set of previous experiments. DUFFY and FELTOVICH (2002, 2006) study how information about previous-round behavior affects play in 2x2 games and compare such information about the past with (cheap talk) promises about the future. Unpredicted by theory, both double efficiency levels in prisoners' dilemma games. In the context of trust games, feedback mechanisms and reputation building have been studied by KESER (2002) and BOLTON, KATOK, and OCKENFELS (2004). Both studies show that the provision of feedback about sellers' past has substantial positive effects on efficiency. Similar observations are made by BOHNET and HUCK (2004) who study changes in informational institutions as well as by BOHNET, HARMGART, HUCK, and TYRAN (2005) who examine how sellers can learn the benefits of reputation building from each other. However, all these studies employ exogenous matching schemes and do not touch the issue of competition.³

Experiments with endogenous matching are quite rare.⁴ PAGE, PUTTERMAN, and UNEL (2005) study endogenous group formation in symmetric public good games. Knowing past average contributions subjects could ever so often express preferences about with whom they would like to play which were then fed into a matching algorithm. Endogenous matching led

³ Varying feedback information has also been used to study different learning rules. See, for example, HUCK, NORMANN, and OECHSSLER (1999) and OFFERMAN, POTTERS, and SONNEMANS (2002) for Cournot games; HUCK, NORMANN, and OECHSSLER (2000) for price competition; and DUFFY and HOPKINS (2005) for market entry games.

⁴ Of course, one might view double auctions as games with endogenous matching (to exchange a good) in which case the literature would be vast. However, we have in mind a narrower class of settings where prior to playing a well-defined n -player game, agents decide with whom they want to play this game. See JACKSON and WATTS (2005) for a theoretical approach to this class of games.

to sorting of subjects and significantly increased contributions and efficiency.⁵ RIEDL and ULE (2002) show, quite similarly to PAGE *et al.*, that voluntary association improves efficiency in dilemma games. In CORICELLI, FEHR, and FELLNER (2004) partners for playing 2-player public good games were auctioned off. Despite free riders often making high bids for getting the right to play with a cooperator, overall efficiency increased. In BOHNET and KÜBLER (2005) subjects could bid for playing in an “improved” version of a prisoners’ dilemma where costs of unilateral cooperation were reduced. This, too, led to some sorting and, at least initially, higher rates of cooperation in the prisoners’ dilemma with insurance. There is also a small set of related papers studying experiments on endogenous network formation where subjects form links to play subsequently some simple games with all their “neighbors;” see KOSFELD (2004) for a survey.⁶

While we examine competition in a market suffering from moral hazard, CABRALES, CHARNESS, and VILLEVAL (2006) study the effects of choice of partners in a hidden-information framework. However, the paper that is most closely related to ours is BROWN, FALK, and FEHR (2004) who study endogenous relational contracts in markets that rely on gift-exchange and, thus, suffer from moral hazard.⁷ They observe the endogenous emergence of long-term relations with generous rent sharing and high effort. There are, however, a number of crucial differences between their study and ours. First, in their study prices are endogenous which prevents them from identifying the pure effect of freedom of choice on which we focus here. Second, there is no random matching benchmark as agents’ matching is *always* endogenous. Again, this implies that they cannot say anything about the effect of freedom of choice. Third, agents from both market sides have to form *pairs* which removes an important aspect of competition in many markets, namely competition for market shares. Fourth, they do not examine different informational environments but focus on a setting where agents just remember their private experiences.

The comparison with BROWN *et al.* raises one further important issue. On the surface, one might think that our design cannot disentangle pure competition effects from the effects of

⁵ An earlier study on endogenous group formation in public good games is EHRHART and KESER (1999). There, however, the choice of partners is indirect since subjects choose locations (“islands”) where they want to play and all subjects choosing the same location form then a group.

⁶ There is also a literature on buyer-seller networks (theoretically developed by KRANTON and MINEHART (2001)). While this sounds as if it should be related to our study, in essence, it is not since there the network as such changes cost structures while in our study the underlying game itself is never affected by the choice of partners.

⁷ A similar study for search goods can be found in KIRCHSTEIGER, NIEDERLE, and POTTERS (2005).

long-term relations. However, this is not true. The crucial point is that, both, from a game theoretic point of view as well as from a behavioral angle, an environment with full information and random matching is, for all practical purposes, identical to a treatment with repeated interaction (see BOHNET and HUCK (2004) for a comparison of very similar trust games under such treatments as well as HSU (2003) who shows how the equivalence depends on experience and the amount of information made available). Hence, any difference that we observe between full information with random matching and full information with endogenous matching can be attributed to the introduction of competition.

The theory on reputation building that goes back to SCHELLING (1960) and was formally developed in seminal contributions by KREPS and WILSON (1982) and KREPS, MILGROM, ROBERTS, and WILSON (1982) is by now well understood and it is obvious how the presence of feedback information enables trust as an equilibrium phenomenon in such models. There is far less theory about endogenous matching of agents to play games. Most of the network literature has focused on a reduced-form approach where the value of links between neighbors is immediately given through a value function that maps network configurations into payoffs. Exceptions include studies of network formation in coordination games (see, for example, GALEOTTI *et al.* 2006). More general classes of non-cooperative games are studied in a recent paper by JACKSON and WATTS (2005) who introduce the notion of “social games”. They use a non-cooperative approach to study the interaction once a match has been formed, but rely on cooperative methods to analyze the stability of matches.

While we do not attempt a detailed formal analysis of the game we implement in the experiment, some observations are straightforward. As mentioned above, the basic game is a binary 2-player trust game where the first mover (the buyer) can either buy or not. If he decides to buy, the second mover (the seller) chooses whether to deliver high or low quality. This game has a *unique* Nash equilibrium where the seller would deliver low quality and the buyer does not buy. This no-trade result remains robust if the game is repeated finitely many times, *without* the need to invoke subgame perfection. Indeed, any Nash equilibrium of the finitely repeated trust game entails no trade in all periods. This is due to the fact that there is no way for the buyer to reward the seller for high quality along the equilibrium path. The same holds true in our experimental version with random matching. In this setup there is one further type of outcome though that needs to be considered, the outcome where a seller has no buyer because at least one seller was matched with two buyers. We decided to make this the worst outcome for the seller. The economic justification for this is that sellers might carry a

number of products, the experience good that is at the core of our study, and some other goods of obvious quality that the consumer purchases on the way, regardless of whether he buys the experience good. Without competition, the possibility that sellers might face an empty store does not affect the equilibrium outcome. Again, there is no trade in all periods since buyers still cannot reward sellers for having chosen high quality in the past.

Things change dramatically when competition is introduced. Even without incomplete information and reputation building there are now subgame perfect equilibria with trade. The reason for this is that buyers can now reward sellers for high quality, simply by *choosing them again*. The (credible) promise to return to a seller can induce sufficiently high incentives for sellers to choose high quality.⁸ However, the no-trade equilibria do not disappear. They co-exist and are also subgame perfect.

In the experiment, of course, subjects are able to move to better outcomes and our analysis of patterns of individual play suggests that they do indeed understand the equilibrium mechanics of rewarding high-quality sellers with repeat custom.

2. Experimental Design

Section 2.1 introduces the binary-choice trust game used in all treatments of our experiment. Section 2.2 describes our experimental design and procedures.

2.1 The trust game

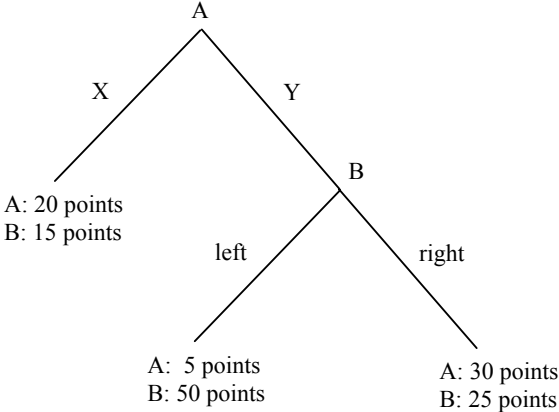
We consider a binary-choice trust game with two players as shown in Figure 1. The first mover (A, the buyer) chooses between “Y” (buy, trust) and “X” (don’t buy, don’t trust). After “Y”, i.e., if being trusted, the second mover (B, the seller) can choose between “right” (honor trust, choose high quality) and “left” (exploit trust, choose low quality).

The payoffs were chosen to reflect a conflict of interest: buyers prefer high quality over low quality ($30 > 5$) but it is cheaper for sellers to provide low quality ($50 > 25$, see Figure 1). While buyers prefer not to buy over getting low quality ($20 > 5$), sellers would rather provide high quality than not selling at all ($25 > 15$). Payoffs are intentionally asymmetric to create a

⁸ The total trade volume that can be sustained in a subgame perfect equilibrium depends, of course, on the exact specification of payoff parameters.

difficult moral hazard problem.⁹ Assuming that subjects maximize their own monetary income, the game has a unique Nash equilibrium in which the buyer does not trust and the seller does not honor trust if being trusted.

Figure 1: The trust game



Notice that, as long as the buyer trusts, the sum of payoffs is maximal. Only the distribution of payoffs depends on the sellers’ quality choice. Nevertheless, we shall in what follows refer to (Y, right), the outcome with successful trade, as “the efficient outcome” as both the buyer and the seller are better off with a high-quality trade than without trade.¹⁰

In an earlier study, BOHNET, HARMGART, HUCK, and TYRAN (2005) examined the same game, always with one-on-one random matching. They found that the payoffs are such that when sellers cannot be identified (since they do not have labels), i.e., in the absence of any information about the history, trust and honor rates are extremely low. In their “NoInfo” treatment subjects reach efficient trade in only 5% of all cases. Hence, the game serves as an extremely good benchmark to examine the effect of various institutional aspects designed to overcome the moral hazard problem. We replicate BOHNET *et al.*’s finding in our *control* treatment achieving an efficiency level of 8%.

⁹ With symmetric payoffs subjects find it much easier to achieve efficiency already in one-shot games (see, for example, BACHARACH, GUERRA, and ZIZZO (2004) or BOLTON, KATOK, and OCKENFELS (2004)).

¹⁰ This can be viewed as short for “the unique individually rational efficient outcome”, as individual rationality would be violated in (Y, left) where the buyer would clearly prefer to have played X instead.

2.2 Experimental Design and Procedures

In all treatment conditions, the following basic structure is implemented. The stage-game described in the previous section is repeated 30 times, and this is announced at the beginning of the experiment. Subjects are either a buyer or a seller for the entire experiment. A market consists of eight participants: 4 buyers and 4 sellers.

Treatments differ in two dimensions: the information structure and the matching procedure (see Table 1). Matching can be exogenous (with each buyer being randomly assigned to a seller in each round) or endogenous (with buyers choosing sellers in each round).¹¹ We refer to the latter treatments as the treatments “with competition” since sellers compete for buyers or market shares in this case. We also vary the amount of information buyers obtain about sellers’ past. In the absence of information about seller identity and, hence, their history, reputation formation is not possible and buyers cannot make a meaningful choice among buyers in this case (see treatment *control* in Table 1). With “private information” (*pi* in Table 1) all subjects (buyers and sellers) are only informed about their own past: in each period they can access a complete history of their *own* past interactions. With “full information” (*fi* in Table 1), all subjects have access to the history of *all* sellers in the market. It is important to notice that, in asymmetric trust games considered here, treatments with full information and random matching have been shown to generate outcomes that are virtually identical to treatments with repeated interaction (see HSU 2003 as well as BOHNET and HUCK 2004). Hence, the comparison between treatments *fi-nc* and *fi-c* simply measures the impact of competition.

In both treatments, feedback information about the past was provided by a visual interface using colored hash (#) signs to summarize outcomes in a “history window”.¹² This window shows four columns of hash (#) signs, each column representing one seller and each row representing one period.¹³ Initially, each column consists of thirty *white* hash signs. Then, period by period, hash signs change their color according to what happens in the game: a hash turns *grey* if a subject did not receive any information about a particular seller in the previous

¹¹ Buyers always had to choose a seller even if they did not want to buy at all. This was just for convenient interface design (see below) and to keep expected payoffs in case of no trust constant across treatments. However, we shall interpret the choice of any particular seller, combined with the decision not to buy, simply as the decision not to buy from any seller.

¹² See Appendix B for examples of screenshots.

¹³ In the private information treatments, sellers see only one column in the history window, the one representing their own history.

period, it turns *black* if a seller had a non-trusting buyer, it turns *red* if a seller had at least one buyer and chose low quality, and, finally, it turns *green* if a seller had at least one buyer and chose high quality.¹⁴ In case a transaction occurred in the treatments with full information, (red or green) hash signs were followed by a positive number indicating how many buyers had chosen this seller.

Table 1: Treatments

	<i>no competition</i>	<i>competition</i>
<i>full information</i>	<i>fi-nc</i>	<i>fi-c</i>
<i>private information</i>	<i>pi-nc</i>	<i>pi-c</i>
<i>no information</i>	<i>control</i>	

Note that with both matching procedures sellers can have between zero and four buyers in each period. But sellers can only choose one quality level, i.e., they cannot discriminate between different buyers to keep the experiment simple.¹⁵ A seller's total payoff in a particular period is the sum of payoffs that results from interacting with all his buyers. If a seller does not have a buyer in a period, his payoff in that period is zero. Buyers can have at most one interaction per period and earn the payoffs shown in Figure 1.

The sequence of actions is as follows. Depending on the matching procedure, buyers are either assigned to a seller or choose one at the beginning of each period. Then, buyers choose to buy or not. Sellers learn how many buyers have chosen them as well as how many decided to buy from them. Next, sellers have to decide about the quality of the good, provided they have at least one buyer. At the end of each period, buyers who decided to buy are informed about the choice of their seller and the resulting payoff, while sellers are simply reminded of the number of buyers who bought from them, their own choice, and the resulting payoff.

The experiment was conducted in the experimental laboratory at the University of Erfurt. All sessions were computerized and the experimental software was developed with FISCHBACHER's (1999) z-Tree. In total, 336 students of various fields participated in the experiment. Every subject only participated in one session. Subjects were recruited over the

¹⁴ Our suggestive color coding serves the purpose of making the complex history information easy to read. Of course, their might have been problems with color-blind subjects but fortunately there were none.

internet.¹⁶ Before the start of the experiment, subjects received written instructions.¹⁷ The language was kept neutral. Neither did we use terms like “trust” or “trustworthiness”, nor terms like “buyer” and “seller.” Instead, “A”-participants had to choose between “X” and “Y” and “B”-participants could decide between “left” and “right”, as labeled in the game tree shown in Figure 1. All subjects correctly answered a series of control questions before we started the experiment.

We ran three (two) sessions per main treatment (in the control treatment), each with 24 subjects, in three markets of 8 subjects. Thus, we have 9 independent markets in the main treatments, and 6 markets in the control treatment. Subjects earned 14.08 Euros on average within 60 to 90 minutes, including a lump sum depending on their role in the experiment.¹⁸

3. Results

We begin our discussion of results by examining some descriptive statistics of aggregate play. We then proceed by describing some aspects of the dynamics and close with a brief analysis of how individual reputations affect behavior.

Table 2 shows in the upper part the average trust rate (i.e. the percentage of cases in which buyers chose to buy), the average honor rate (i.e. the percentage of cases in which sellers provided high quality) and the average efficiency rate (i.e. the percentage of cases in which buyers trust sellers *and* sellers deliver high quality) for each treatment aggregated over all 30 rounds. The two rightmost columns show average earnings per period for buyers and sellers, respectively.

In the treatment *control*, trust and honor rates are low (28%), and, as a consequence, the efficiency rate is at a meager 8%.¹⁹ Per-period earnings are low for both buyers and sellers. Average buyer earnings are 17.7, which is even below the outside option since the buyers

¹⁵ Also, since sellers cannot identify buyers, discrimination would not make much sense.

¹⁶ We used GREINER's (2004) *ORSEE*.

¹⁷ A translation of the instructions can be found in Appendix A.

¹⁸ Buyers received fewer points than sellers to guarantee that sellers who did not very often interact with buyers got a reasonable payoff.

¹⁹ Note that the trust rate is the percentage of times “Y” was chosen ($= \#Y / (\#X + \#Y)$), the honor rate is percentage of “right” given that a buyer chose “Y” ($= \#“right” / \#“Y”$). The efficiency rate is the percentage of all choices that end in the lower right node of Figure 1, and is therefore the product of the two other rates ($= \#“right” / (\#X + \#Y)$).

who trust get low quality most of the time. This treatment generates an outcome that is quite close to the game theoretical prediction of no trade and is in line with earlier findings.²⁰

Table 2: Overview of aggregated results

	trust rate	honor rate ²¹	efficiency rate	buyer earnings	seller earnings
<i>control</i>	0.28 (0.09)	0.28 (0.15)	0.08 (0.04)	17.7 (1.23)	23.0 (2.79)
<i>pi-nc</i>	0.50 (0.12)	0.72 (0.08)	0.36 (0.11)	21.5 (1.20)	23.5 (1.86)
<i>fi-nc</i>	0.51 (0.16)	0.73 (0.16)	0.37 (0.19)	21.6 (2.54)	23.6 (1.48)
<i>pi-c</i>	0.86 (0.08)	0.92 (0.05)	0.79 (0.07)	26.8 (1.03)	25.4 (1.65)
<i>fi-c</i>	0.90 (0.07)	0.94 (0.03)	0.85 (0.07)	27.7 (0.90)	25.3 (0.95)
Effect of competition					
<i>fi-c</i> vs <i>fi-nc</i>	$p = 0.000$	$p = 0.000$	$p = 0.000$	$p = 0.000$	$p = 0.012$
<i>pi-c</i> vs <i>pi-nc</i>	$p = 0.000$	$p = 0.000$	$p = 0.000$	$p = 0.000$	$p = 0.031$
Effect of information					
<i>control</i> vs <i>pi-nc</i>	$p = 0.000$	$p = 0.000$	$p = 0.000$	$p = 0.000$	$p = 0.012$
<i>fi-c</i> vs <i>pi-c</i>	$p = 0.115$	$p = 0.149$	$p = 0.032$	$p = 0.033$	$p = 0.281$
<i>fi-nc</i> vs <i>pi-nc</i>	$p = 0.490$	$p = 0.441$	$p = 0.475$	$p = 0.475$	$p = 0.390$

Standard deviations are given in parentheses. The p -values refer to treatment differences according to Mann-Whitney U-tests (one-tailed).

With private information (*pi-nc*), market outcomes are much better than in *control*. In comparison to *control*, the trust rate almost doubles (50%), the honor rate more than doubles (72%), and the efficiency rate more than quadruples (36%). The possibility of reputation formation benefits both buyers and sellers. Average payoffs are significantly higher for both types according to Mann-Whitney U-tests. In fact, the treatment differences between *control* and *pi-nc* are highly significant for all measures (see the bottom part of Table 2).

What is the effect of having full information rather than private information in the absence of competition? Surprisingly, there is no effect. The aggregate results in treatment *fi-nc* are almost identical to *pi-nc* on all accounts.

²⁰ As mentioned earlier, Bohnet *et al.*'s (2005) "NoInfo" treatment is – except for the matching procedure – the same as *control* here. They find for UK-subjects a trust rate of 21%, a honor rate of 19%, and an efficiency rate of 5%.

²¹ The reported honor rate weighs according to the number of buyers a seller had in a given period. Thus, if a seller had $n \leq 4$ buyers he is considered as being trustworthy n -times. However, results do not change much without this weighting: *fi-c* 0.94, *fi-nc* 0.73, *pi-c* 0.89, *pi-nc* 0.73.

On the other hand, the effect of competition is very powerful. Given private information, competition boosts the trust rate from 50 to 86% and the honor rate from 72 to 92%. As a consequence, the efficiency rate more than doubles from 36 to 79% (compare *pi-nc* and *pi-c* in Table 2). Buyers benefit strongly from competition. Their flexible earnings (i.e. earnings in excess of the minimum of 5 per period) increase by 32% (from 16.5 to 21.8). But, amazingly, sellers also benefit from competition. Their flexible earnings²² increase by 47% (from 11.3 to 16.6). Sellers benefit from competition since it helps to overcome the moral hazard problem straining all agents.

The effects of competition are equally impressive given full information: The trust rate is at 90% in *fi-c* rather than 51% in *fi-nc*, and the honor rate is at 94% in *fi-c* rather than 73% in *fi-nc*. The efficiency rate is boosted from 37% in *fi-nc* to 85% in *fi-c*. Effects on earnings are also similar to the partial information case.

Providing full information in a competitive market has a weak effect at best. A comparison of the 4th and 5th line in Table 2 shows that the levels are overall similar. According to Mann-Whitney U-tests, the efficiency rate and buyer earnings are weakly significantly higher with full information, while the other rates and seller earnings are not.

Overall, competition and information together more than triple the trust and honor rates, and improve efficient high-quality trade by a factor of 10 compared to the control treatment. Approximately half of this overall effect can be attributed to providing market participants with some (i.e. private) information, allowing them to build and observe reputations, while the other half can be attributed to competition. However, providing more than private (i.e. full) information and, thus, increasing market transparency, yields weak improvements at best. Apparently, full information is not necessary for these markets to provide good results.

Market outcomes over time

Figure 2 shows that average trust rates start out at similar levels of around 50% in all treatments. Over the next 10 periods, however, remarkable differences between some treatments are observed. In the two treatments with competition (marked with triangles), trust rates increase to around 90%, while they remain at around 50% in the two treatments without

²² Sellers fixed expected earnings under random matching are computed by taking into consideration that each buyer is assigned with probability $\frac{1}{4}$ to each seller. A seller's minimal payoff from a match is 15. In expectation, the fixed income component under random matching is 10.25.

competition (marked with squares). In the control treatment (marked with circles and a dotted line), trust falls slightly, to approximately 40%, by period 10. Over the rest of the game, these trends become more pronounced. In treatments with competition, trust rates further increase to reach almost 100%, in the treatments without competition a weakly declining trend is apparent, and in *control* the trust rates fall further to around 10%. Endgame effects occur literally in the last period only, and are most pronounced in the treatments with competition. Figure 2 also illustrates the fact that there was essentially no difference between private and full information. In fact, the two lines with triangles and the two lines with squares exhibit almost identical dynamic properties, respectively.

Figure 2: Average trust rates over time

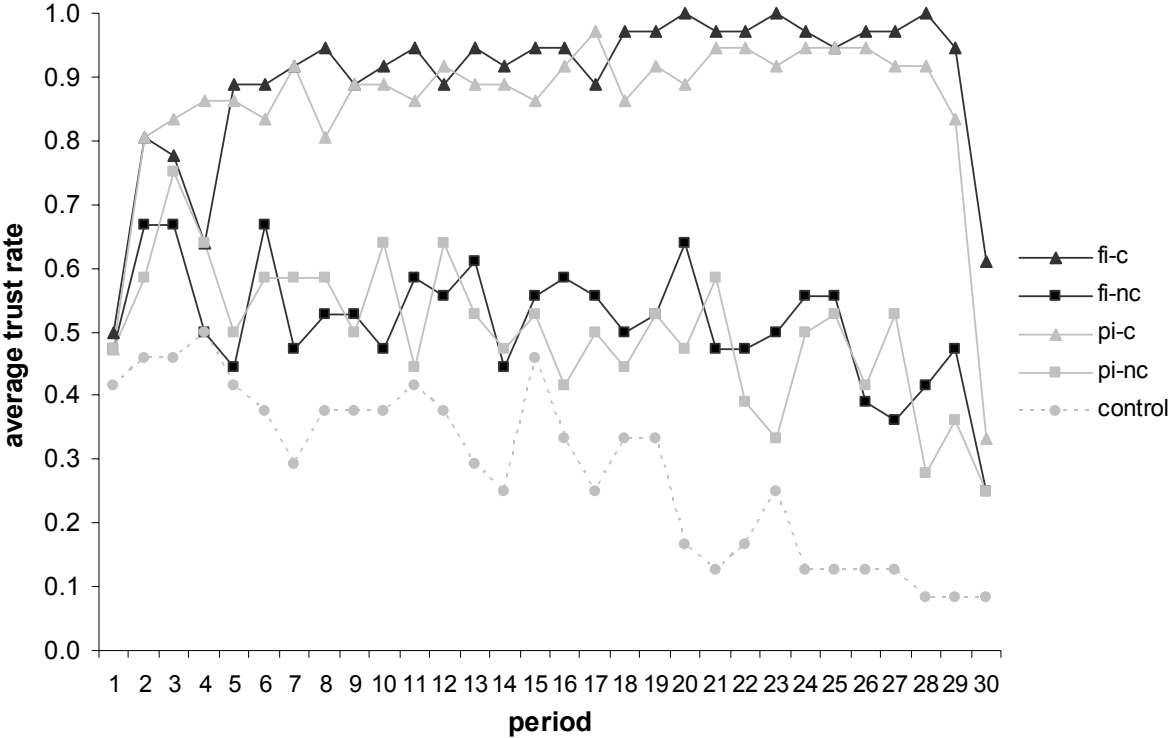
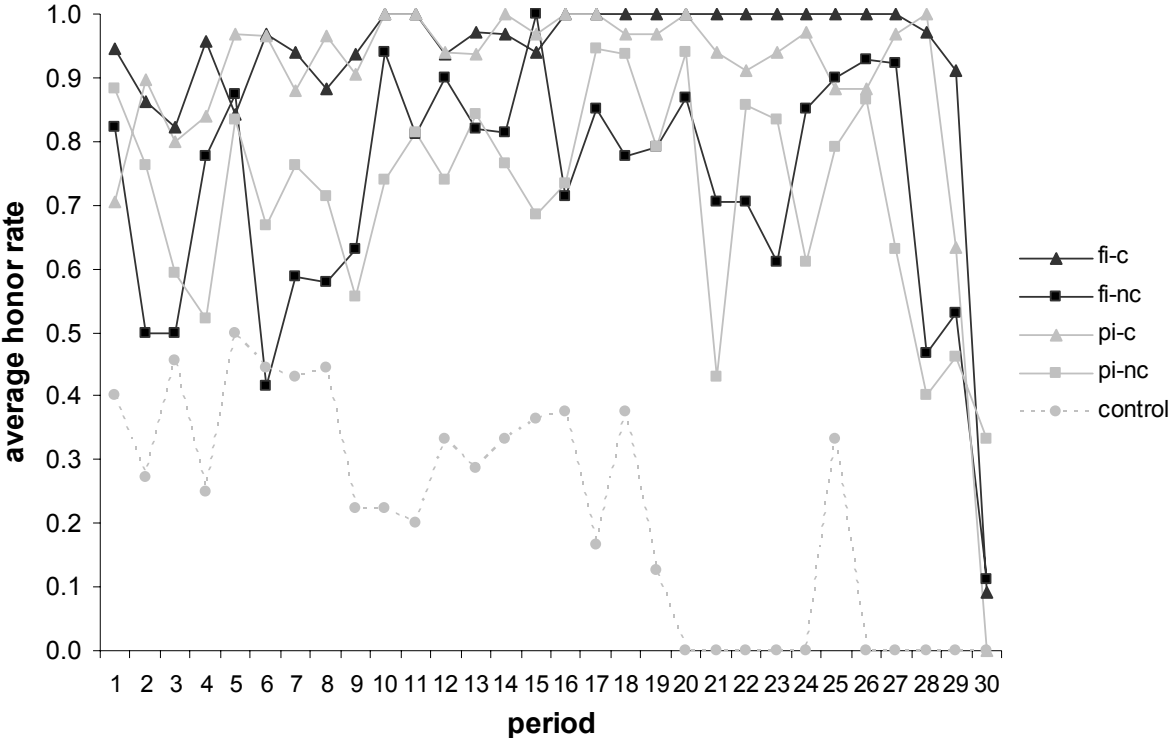


Figure 3 shows average honor rates over time and provides a somewhat noisier picture than Figure 2. The reason is that the denominator of this rate (i.e. number of times buyers choose to trust) is small in the treatments without competition. Taking this into consideration, we note the following: Honor rates start out high, at levels of 70 to 90% in all treatments except *control*. These high initial levels of trustworthiness are remarkable given the relatively low levels of trust of around 50% shown in Figure 2. Possible explanations for these comparatively low rates of trust are that buyers underestimate the sellers’ trustworthiness or,

in the treatments with full information, want to free-ride on other buyers' experience in the first few periods of the game. In treatments with competition (see lines with triangles in Figure 3), the honor rates increase even further to reach values of close to or at 100% in the second half of the experiment. Notice, however, that some of this remarkable increase in trustworthiness is simply a consequence of selection as, with identifiability of sellers, less trustworthy sellers are trusted less often.

Figure 3: Average honor rates over time



It is only in the last (and next to last) period, where we observe pronounced endgame effects. In fact, the honor rate in the very last period is 5% in *fi-c* and exactly zero in *pi-c*. In the treatments without competition, honor rates are more fickle throughout but essentially remain at the initial level until close to end of the game. In control, however, honor rates decline after period 8 and are essentially zero over the last 10 periods.

Overall, the possibility of reputation building (i.e. the provision of private information) has a more pronounced effect on the average honor rate than on the trust rate, while competition seems to have a stronger effect on the trust rate than on the honor rate (compare Figures 2 and 3).

Figure 4 shows efficiency rates and provides a very clear and easy to interpret picture. Efficiency rates start out at similar levels of around 30 to 50% in the four main treatments. But a separation occurs quickly. In treatments with competition, efficiency rates increase steadily to reach levels of around 90 to 100%. In the last third of the experiment, the treatment with full information and competition seems to outperform the one with private information, as there is a consistent gap of about 10 percentage points (Mann-Whitney U-test: *fi-c* vs. *pi-c*: 0.000, one-tailed). No such difference due to the level of information is apparent in treatments without competition. The two lines marked with squares in Figure 4 are essentially identical, and reveal no increasing trend over time. In the control treatment, efficiency is low, at around 10 to 20% in the first half, and falls to essentially zero in the last third of the game.

Figure 4: Average efficiency rates over time

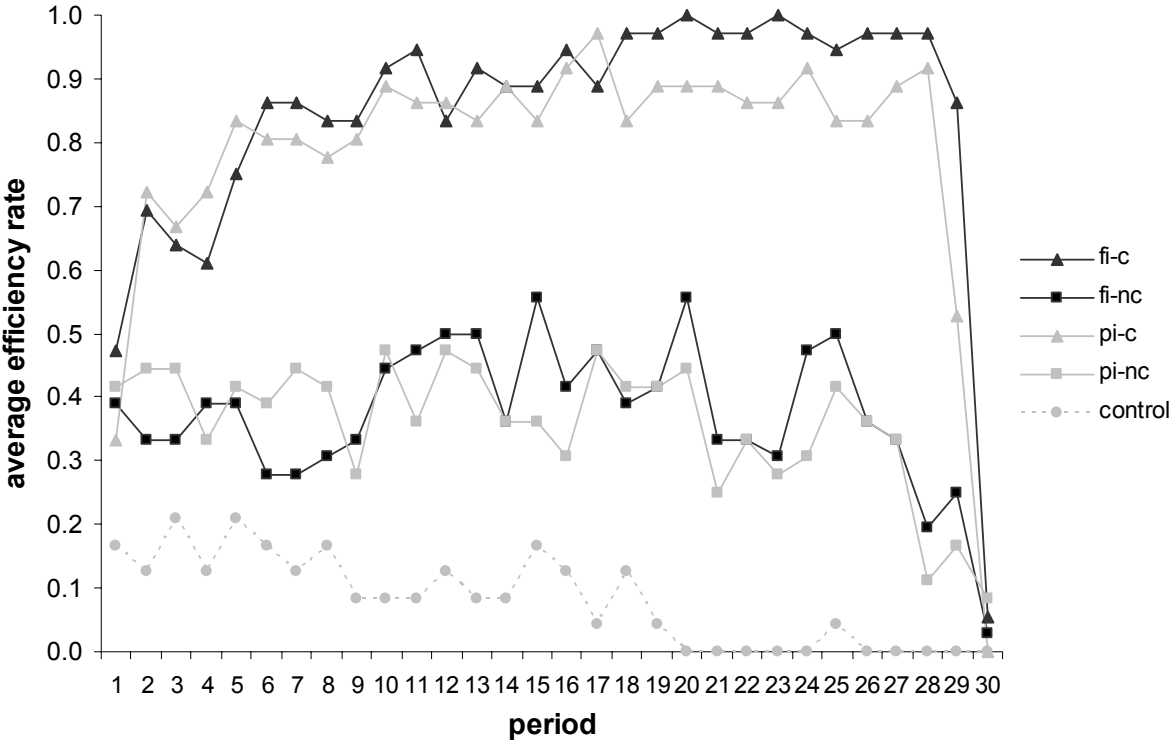


Figure 4 very clearly demonstrates that the combined effect of reputation and competition is dramatic. The combined effect literally makes the difference between a rate of zero and 100% efficiency in the last third of the game (compare *control* and *fi-c* in Figure 4). The figure also shows that a bit less than half of this impressive combined effect is due to the availability of private information (vs. *control*), and a bit more than half of the combined effect is due to competition.

The role of reputations

From what has been said above, it is evident that providing minimal (private) information to build some reputation considerably improves market performance compared to anonymous markets, and that competition for buyers massively improves market performance, both, with private and full information. The reason why we expect reputations to be effective in raising market efficiency and expect competition to multiply these effects is that buyers follow specific behavioral rules. If buyers systematically shop from sellers with a clean past and shun sellers with a history involving low quality, sellers have an incentive to invest in and maintain a good reputation, at least in early periods of the game. The existence of such a behavioral pattern necessitates that buyers extrapolate sellers' behavior into the future. That is, they take a seller's history as an indicator of his future trustworthiness.

In order to investigate both these issues we run two simple regressions. The first simply investigates whether good reputations do indeed drive consumers' choices to buy. Focusing on the full information treatments we regress the buyer's choice whether to trust or not on the average honor rate of the seller so far (*greenshare*), a competition dummy (equal to one in *fi-c*) and an interaction term of these two variables. The second regression investigates whether high quality choices in the past are a reliable indicator for high quality in the future, simply by regressing the seller's quality choice on the same set of variables. Estimation results for linear probability models (with clustering on the matching group level) are shown in Table 3. Probits give identical results.

The estimation results confirm our intuition. A seller's average quality in the past is a good predictor for both, a buyer's trust *and* the seller's next quality choice. In *fi-c* buyers react a little less sensitively to variations in past quality but trust generally more. However, the second regression shows that past quality predicts future quality equally well in both full information treatments. (Similar results are obtained for treatments with partial information.)

Table 3: Regression results in treatments with full information

<i>variable</i>	<i>dependent variable:</i> <i>buyer's trust</i>		<i>dependent variable:</i> <i>seller's high quality choice</i>	
	<i>coefficient</i>	<i>std. error</i>	<i>coefficient</i>	<i>std. error</i>
<i>greenshare</i>	0.669***	0.081	0.689***	0.057
<i>competition</i>	0.573**	0.131	0.314	0.117
<i>greenshare*competition</i>	-0.447*	0.143	-0.096	0.143
<i>constant</i>	0.130*	0.049	-0.014	0.023
	matching groups: 9		matching groups: 9	
	observations: 2,015		observations: 2,015	
	R ² 0.383		R ² 0.439	

*** $p \leq 0.001$, ** $p \leq 0.01$, * $p \leq 0.05$

The regressions nicely illustrate the role of reputations. Sellers have a strong incentive to invest in good reputations as this generates repeat custom. But once sellers have a good reputation they also face the temptation of exploiting it at some stage. *Ceteris paribus*, this temptation is much stronger when there are more customers in a seller's store and, in the absence of competition, we do find that average quality dramatically falls in crowded stores. In the presence of competition, however, we find that this reversed. If buyers can choose from whom they want to buy in the future, the costs of alienating many buyers overcompensate for the increased temptation. Table 4 shows average rip-off rates (i.e. the percentage of cases in which sellers provided low quality) for sellers with up to or more than two sellers. The effects of competition are dramatic. With competition, buying in crowded stores is tremendously safe. In the absence of competition, on the other hand, buying in crowded stores is a terrible idea. Our guess is that this pattern could also be nicely tested in the field. Crowded stores at market places mainly frequented by locals (who actually *chose* to buy in these stores) may function quite differently from crowded stores at market places frequented by tourists (who more or less randomly picked a seller).

Table 4: Average rip-off rates dependent on the number of customers

	<i>number of customers</i>	
	<i>1 or 2</i>	<i>3 or 4</i>
<i>fi-c</i>	0.06	0.05
<i>pi-c</i>	0.11	0.01
<i>fi-nc</i>	0.26	0.35
<i>pi-nc</i>	0.26	0.51

Customer loyalty and competition

When buyers can choose from whom to buy customer loyalty becomes an issue. In general, we find that buyers tended to be very loyal. To provide a measure of customer loyalty, we calculate for each buyer the longest uninterrupted sequence of interaction with the same seller. On average, these sequences extend through more than half of the game in both information conditions (*fi-c* 16.47, and *pi-c* 15.89 periods), much longer than the corresponding value resulting from random matching (*fi-nc* 2.36 and *pi-nc* 2.11).

We find that competition for buyers is a discipline device because buyers systematically reward and sanction sellers for their past behavior in their *own* interaction. For example, the average buyer remains loyal to a seller if he sold him high quality in the previous period in 72% of the cases in both *fi-c* and *pi-c*. In contrast, after having received low quality, buyers *always* switch to another seller in *fi-c* and do so 79% of the time in *pi-c*.

While buyers shun sellers who have disappointed them, we next ask whether buyers are tempted to switch to sellers with particularly good reputations. It turns out that they do, but that reputation earned in one's own transactions is given more weight than reputation a seller has earned in transactions with other buyers. To illustrate, we ask how often buyers buy from the seller with the relatively best reputation, i.e., the seller with the highest share of green hashes among all sellers. While in *pi-c* buyers choose the seller with the best history in 78% of all cases, buyers do so only 59% of the time in *fi-c*. This highly significant difference (Mann-Whitney U-test: *fi-c* vs. *pi-c*: 0.016, one-tailed) is startling. Why do buyers have less trust in the available information in treatment *fi-c*? The answer can be found by computing how often buyers in *fi-c* bought from the seller with whom they *themselves* made the best

experience. Specifically, we compute how often a buyer in *fi-c* bought from the seller who has the relatively best reputation vis-à-vis *this* buyer. Remarkably, this number is almost identical to the one we found in *pi-c*. In 76% of all decisions buyers in *fi-c* buy from the seller with whom they had the best experience in the past (*pi-c*: 78%). In other words, it seems that buyers largely ignored the additional information they got in *fi-c* and simply concentrate on their own experience.²³

Finally, we can ask how forgiving buyers were. That is, do they return to a seller who provided them with low quality? It turns out that buyers are more forgiving in *fi-c* than in *pi-c*. In *fi-c*, 50.0% (17/34) of the buyers “forgive” a seller for having provided them with low quality at some point (if they return, this happens after 9.9 periods on average). In *pi-c*, buyers return in only 25.0% (16/64) of the cases (after 5.4 periods on average). This finding indicates that buyers made some use of other buyers’ experience. A seller is more likely to “get a second chance” after providing a buyer with low quality if this buyer can observe that the seller provided high quality to other buyers. This, of course, is only possible in treatments with full information.

As a consequence of the regularities in buyers’ choices discussed above, market shares are quite unequal in the presence of competition. Table 5 shows average market shares of more and less successful firms for all four main treatments.

Table 5: Average market shares

	<i>highest</i>	<i>seller with the</i>		<i>lowest</i>
		<i>2nd highest</i>	<i>3rd highest</i>	
		<i>market share</i>		
<i>fi-c</i>	0.51	0.23	0.15	0.11
<i>pi-c</i>	0.57	0.24	0.11	0.08
<i>fi-nc</i>	0.32	0.26	0.23	0.19
<i>pi-nc</i>	0.28	0.27	0.24	0.21

²³ This overweighting of own information and/or experiences is also found in different contexts. For example, subjects in information cascade experiments tend to overweigh their own private signal (see, for example, GOEREE, PALFREY, ROGERS, and MCKELVEY (2006)). A finding that resembles our results even more is reported in HUCK and JEHIEL (2004) who show that subjects tend to predict other’s behavior largely based on own experience even when more accurate public statistics are available.

In the future we also plan to study experiments where buyers can only observe market shares of different sellers to examine whether success can signal quality.

4. Conclusion

This study suggests that reputation mechanisms can greatly improve the functioning of markets for experience goods, and this finding is in line with the intuitions of many economists and with earlier experimental results. However, the effects of reputation on its own are limited. Only if some minimal feedback information – enabling the mechanics of reputation building – is coupled with competition we observe that the market-crippling moral hazard problem is fully overcome.

There are a couple of surprising aspects of our findings. First, the sheer size of the competition effect is stunning. While, given the size of our data set, this is clearly not a fluke one wonders about the robustness of these results with respect to the implicit market price²⁴ and the precision of the feedback information. Both these questions provide obvious and interesting avenues for further research. Second, we find that full public information is no better than private experience only. While this is not very surprising in the treatments with competition where private information is already sufficient to reach an effective ceiling, this is puzzling in the random matching treatments. Clearly, delivering low quality should be more costly if this is observed by the entire population than just by a single customer. This is why we expected to observe higher efficiency levels with full information. One reason for why this does not happen is that buyers apparently attach a much higher weight to their own experience than to the experience of others, a phenomenon not limited to our study.²⁵

The main contribution of this study is clearly to demonstrate for the first time a strong link between trust and competition. As discussed in the introduction, trust has been shown to be an important element of economic growth and lack of trust (or more generally low levels of social capital) can impede economic development. What we show here is that *competition fosters trust*. In our environment the efficiency gains from competition stem from this link.

²⁴ In this study we have focussed on one aspect of competition, free choice of trading partners. In most markets this element is complemented by some form of price competition. In a follow-up study we plan to investigate the consequence of an additional first stage where sellers can post a price before buyers choose a seller.

²⁵ Another possible reason is the rather small size of our market. In a much larger market (the market for restaurants in Los Angeles) Jin and Leslie (2003) do find positive effects of public feedback about quality.

With competition, there are stronger incentives for resisting the short-term temptation of exploiting trading partners. And these incentives are increasing in past success. The more trading partners you have, the more costly a deviation. This is why in our experiments with competition buyers are extremely safe in crowded stores. In that sense we observe how success breeds success — trust being the key catalyst in this process.

References

- ANDERHUB, Vital; ENGELMANN, Dirk and GÜTH, Werner (2002): An Experimental Study of the Repeated Trust Game with Incomplete Information, *Journal of Economic Behavior and Organization* 48, 19-216.
- BACHARACH, Michael; GUERRA, Gerardo and ZIZZO, Daniel (2004): The Self-Fulfilling Property of Trust: An Experimental Study, *mimeo*, University of Oxford.
- BERG, Joyce; DICKHAUT, John and MCCABE, Kevin (1995): Trust, Reciprocity, and Social History, *Games and Economic Behavior* 10, 122-142.
- BOHNET, Iris; HARMGART, Heike; HUCK, Steffen and TYRAN, Jean-Robert (2005): Learning Trust, *Journal of the European Economic Association* 3(3), 322-329.
- BOHNET, Iris; FREY, Bruno S. and HUCK, Steffen (2001): More order with less law: On contract enforcement, trust, and crowding, *American Political Science Review* 95, 131-144.
- BOHNET, Iris and HUCK, Steffen (2004): Repetition and Reputation: Implications for Trust and Trustworthiness when Institutions Change, *American Economic Review* 94(2), 362-366.
- BOHNET, Iris and KÜBLER, Dorothea (2005): Compensating the Cooperators: Is Sorting in the Prisoner's Dilemma Possible?, *Journal of Economic Behavior and Organization* 56(1), 61-76.
- BOLTON, Gary; KATOK, Elena and OCKENFELS, Axel (2004): How Effective are Electronic Reputation Mechanisms? An Experimental Investigation, *Management Science* 50(11), 1587-1602.
- BROWN, Martin; FALK, Armin and FEHR, Ernst (2004): Relational Contracts and the Nature of Market Interactions, *Econometrica* 72, 747-780.
- CABRALES, Antonio; CHARNNESS, Gary and VILLEVAL, Marie-Claire (2006): Competition, Hidden Information and Efficiency: An Experiment, *IZA Discussion Paper* No 2296.
- CORICELLI, Giorgio; FEHR, Dietmar and FELLNER, Gerlinde (2004): Partner Selection in Public Goods Experiments, *Journal of Conflict Resolution* 48(3), 356-378.
- DUFFY, John and FELTOVICH, Nick (2002): Do Actions Speak Louder than Words? Observations vs. Cheap Talk as Coordination Devices, *Games and Economic Behavior* 39, 1-27.
- DUFFY, John and FELTOVICH, Nick (2006): Words, Deeds and Lies: Strategic Behavior in Games with Multiple Signals, *Review of Economic Studies* 7, 669-688.

- DUFFY, John and HOPKINS, Ed (2005): Learning, Information and Sorting in Market Entry Games: Theory and Evidence, *Games and Economic Behavior* 51, 31-62.
- EHRHART, Karl-Martin and KESER, Claudia (1999): Mobility and Cooperation: On the Run, *CIRANO Working Paper 1999s-24*.
- FISCHBACHER, Urs (1999): z-Tree: Zurich Toolbox for Readymade Economic Experiments, *Working Paper No. 21*, Institute for Empirical Research in Economics, University of Zurich.
- GALEOTTI, Andrea; GOYAL, Sanjeev; JACKSON, Matthew O.; VEGA-REDONDO, Fernando and YARIV, Leeat (2006): Network Games, *mimeo*.
- GLAESER, Edward L.; LAIBSON, David I.; SCHEINKMAN, Jose A. and SOUTTER, Christine L. (2000): Measuring Trust, *Quarterly Journal of Economics* 115, 811-846.
- GOEREE, Jacob K.; PALFREY, Thomas R.; ROGERS, Brian W. and MCKELVEY, Richard D. (2006): Self-correcting Information Cascades, *Working paper*, California Institute of Technology.
- GREINER, Ben (2004): The Online Recruitment System ORSEE 2.0 – A Guide for the Organization of Experiments in Economics, *Working Paper Series in Economics 10*, University of Cologne.
- GÜTH, Werner; OCKENFELS, Peter and WENDEL, Markus (1997): Cooperation Based on Trust: An Experimental Investigation, *Journal of Economic Psychology* 18, 15-43.
- HSU, Li-Chen (2003): Random Matching with Observable Past Actions: Experimental Tests of the Reputation Effect, *mimeo*.
- HUCK, Steffen and JEHIEL, Philippe (2004): Public Statistics and Private Experience: Varying Feedback Information in a Take-or-Pass Game, *Working paper*, University College London.
- HUCK, Steffen; NORMANN, Hans-Theo and OECHSSLER, Jörg (1999): Learning in Cournot Oligopoly: An Experiment, *Economic Journal* 109, 80-95.
- HUCK, Steffen; NORMANN, Hans-Theo and OECHSSLER, Jörg (2000): Does Information about Competitors' Actions Increase or Decrease Competition in Experimental Oligopoly Markets?, *International Journal of Industrial Organization* 18, 39-57.
- JACKSON, Matthew O. and WATTS, Alison (2005): Social Games: Matching and the Play of Finitely Repeated Games, *mimeo*, Caltech and Southern Illinois University.
- JIN, Ginger and LESLIE, Phillip (2003): The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards, *Quarterly Journal of Economics* 118, 409-451.
- KESER, Claudia (2002): Trust and Reputation Building in E-Commerce, *CIRANO Working Paper 2002s-75*.
- KIRCHSTEIGER, Georg; NIEDERLE, Muriel and POTTERS, Jan (2005): Endogenizing Market Institutions: An Experimental Approach, *European Economic Review* 49, 1827-1853.
- KOSFELD, Michael (2004): Economic Networks in the Laboratory: A Survey, *Review of Network Economics* 3, 20-41.
- KRANTON, Rachel and MINEHART, Deborah (2001): A Theory of Buyer-Seller Networks, *American Economic Review* 91(3), 485-508.

- KREPS, David M. and WILSON, Robert B. (1982): Reputation and Imperfect Information, *Journal of Economic Theory* 27(2), 253-279.
- KREPS, David M.; MILGROM, Paul R.; ROBERTS, John and WILSON, Robert B. (1982): Rational Cooperation in the Finitely Repeated Prisoners' Dilemma, *Journal of Economic Theory* 27(2), 245-252.
- OFFERMAN, Theo; POTTERS, Jan and SONNEMANS, Joep (2002): Imitation and Belief Learning in an Oligopoly Experiment, *Review of Economic Studies* 69, 973-977.
- PAGE, Talbot; PUTTERMAN, Louis and UNEL, Bulent (2005): Voluntary Association in Public Goods Experiments: Reciprocity, Mimicry and Efficiency, *Economic Journal* 115, 1032-1053.
- RIEDL, Arno and ULE, Aljaž (2002): Exclusion and Cooperation in Social Network Experiments, *mimeo*, University of Amsterdam.
- SCHELLING, Thomas C. (1960): *The Strategy of Conflict*, Harvard University Press, Cambridge.

Appendix A: Instructions (Treatment fi-c)

(Original instructions were in German. They are available from the authors upon request. In treatments with no competition A-participants were randomly assigned to one of the four B-participants. In treatments where only partial information was provided A-participants could only keep track of the history of the B-participant they interacted with in a particular period while B-participants only knew their own history. Here, A-participants were not provided with the information on the number of A-participants interacting with the same B-participant they were who chose Y.)

Welcome to the experiment!

Please read these instructions carefully! Do not speak to your neighbors and keep quiet during the entire experiment! In case you have a question raise your hand! We will then come to you.

At the beginning of the experiment you are randomly separated into **subpopulations of 8 participants**. During the experiment you solely interact with the participants of your subpopulation.

In this experiment you will repeatedly make decisions. Doing this you can earn points. Your total sum of points plus a show-up fee, which is dependent on your role in the experiment, will be converted into Euros at the end of the experiment and paid to you in cash. The show-up fee amounts to 150 points for A-participants and 330 points for B-participants. Following rule applies to the conversion of points into Euros:

$$\mathbf{1 \text{ point} = 0.015 \text{ €}}$$

How much you earn depends on your decisions and on the decisions of other participants in your subpopulation. All participants receive the same instructions. All decisions are made anonymously. No other participant will get to know your name and your payoff.

Altogether there are (in your subpopulation) eight participants. At the beginning of the experiment all participants are randomly assigned one of two roles (**A** or **B**) which is displayed on the computer screen. There are four A-participants (*A1, A2, A3 and A4*) and four B-participants (*B1, B2, B3 and B4*). All participants keep their role and the number assigned to them throughout the experiment.

At the beginning of each round each A-participant **chooses** one of the four B-participants with whom he interacts in this period, this means that A-participants decide whether they want to interact with B1, B2, B3 or B4. Therefore, a B-participant can interact with zero to four A-participants in one round. This process is repeated in the next round.

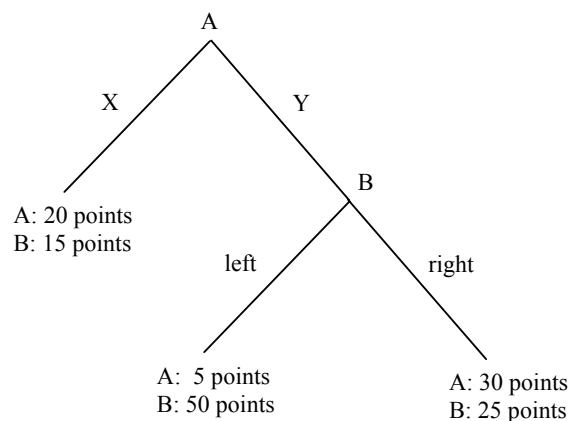
B-participants only learn the number of A-participants who chose them, but not who has chosen them.

After the A-participants have chosen the B-participants, it is each A-participants turn to make a decision. More specifically, each A-participant has to choose between option **X** and **Y** (see Figure). If he picks option X, the A-participant will earn 20 points and the B-participant chosen by him will earn 15 points from this interaction. If he picks option Y, the payoffs depend on the choice of the B-participant chosen by him who has to decide whether he wants to go “**left**” or “**right**”. If he decides to pick “left”, the A-participant will earn 5 points and the

B-participant will earn 50 points from this interaction. If he decides to pick “right”, the A-participant will earn 30 points and the B-participant will earn 25 points.

If a B-participant is not chosen by at least one A-participant or if all A-participants who have chosen him picked X, he does not have to make a decision. Otherwise the B-participant learns the number of A-participants who have chosen him and picked Y and has to make the **same** decision for all these A-participants, that is whether he wants to go “left” or “right” for all A-participants. A B-participant who has not been chosen by at least one A-participant earns 0 points.

Payoffs for an A-participant and a B-participant from one interaction.



(B-participants can receive payoffs from several interactions in one round.)

The experiment consists of **30 rounds**. After each round you will be informed about what has happened and you will be reminded of your payoff and your total sum of points so far.

Moreover, **all** (A- and B-) **participants can keep track of the entire history of B-participants**. For all participants there will be a screen depicting the history of all B-participants. For each round and each B-participant there will be a colored little # (hash) with a number behind.

- A **black** # indicates that the B-participant had nothing to decide because either no A-participant chose him or *all* A-participants who chose him picked X.
- A **red** # indicates that the B-participant picked “left”.
- A **green** # indicates that the B-participant picked “right”.
- The integer behind the # shows the number of A-participants who picked Y after choosing this B-participant.

These are the rules. You can trust us that everything will happen exactly according to these rules. Take your time going over these instructions again and feel free to ask questions. But don’t shout! Simply raise your hand.

Appendix B: Screenshot (Example: fi-c)

A-participant's screen

Period 17 of 30 remaining time [sec]: 11

Round	B1	B2	B3	B4
30	#0	#0	#0	#0
29	#0	#0	#0	#0
28	#0	#0	#0	#0
27	#0	#0	#0	#0
26	#0	#0	#0	#0
25	#0	#0	#0	#0
24	#0	#0	#0	#0
23	#0	#0	#0	#0
22	#0	#0	#0	#0
21	#0	#0	#0	#0
20	#0	#0	#0	#0
19	#0	#0	#0	#0
18	#0	#0	#0	#0
17	#0	#0	#0	#0
16	#3	#1	#0	#0
15	#2	#1	#1	#0
14	#3	#1	#0	#0
13	#3	#0	#0	#0
12	#1	#2	#0	#1
11	#3	#1	#0	#0
10	#2	#2	#0	#0
9	#2	#2	#0	#0
8	#3	#1	#0	#0
7	#3	#1	#0	#0
6	#2	#0	#1	#0
5	#0	#2	#0	#1
4	#1	#1	#0	#1
3	#1	#1	#2	#0
2	#2	#1	#0	#0
1	#2	#0	#0	#0

Choose a B-participant: B1
 B2
 B3
 B4

Choose which way to go: X
 Y

OK

Key:
Not done yet
Didn't participate
Went left
Went right

The number next to the hash shows the number of A-participants who chose Y.

B-participant's screen

Period 17 of 30 remaining time [sec]: 19

Round	B1	B2	B3	B4
30	#0	#0	#0	#0
29	#0	#0	#0	#0
28	#0	#0	#0	#0
27	#0	#0	#0	#0
26	#0	#0	#0	#0
25	#0	#0	#0	#0
24	#0	#0	#0	#0
23	#0	#0	#0	#0
22	#0	#0	#0	#0
21	#0	#0	#0	#0
20	#0	#0	#0	#0
19	#0	#0	#0	#0
18	#0	#0	#0	#0
17	#0	#0	#0	#0
16	#3	#1	#0	#0
15	#2	#1	#1	#0
14	#3	#1	#0	#0
13	#3	#0	#0	#0
12	#1	#2	#0	#1
11	#3	#1	#0	#0
10	#2	#2	#0	#0
9	#2	#2	#0	#0
8	#3	#1	#0	#0
7	#3	#1	#0	#0
6	#2	#0	#1	#0
5	#0	#2	#0	#1
4	#1	#1	#0	#1
3	#1	#1	#2	#0
2	#2	#1	#0	#0
1	#2	#0	#0	#0

2 A-participants chose you.
Thereof 2 A-participants chose Y.

Choose which way to go: left
 right

OK

Key:
Not done yet
Didn't participate
Went left
Went right

The number next to the hash shows the number of A-participants who chose Y.

Appendix B: Screenshot (Example: pi-c)

A-participant's screen

Period 18 of 30 remaining time [sec]: 19

Round	B1	B2	B3	B4
30	#	#	#	#
29	#	#	#	#
28	#	#	#	#
27	#	#	#	#
26	#	#	#	#
25	#	#	#	#
24	#	#	#	#
23	#	#	#	#
22	#	#	#	#
21	#	#	#	#
20	#	#	#	#
19	#	#	#	#
18	#	#	#	#
17	#	#	#	#
16	#	#	#	#
15	#	#	#	#
14	#	#	#	#
13	#	#	#	#
12	#	#	#	#
11	#	#	#	#
10	#	#	#	#
9	#	#	#	#
8	#	#	#	#
7	#	#	#	#
6	#	#	#	#
5	#	#	#	#
4	#	#	#	#
3	#	#	#	#
2	#	#	#	#
1	#	#	#	#

Choose a B-participant: B1
 B2
 B3
 B4

Choose which way to go: X
 Y

OK

Key:

- # Not done yet
- # Didn't participate
- # Went left
- # Went right
- # No information

B-participant's screen

Period 18 of 30 remaining time [sec]: 26

Round	B2
30	#0
29	#0
28	#0
27	#0
26	#0
25	#0
24	#0
23	#0
22	#0
21	#0
20	#0
19	#0
18	#0
17	#0
16	#0
15	#0
14	#0
13	#1
12	#0
11	#0
10	#0
9	#0
8	#1
7	#0
6	#0
5	#0
4	#1
3	#1
2	#1
1	#0

1 A-participant chose you.
 Thereof 1 A-participant chose Y.

Choose which way to go: left
 right

OK

Key:

- # Not done yet
- # Didn't participate
- # Went left
- # Went right

The number next to the hash shows the number of A-participants who chose Y.