How Effective are Electronic Reputation Mechanisms?

An Experimental Investigation

Gary E Bolton*

Smeal College of Business, Penn State University and Harvard Business School, USA

Elena Katok

Smeal College of Business, Penn State University and Harvard Business School, USA

Axel Ockenfels

Max Planck Institute for Research into Economic Systems, Jena, Germany

March 2003

Electronic reputation or "feedback" mechanisms aim to mitigate the moral hazard problems associated with exchange among strangers by providing the type of information available in more traditional close-knit groups, where members are frequently involved in one another's dealings. In this paper, we compare trading in a market with electronic feedback (as implemented by many Internet markets) to a market without, as well as to a market in which the same people interact with one another repeatedly (partners market). We find that, while the feedback mechanism induces quite a substantial improvement in transaction efficiency, it also exhibits a kind of public goods problem in that, unlike the partners market, the benefits of trust and trustworthy behavior go to the whole community and are not completely internalized. We discuss the implications of this perspective for improving these systems.

---

* Correspondence to Gary Bolton, 310 Beam, Smeal College of Business, Penn State University, State College, PA 16870; gbolton@psu.edu. Ockenfels gratefully acknowledges the support of the Deutsche Forschungsgemeinschaft through the Emmy Noether-program.

1. Introduction

Many online markets rely on electronic reputation systems to promote trust in transactions. The reliance on reputation *per se* does not distinguish online markets from traditional markets, legal enforcement being in all cases expensive.[1] On balance, however, online markets appear to have more problems with fraud. A recent report by the research group GartnerG2 concludes that "Internet transaction fraud is 12 times higher than in-store fraud." A U.S. Department of Justice survey also cites high levels of online fraud, pointing especially at frauds common on auction sites (many with electronic reputation systems) that "induce their victims to send money for the promised items, but then deliver nothing or only an item far less valuable than what was promised (e.g., counterfeit or altered goods)."[2]

The power of reputation to promote trust in business transactions is closely associated with networked communities, places where there is a good deal of interpersonal communication as well as exchange. Online and traditional markets are networked in different ways. The laboratory experiments we present in this paper illustrate how differences in the flow of information and in the pattern of interaction can differentially influence trust. In discussing marketplace trust, Granovetter (1985) argues that people put more stock in information acquired "from one's own past dealings" than in information from outside sources. Our data, from experiments that simulate electronic markets, confirms this hypothesis, although for a reason different from those Granovetter put forward having to do with differing social ties or differing cost or quality of information. None of these are a factor in our markets. Instead, we find that the *use value* of reputational information depends on the pattern of interaction (and in ways that standard game theory does not anticipate). The important implication is that online markets

[1] Use of reputation to enforce trustworthiness is perhaps as old as human social interaction. Alexander (1987) argues that reputation mechanisms are at the base of human moral systems. Greif (1989) investigates the reputation systems used by certain groups of Mediterranean traders during the Middle Ages to facilitate trust among strangers. Milgrom et al. (1990) provide a historical as well as theoretical account of the role of fair for disseminating reputation information during the Middle Ages.

[2] For the GartnerG2 report, see http://www.gartnerg2.com/rpt/rpt-0102-0013.asp. The Department of Justice report is at http://www.internetfraud.usdoj.gov/#What Are the Major Types of Internet. Other evidence comes from the U.S. Internet Fraud Complaint Center (IFCC), a partnership of the U.S. White Collar Crime Center and the Federal Bureau of Investigation, that recently reported that 63 percent of the 49,711 formal complaints received in 2001 involved either non-delivery of merchandise, non-payment, or auctions, and that these numbers are rising rapidly.

might improve trust on their platforms by increasing the value of reputational information on their sites (more on this later).

In traditional business communities, the patterns of information flow and contact that promote trust often interact in subtle ways. Vietnam's free market reform in the mid-1980's gives an example that highlights the effectiveness of informal reputational controls: While there was little in the way of legal protection against exchange malfeasance, markets nevertheless flourished. According to McMillan (2002), "People in the same line of business would meet each other every day in teahouses and bars …to discuss the reliability of particular customers. … About half of a sample of entrepreneurs said that they had had no prior connections with the businesses that were to become long-standing trading partners." Thus, in traditional networked communities, interaction between members promotes trust in two ways. For one, information about individual reliability is transmitted to third parties, some of whom are prospective future trading partners. Second, the pattern of interaction promotes long-term relationships; a business partner whose trust has been validated is, all things equal, likely to return to do future business.

Online business platforms bridge spatial divides to broaden the pool of potential trading partners. Taking advantage of new matching opportunities, however, requires lesser reliance on long-term trading relationships; in data collected from EBay, Resnick and Zeckhauser (2001) found that 89 percent of all encounters during a five month period were one-shot. To promote the exchange of information on the reliability of individual traders, online business platforms, including Amazon, Cnet, Ebay, Half and Yahoo, have instituted electronic reputation mechanisms, known as "feedback" systems. Recent field studies of online auction platforms find that feedback systems have at least some of the desired economic effect in the sense that reputable sellers are more likely to sell their items (Resnick and Zeckhauser, 2001), and can expect price premiums (e.g., Lucking-Reiley et al., 1999).[3] The sale of used books on Amazon's

[3] Analogous results come from Ba and Pavlou (2002), Houser and Wooders (2001), Melnik and Alm (2001), and Ockenfels (2003); see Resnick et al. (2002) and Dellarocas (2003) for recent surveys. Brynjolfsson and Smith (2000) compared pricing behavior at 41 Internet and conventional retail outlets and also identify internet sellers' trustworthiness as one important factor that affects market outcomes.

site provides a simple but illustrative example of how these systems work. Amazon provides a market platform for independent dealers, brick-and-mortar bookstores as well as private individuals, to sell used books. Sellers state the price and describe the book's condition, and Amazon posts this information. Buyers pay through Amazon, who takes a percentage, but sellers ship directly to buyers. The moral hazard problems inherent in the seller's side of the deal – stipulating the book's condition and the shipping – is addressed through a system in which buyers are invited to post feedback on the transaction. Future buyers see a seller's feedback when deciding whether to make a purchase, and this provides an incentive for honest dealing.

Another way of stating the difference between online and traditional reputation networks is to note that they emphasize different kinds of reciprocity. In all cases, information about reputation enforces trust by inducing a reciprocal response: past trustworthiness is a prerequisite to future business. But there are two types of reciprocity: direct reciprocity, essentially long-term business relationships, and indirect reciprocity, business based on 'word-of-mouth' third party recommendations. Traditional markets are more conducive to direct reciprocal relationships; the repeated interaction common in these markets gives rise to long-term relationships. Online markets broaden the trading pool by lowering transaction costs; taking advantage of these new opportunities means moving away from long-term relationships. On the other hand, electronic feedback systems arguably improve the flow of reputational information necessary to indirect reciprocity; accessing online reputation information neither requires connections in the trading community nor direct contact with other traders. Later, we explain that standard economic theory implies indirect reciprocity can be just as effective as direct reciprocity, so long as traders in indirect reciprocal relationships have access to sufficient information about traders' reputations. Putting our investigation in these terms, then, we look at how well markets that rely on indirect reciprocal relationships build and sustain trust in comparison to markets that rely on direct reciprocal relationships.[4]

[4] See Brynjolfsson and Smith (2000), Dellarocas (2001), and Resnick and Zeckhauser (2001), for detailed comparisons of electronic and conventional dissemination of information about reputation.

Our experiment focuses on the role differing patterns of information flow and interaction have in these markets, and so our design carefully excludes other potentially confounding factors. For example, we control for the noise in feedback production (always truthfully provided in our experiment), we control the distribution of individual valuations and knowledge of these valuations in all markets, and we focus on the effect of reputation on the probability of trade keeping the price fixed across markets. Many of these factors, and their influence on trust, are interesting in their own right and deserving of broader investigation. For our purposes, however, it is desirable to hold them fixed. We add that an experiment allows us to separate effects in a systematic way (our experiment provides sufficient framework to take these other factors up in future investigations; see the summary section), and in a way that might be difficult to do in the field; see, for example, Dellarocas (2001) on the difficulty of accounting for the online reputation mechanisms create to manipulate feedback and to secretly switch identities.[5]

The moral hazard problem in the trading environment we study concerns shipping. In the experiment, a buyer chooses whether to purchase an item at a fixed price. If a purchase order is sent, the seller decides whether to ship or simply keep the buyer's money. In one market, we introduce a feedback system that tracks seller histories of shipping decisions and provides this information to prospective buyers (reputation market). To determine the value added by feedback we benchmark against the amount of commerce observed when no feedback is available (strangers market). Comparing strangers and reputation markets provides a measure of the improvement in trust (decision to buy) and trustworthiness (decision to ship) feedback brings about. We then compare the performance of the reputation market with that of a market where the same two people interact repeatedly (partners market).

The comparison between the strangers and the reputation market will confirm that the information about past seller behavior alone is sufficient to produce economic benefits. One might have thought that information about reputation would increase trading efficiency only if,

[5] For more discussion on the relationship between experimental and field studies see Roth (2002) and Ariely et al. (2002).

for example, reliable sellers were able to charge higher prices; our test indicates that this is not the case. At the same time, the comparison between the reputation and the partners market (a comparison that would be difficult, if not impossible, to conduct in the field) suggests that reputation mechanisms suffer a kind of public goods problem, in that the benefits of trust and trustworthiness induced by the reputation mechanism go to the whole community and are not completely internalized. Later, we discuss the managerial implications of this finding.

2. Experimental design and some theoretical considerations

2.1 The basic market platform

The same set of buyer-seller encounters is common to all three markets of the experiment. Each time there are 16 traders interacting over a period of 30 rounds. At the beginning of each round, each trader is matched into a pair, with one person assigned the role of buyer and the other the role of seller. The matching process varies with market, as does the information the buyer has about the seller (variations are taken up below). The role assignment in all markets is random under the commonly known restrictions that each trader is a buyer half the time and a seller half the time.

Figure 1. The buyer-seller encounter

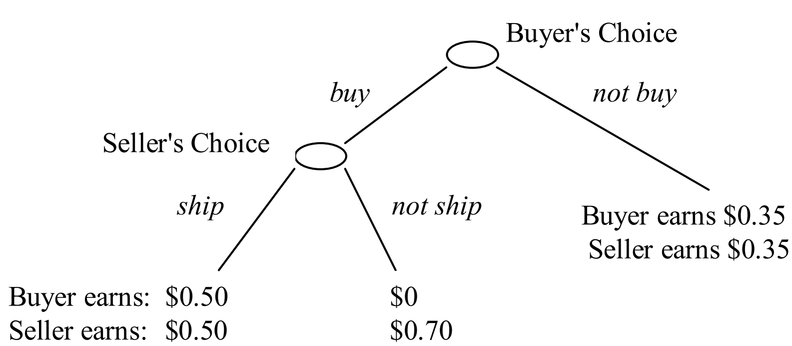

Figure 1 illustrates the buyer-seller encounter. Game moves and payoffs can be interpreted in the following way. Both the seller and the buyer are endowed with $0.35, which is the payoff when no trade takes place. The seller offers an item for sale at a price \$0.35 which has a value of \$0.50 to the buyer. The seller's cost of providing the buyer with the item, costs associated with executing the trade, shipping, handling etc., as well as production costs,[6] is \$0.20. So each successfully completed trade increases efficiency by creating a consumer surplus of \$0.15 and a \$0.15 profit for the seller. If the buyer chooses to *buy* the item, he sends his endowment of \$0.35 to the seller, who then has to decide whether to ship the item, or whether to keep both the money and the item. If the seller does not ship the item he receives the price plus his endowment of \$0.35 for a total of \$0.70. If he ships, he receives the price minus the costs plus his endowment for a total of \$0.50, whereas the buyer receives his value of the item. If the buyer chooses not to buy the item, no trade occurs. In this sense, the buyers can choose with whom to trade and with whom not to trade.

Under the assumption that the buyer-seller encounter is one-shot, the seller, once he receives the money from the buyer, has no pecuniary incentive to be trustworthy and to ship the item. Anticipating this, the buyer may not trust the seller, so that trade does not take place, even though it would make everybody better off. This is the basic dilemma online reputation mechanisms are designed to solve.

2.2 How the three markets differ and the experimental protocol

In the strangers market, buyers and sellers are randomly paired in each round under the commonly known restriction that nobody is matched with the same player in the same role more than once. Buyer and seller are anonymous to one another, and no information about one another's history of action is made available. The random matching scheme is public knowledge.

Matching is done the same way in the reputation market (we used the same pairing rotation used in the strangers market; in particular, no buyer is matched more than once with a

[6] Production costs where either the seller only produces the item once he knows the demand, or the product is produced before the buyer's decision is known but costs are not sunk (e.g., when the item can be resold at a price equal to production costs).

specific seller), but here buyers are given the feedback on the seller's shipping decisions prior to choosing whether to buy. The feedback includes a summary of the number of times the seller shipped in the past, as well as a round-by-round history of their shipping decisions. Appendix A includes a typical buyer's computer screen.

In the partners market, matching is fixed; that is, the same two subjects are always paired together, and this is public knowledge. The same kind of feedback information available to reputation market buyers is available to partner market buyers. In this way, the information structure is parallel across the two markets, albeit the information is redundant in the sense that it is telling the buyers things the buyer has himself experienced.

Each of the three markets consists of three sessions. There are 16 subjects per session (48 per market) for a total of 144 participants in the experiment. All sessions were conducted in March and April of 2002[7] at the Laboratory for Economic Management and Auctions in the Smeal College of Business, Penn State University. Subjects were Penn State University students, mostly undergraduates, from various fields of study who volunteered through an on-line recruitment system. Cash was the only incentive to participate. Upon arrival at the laboratory participants were seated at the computers, separated by partitions. They were asked to read the instructions. (See Appendix B for the written instructions given to subjects, and illustrative examples.) To create public knowledge, the monitor read instructions to subjects out loud, after which consent forms were signed and collected. Subjects then played several practice games in a sequence of roles that was chosen at random, with the computer as partner making its moves at random.[8] Once familiar with the game interface, subjects played the 30 actual rounds. Upon completion of the session, each subject was privately paid his or her earnings in cash plus a $5 show-up fee.

---

[7] The partners sessions were held on 3/27/2002 1 PM, 3/27/2002 2 PM and 3/39/2002 10 AM; the reputation sessions were held on 3/15/2002 1 PM, 3/15/2002 2 PM, and 4/08/2002 4 PM; the strangers sessions were held on 3/21/2002 9 AM, 3/21/2002 10 AM, and 3/29/2002 11 AM.

[8] In order to encourage subjects to explore the features of the game interface (the point of practice), practice game payoffs were displayed as the Marx brothers: Chico, Groucho, Harpo and Zeppo.

2.3 Theoretical considerations

The comparative static predictions across markets depend on the complexity of the underlying model. Standard models suggest that cooperation rates should not differ across treatments unless the amount of information differs. The simplest (arguably too simple) Nash equilibrium analysis assumes complete information, and suggests no trade in any of the three markets: Consider again Figure 1 and suppose that we are in the 30th and final round of play. In terms of pecuniary rewards, the seller's optimal action, regardless of the game's history, is *not ship*, and so the buyer's optimal action is to exercise the outside option. This remains true regardless of what seller feedback the buyer has, or of whether the buyer and seller are randomly matched or partnered. Backing through the game then shows that, in Nash equilibrium, there can be no trading in any of the 30 rounds: Since there is no trading in the 30th round, there is no incentive to ship in the 29th round, again independent of feedback or matching considerations, and so again the buyer's optimal action is the outside option, and so on back to the 1st round.

There are strong theoretical, as well as behavioral, reasons to believe that this simple analysis overstates the difficulty of trading.[9] For one, the no-trade equilibrium is not robust to minor perturbations with respect to plausible player beliefs about others' behavior and payoffs – at least for the markets were information about seller past histories is available. Kreps et al. (1982) demonstrate that if each player assesses a (small) positive probability that his partner is 'cooperative' (i.e., he prefers to cooperate (defect) if the other cooperates (defects)), then sequential equilibria exist wherein perfectly rational players cooperate until the last few stages. Intuitively, if rational players believes there are even a small number of cooperative-types, then they have an incentive to cooperate themselves in order to encourage cooperative-types to cooperate in future rounds. Specific to our game, suppose buyers believe there are a small

[9] In fact, as we will see in the next sections, there is a lot of trading in all markets that systematically responds to the different available information channels across markets. Similar deviations from the simple theory have been observed in numerous finitely repeated prisoners' dilemmas where cooperation has been shown to be quite robust (ex., Andreoni and Miller, 1993, Selten and Stöcker, 1986, Ledyard, 1995).

proportion of sellers who will ship, say, out of a sense of social obligation.[10] They buy only if the seller has shipped for orders received in the past; giving sellers incentives to ship in order to continue to receive buy orders in the future. Note that this argument requires the buyer to be able to observe what the seller has done in the past – a condition that holds in the reputation and partner market, but not strangers (the same is true of Kreps et al.'s model). This suggests, then, more trading in reputation and partners markets than in strangers, but offers little reason to think that trading levels will differ across reputation and partners.[11]

A second reason to believe there would be more trading than suggested by the complete information Nash argument is that people have been shown to be more myopic than assumed by the backwards induction argument (e.g., Selten and Stöcker, 1986). Rather than go to the end of a long game and reason back, most people appear to have trouble reasoning more than one or two steps ahead. Myopia has less bite with the strangers market because in this case all a seller need realize is that the decision made now will not be observable to future partners, making each encounter, in essence, a separate game. In the reputation and partners markets, however, shipping behavior is reported to future partners, information the seller can dismiss only if he does the complete backwards induction – and is confident that future buyers will do so as well. Consideration of myopic behavior, then, suggests the same pattern of trading as consideration of incomplete information.[12]

---

[10] Bolton and Ockenfels (2000) and Fehr and Schmidt (1999) show that a model in which people are assumed to care about relative as well as pecuniary payoffs can explain data patterns across a wide variety of experimental games, including bargaining, social dilemma, and market games. These models suggests that some (possibly very few) sellers will ship when they are trusted, even in the 30th round. Explicit connection to the Kreps et al. model is made in Bolton and Ockenfels (2000, Section VI.D). Güth and Ockenfels (forthcoming) review the economic theory literature on the evolution of preferences in trust games not unlike the one we study. They show that depending on the details of the institutional environment trustworthy behavior may survive evolutionary competition. These two complementary lines of research suggest that trustworthiness and trust can also have non-strategic causes and can be exhibited even in one-shot encounters or in the last round of repeated interaction. Our experiments measure the extent to which trust and trustworthiness is intrinsically or strategically motivated.

[11] As with Kreps et al.'s model, the formal incomplete information model would be technically demanding, and we will not attempt to work such out here, but see Bolton and Ockenfels (2003).

[12] To give a more precise example, assume that myopic subjects play the game as if it is infinitely repeated. It is easy to see that in this case simple trigger strategies that call from buyers to punish untrustworthy sellers by never trusting them again support cooperative equilibria in the partners and reputation markets if the discount rate is sufficiently high (see e.g. Ockenfels, 2003, for the theory in a related scenario), but not in a strangers market when no information about the opponent's history is available (regardless of the players' discount rates).

None of the arguments mentioned so far suggests a difference of trading patterns between our experimental partners and reputation markets. The reason is that information about past behavior − and not the matching scheme − drives future cooperation in these models, and that our experiment provides accurate information both in the reputation and in the partners markets: the complete reputation profile of the sellers is always truthfully provided to buyers. We will see that this prediction stands in stark contrast to our experimental result; traders on the partners markets generate substantially higher cooperation rates than traders on the reputation markets. As we will argue, this may be due to a kind of public goods problem on reputation markets in that the benefits of trust and trustworthy behavior go to the whole community and are not completely internalized. In Section 3.6 we explain that a more complex model than the ones mentioned here may capture this phenomenon.

One final theoretical consideration: That our game has a finite number of rounds makes for a stringent test of reputation systems to induce reputation building, because the strategic value of having a good reputation will diminish as the last round approaches (even if all or some players are quite myopic). We could have designed a finite game to simulate an infinitely repeated game by inserting some probability of stopping after each round of play.[13] In these games, there are equilibria in which rational actors – assuming none are too risk averse – cooperate every round (although not cooperating in *any* round is always an equilibrium as well). But now the market structure overstates things in the opposite direction: In reality, at any given time, there are market actors who are in the market short term, perhaps for the last time. In our set-up, all traders exit the market at a publicly stated time, and in this sense it is the tougher test.

3. Experimental Results

We first describe the basic treatment effects we observed. Theory suggests that the key to understanding differences in transaction behavior across markets is how buyers react to

[13] While this would eliminate the stopping issue, it would also introduce risk-over-stopping as a factor, the influence of which is not directly observable and so difficult to measure. Alternatively, we might not reveal the ending round to subjects, but this would have the same drawbacks.

information about the sellers they are paired with. If buyers discriminate between sellers who have been trustworthy in the past and those who have not been, then sellers will have an incentive to be trustworthy. For this reason, we look at buyer behavior in some depth.

3.1 Treatment effects

The major treatment effects have to do with trading patterns. These can be measured three ways: *efficiency* or the percentage of potential transaction completed (Figure 2), *trust* or the percentage of buy orders given (Figure 3), and *trustworthiness* or the percentage of shipped items, conditioned on buy orders (Figure 4). In all three figures, the treatment data has been aggregated across sessions.

Figure 2. Efficiency measured as how often the gain from trade is realized, by round

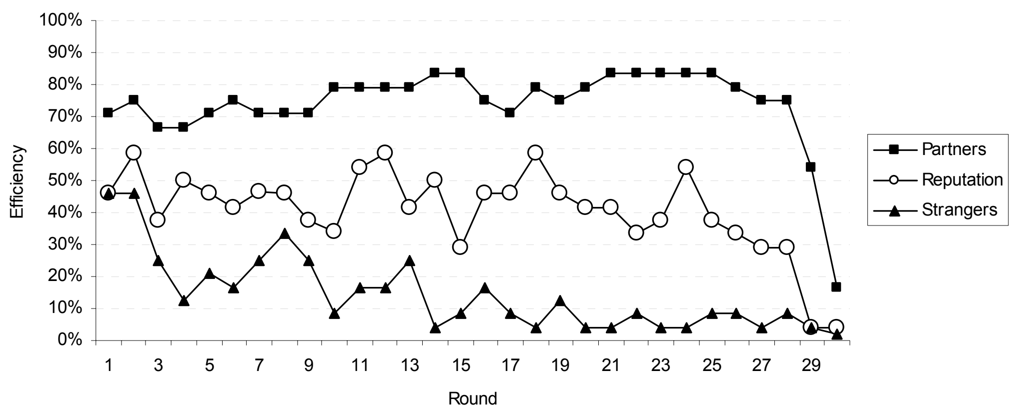

Figure 3. Trust measured as the percentage of buying per round

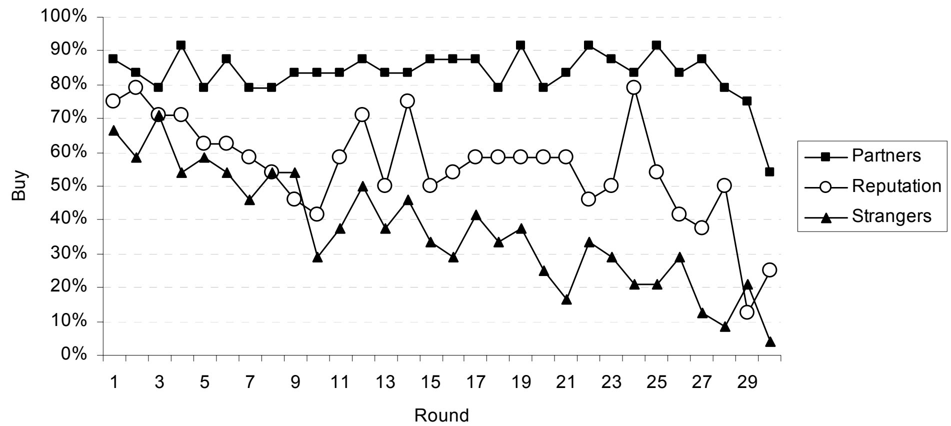

Figure 4. Trustworthiness measured as percentage of shipping per round

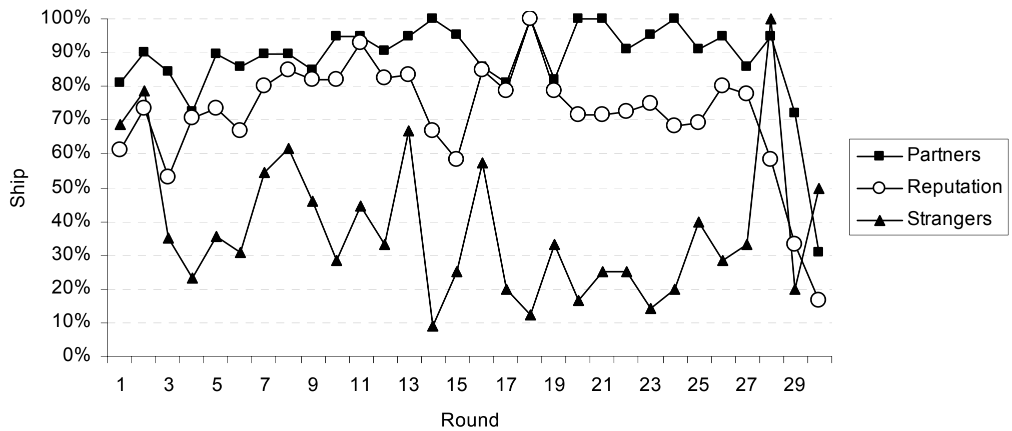

The same pattern is evident in all three figures: There is the least efficiency, trust and trustworthiness in the strangers market, more of all three in the reputation market, and more still in the partners market. For instance, averaged over all rounds, reputation yields 2.8 times the efficiency of strangers, and partners yields 1.8 times the efficiency of reputation. Pair-wise *t*-tests, in which sessions are taken as the individual observations, verify these observations ($p < 0.025$ in all cases save for comparing average buying between strangers and reputation, $p = 0.08$, one-tail tests, assuming equal variances).[14]

[14] A non-parametric rank test on session observations yields similar results, with $p = 0.05$ for all but the strangers-reputation comparison where $p = 0.10$ (one-tail tests).

Table 1. Percentage of efficiency, buying, and shipping, by session

Treatment	Session	Efficiency	Buy	Ship*
Strangers	I	0.158	0.391	0.411
	II	0.046	0.204	0.224
	III	0.225	0.517	0.435
	Mean	0.143	0.371	0.357
	coef. var.	0.633	0.424	0.323
Reputation	IV	0.517	0.671	0.767
	V	0.367	0.538	0.682
	VI	0.338	0.458	0.736
	mean	0.407	0.556	0.728
	coef. var.	0.236	0.193	0.059
Partners	VII	0.738	0.833	0.885
	VIII	0.671	0.792	0.847
	IX	0.808	0.875	0.924
	mean	0.739	0.833	0.885
	coef. var.	0.093	0.050	0.043

* conditioned on buying

Table 1 breaks the data out by sessions and reveals how the markets' differences in matching schemes and information flows affect the variability of trading outcomes. For all gauges (efficiency, buy, ship), the coefficient of variation is highest for strangers, less than half as large for reputation, and then smaller still for partners. Thus, on average, trading patterns are less reliable in strangers markets, more so when a reputation mechanism is available, and more so still in partners relations.[15]

There are also marked differences in the trading dynamics across treatments. Observe from Figures 2, 3 and 4 that while there is a clear downward trend in efficiency, trust and trustworthiness over all rounds in the strangers treatment, trading volumes appear to be rather stable in reputation and partners, save for the very last two rounds when trading collapses.

[15] Table 1 also shows that trustworthiness levels in reputation and partners are closer to each other than trust levels, which may be due to the self-selection effect caused by the available reputation information. While the numbers in the table measure the trust level of *all* buyers, they only incorporate the trustworthiness of *trusted* sellers. The more successful buyers are in separating trustworthy from untrustworthy sellers, the closer one would expect the trustworthiness levels in partners and reputation markets.

Table 2. Random effects probit models, buyers[a]

Maximum likelihood estimates (and two-sided *p*-values) for buyer behavior
Dependent variable = "1" for *buy*

Independent variable	Model 1	Model 2	Model 3
CONSTANT	0.533	0.347	0.524
	(.0040)	(.0185)	(.0001)
REPUTATION	-0.020	0.200	
= 1 if buyer is from reputation treatment, and 0 else.	(.9473)	(.4347)	
PARTNERS	0.963	1.48	0.852
= 1 if buyer is from partners treatment, and 0 else.	(.0001)	(.0000)	(.0011)
TOTALSHIPreputation			0.0616
= number of seller ships prior to last order.			(.0014)
TOTALNOSHIPreputation			-0.124
= number of seller no ships prior to last order.			(.0144)
SHIPLASTreputation			0.212
= 1 if reputation seller shipped last order, and 0 else.			(.2111)
NSHIPLASTreputation			-0.646
= 1 if reputation seller did not ship last order, and 0 else.			(.0005)
SHIPLASTpartners			1.330
= 1 if seller in partners shipped last order, and 0 else.			(.0000)
NSHIPLASTpartners			-.697
= 1 if seller in partners did not ship last order, and 0 else.			(.0100)
CBSH		0.045	-0.005[b]
= number of past times item was shipped to buyer.		(.0180)	(.8386)
CBNH		-0.412	-0.386[b]
= number of past times buyer bought but not shipped.		(.0000)	(.0000)
ROUNDstrangers	-0.062		
= round in strangers treatment, and 0 else.	(.0000)		
ROUNDreputation	-0.019		
= round in reputation treatment, and 0 else.	(.0006)		
ROUNDpartners	0.006		
= round in partners treatment, and 0 else.	(.4806)		
LAST2ROUNDstrangers	-0.151	-0.390	-0.404
= 1 if round 29 or 30 in strangers treatment, and 0 else.	(.6414)	(.1649)	(.1671)
LAST2ROUNDreputation	-0.903	-0.944	-0.974
= 1 if round 29 or 30 in reputation treatment, and 0 else.	(.0000)	(.0000)	(.0000)
LAST2ROUNDpartners	-1.15	-1.200	-1.322
= 1 if round 29 or 30 in partners treatment, and 0 else.	(.0000)	(.0000)	(.0000)
RHO	0.399	0.456	.444
(random effects)	(.0000)	(.0000)	(.0000)
Number of observations	2160	2160	2160
Log-likelihood	-1087.77	-1056.57	-988.67
χ^2 *p*-value	.0000	.0000	.0000

[a]Analogous estimates for fixed effects linear models are given in Appendix B.
[b]History for Partner's buyers does not include last transaction.

Model 1 in Table 2 supports these observations for the buyer behavior with the help of a random effect probit regression.[16] There is not only much more trust *per se* in partners than in

[16] The more sophisticated Models 2 and 3 examine *individual* trading patterns and will be discussed below. Analogous probit models for the seller behavior are omitted but yield the same qualitative conclusions (see also Table 3 below).

reputation and strangers, as indicated by the treatment dummy PARTNERS, but controlling for end-game effects the trust shown by partners is remarkably stable over time: The ROUNDpartners coefficient is small and not significant. ROUNDreputation is also small, but significantly negative, indicating a slight downward trend of trust in reputation. However, it is still significantly larger than ROUNDstrangers (two-tail p = 0.0106, Wald test). Since the coefficient for the dummy REPUTATION is close to zero and not significant, Model I explains the larger percentage of trust in reputation by the fact that trust in strangers declines more rapidly over time. At the same time, there are large and significant end-game effects in both reputation and partners but not in strangers, indicated by the LAST2ROUND variables.[17] In the next sections we examine the underlying causes for these strong treatment effects by looking at how individual behavior varies across markets.

3.2 Comparing strangers and reputation: The strategic benefits of trust and trustworthiness

Theory suggests a reputation mechanism can help to realize efficiency enhancing trade on large, anonymous market platforms *if* buyers condition their behavior on the shipping history of their sellers, so that sellers have strategic incentives to avoid spoiling their reputation.[18] The data supports this view.

[17] For endgame effects, we focus on the last *two* rounds because, while player roles were switched randomly, they were told they would be a seller (buyer) for half the rounds, so that in round 29 a seller may be in his last round as a seller and thus has no strategic reason to be trustworthy.

[18] As Ebay puts it: “A user's feedback is a key factor people use to determine whether or not they want to trade with that user. What feedback you give or receive is an important part of your trading reputation at ebay. […] If you're a buyer, checking a seller's Feedback Profile before you make a bid is one of the smartest and safest moves you can make. This Feedback Profile answers many questions about how a seller does business.” (http://pages.ebay.com/services/forum/feedback.html).

Figure 5. Marginal trust conditioned on last feedback across treatments*

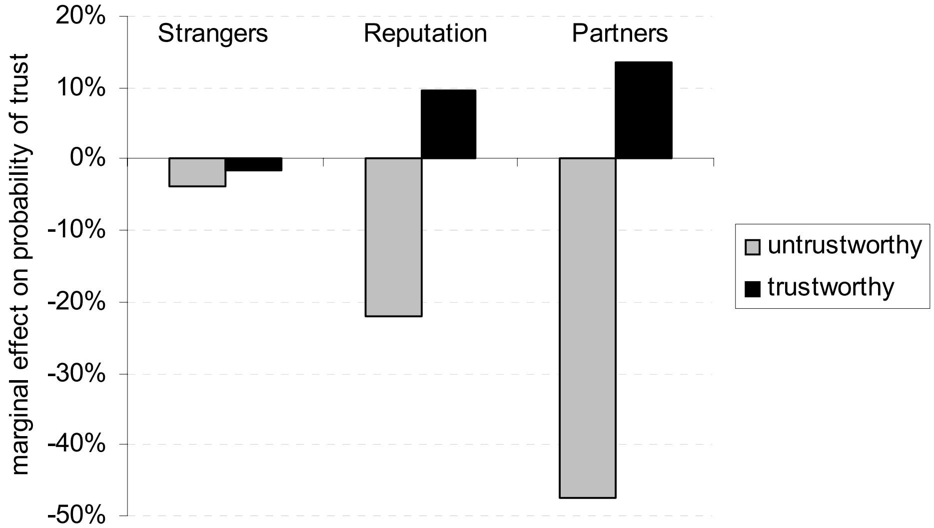

* The base rate (the zero line) is the average buy over all encounters for each treatment separately (37.08 percent in strangers, 55.56 percent in reputation and 83.22 percent in partners).

Figure 5 shows the marginal effects on the probability of trust in all markets depending on whether the seller shipped the *last* order ('trustworthy') or not ('untrustworthy').[19] Strangers cannot distinguish between whether the seller shipped the last order or not, so that the marginal effect is close to zero regardless of the seller's history. (That the effects are negative in both cases is due to the facts that the bars do not include "newbies" − sellers who have not been trusted yet and who are therefore typically encountered in early rounds − and that buy rates in strangers decline rapidly.) Buyers in reputation, on the other hand, strongly condition their behavior on the seller's last feedback. In absolute terms, they trust with probability 33 percent if the seller did not ship the last order, and with about twice this probability, 65 percent, if he

[19] Figure 5 and the next Figure 6 are illustrations of the main effects. A more thorough statistical analysis of buyer behavior across treatments that takes into account the whole feedback history and that controls for other important factors will be discussed in Section 3.4. Keser (2002) independently studied feedback systems in the lab. She found that when feedback is endogenously and voluntarily given by trading partners, efficiency is slightly higher when trading partners are informed of the entire distribution of each other's previous feedbacks than only of the most recent feedback. However, her results are not easily comparable to our results because her experimental design differs in a number of design choices, such as the game itself (e.g., efficiency gains occurred after the buyer's move and were substantially higher), the matching and role assignment scheme (e.g., trading partners were matched more than once), feedback provision (voluntarily and endogenously given by buyers), etc.

shipped the last order. Statistical models in Table 2, to be discussed in a moment, support that trust is strongly conditioned on reputation information.

Conditional buying is sensible since the seller's history has predictive power for his future performance. Table 3 presents a random effect probit for sellers. We can see that shipping the last time both a reputation and partner market seller received a buy order is a significant predictor of whether the seller will do so this time. (The coefficient for LASTSHIPstrangers is significant as well but with a negative sign.) Further, a last decision to ship is more highly predictive of shipping this time in partners than in reputation markets (two-tailed $p = 0.0121$, Wald test).

Table 3. Random effects probit model, sellers

Maximum likelihood estimates (and two-sided p-values) for seller behavior

Dependent variable = "1" for *ship*

Independent variable		Independent variable	
CONSTANT	0.190	SHIPLASTstrangers	-0.366
	(.2818)		(.0713)
REPUTATION	0.278	SHIPLASTreputation	0.350
	(.2666)		(.0168)
PARTNERS	0.557	SHIPLASTpartners	0.898
	(.0267)		(.0000)
ROUNDstrangers	-0.037	LAST2ROUNDstrangers	-0.112
	(.0008)		(.8711)
ROUNDreputation	-0.005	LAST2ROUNDreputation	-1.757
	(.6154)		(.0066)
ROUNDpartners	-0.007	LAST2ROUNDpartners	-1.833
	(.4589)		(.0000)
RHO	0.204	Number of observations	1267
	(.0000)	Log-likelihood	-561.63
		χ^2 p-value	.0000

[a]Variable interpretations are analogous to those given in Table 2.

Finally, the strong end-game effect exhibited by both buyers and sellers in reputation (Tables 1 and 3, respectively) additionally supports the view that it is the strategic incentive to ship created by the 'shadow of the future,' as opposed to, say, an intrinsic preference for being trustworthy, that largely drives the efficiency-enhancing effect of the reputation mechanism.[20]

[20] Of course, the fact that sellers in strangers ship some 36 percent of the time indicates that not all trustworthy behavior can be explained by strategic response to the pecuniary incentives.

Returning to Figure 5, buyers in the partners market respond even more strongly to the sellers' histories than do buyers in partners. Why is it that buyers in reputation markets, who have access to similar history information, are less affected? The next section offers an answer.

3.3 Comparing reputation and partners: The public benefits of trust and trustworthiness

The central thrust of our arguments in this section will be that, unlike in a partner relationship, feedback and past experience do not perfectly overlap in reputation markets. This creates effects of trusting and being trustworthy that are not internalized by the feedback mechanism in the reputation market as they are in the partners market. In this section, we explain this phenomenon with some simple descriptive statistics to illustrate the points. In the next section, we discuss the analogous inferential evidence.

First, observe that a trusting buyer in a reputation market generates valuable feedback information for *other* buyers who meet the same seller in the future. A trusting buyer in a partner relationship generates the same valuable feedback information − but entirely for himself. Thus, the informational benefits from trusting (as opposed to the pecuniary gains from trade) are internalized in trades among partners but not in a reputation market. As a consequence, all other things equal, the overall benefits from trusting are smaller in reputation.[21]

Empirically, this informational dilemma is particularly apparent if a buyer is matched with a newby, a seller with no feedback history yet, because then there is no evidence of whether the opponent can be trusted or not. The average trust in newbies is about 65 percent in strangers, 77 percent in reputation and 93 percent in partners.[22] While the difference between the strangers market and the other markets can be explained by the fact that newbies in the strangers market have no strategic reason to be trustworthy, the difference between reputation and partners may be

[21] There is a related but distinct public goods problem of feedback provision observed Resnick and Zeckhauser (2001), among others: Once the transaction is concluded, buyers have no incentives to provide others with feedback about their experience. (Resnick and Zeckhauser report that on Ebay nevertheless feedback was provided in about half the time.) Since in our experiments feedback was produced automatically, this public good problem was not an issue in the lab (and in this sense our results overestimate the merits of a reputation mechanism that is based on *voluntary* feedback production).
[22] Four buyers in reputation never faced a newby and so are not included in these statistics.

explained by the informational dilemma. In fact, we observe that buying from a newby in the reputation market yielded an expected payoff of 31 cents, less than the 35 cents from not buying. Buying from a seller who shipped the last order, on the other hand, yielded an average payoff of 40 cents, and buying from a seller who did not ship the last order yielded 17 cents.

Thus, a buyer in the reputation market is better off trusting somebody if he or she has already been shown to be trustworthy. But as a matter of empirical fact, a buyer does not have an incentive to generate the information that is needed by the market to discriminate between trustworthy and untrustworthy sellers, because the production of feedback is costly but not beneficial to the buyer. In contrast, trusting buyers in partners markets only produce feedback information for their private use, say, to execute a simple conditional buying strategy. Both costs and benefits are fully internalized.

In reputation markets, there are not only public benefits from trust but also from trustworthiness. The critical evidence for this comes from examining buyers' reactions to their *own* history. Extending the scope of Figure 5, Figure 6 (below) additionally provides some sense of the relative importance of own experience, defined either as 'rewarded' if the buyer's *last* order was shipped, or 'exploited' if not (later, our statistical models will provide more formal evidence). The figure shows that, regardless of the market environment, if a buyer was treated well in the past he is more likely to trust in the future. While strangers cannot distinguish between trustworthy and untrustworthy sellers, they do condition their behavior on whether they were rewarded or exploited in the past: Rewarded buyers trust substantially more often than exploited buyers. In reputation and strangers least trust is shown if the buyer was exploited and faces an untrustworthy seller and most trust is shown if he was rewarded and faces a trustworthy seller.

There is, however, an important difference between reputation and partners markets: In partners markets, histories are perfectly aligned. That is, a buyer's trust is rewarded if and only if his seller is trustworthy, so that the two history effects always cumulate. In reputation, on the other hand, the feedback effect is diluted by the effect from the buyer's own history. A trustworthy seller is less trusted when the buyer's own experience was bad and an untrustworthy seller is more trusted when the buyer's own experience was good. Thus, there is a wedge driven between buyers' and sellers' histories in reputation markets that is (at least partly) responsible for why the marginal response to both the seller's positive and negative feedback is on average weaker in reputation compared to partners (as shown in Figure 5). The wedge does not explain, however, why overall average trust is smaller in reputation. But note that because both the seller's and the buyer's history affect the evolution of trust, trustworthiness creates a public benefit in the reputation market. A seller who ships in the reputation market benefits only through the improved own reputation, but *other* sellers will profit from the induced good history of his buyer. On the other hand, a seller who ships in partners not only contributes to his own good reputation but also contributes to a good history of *his own* future buyer. In other words, the trust-inducing effects of trustworthiness are fully internalized in partners but not in reputation. As a consequence the overall private benefit from trustworthiness is weaker in reputation markets, which reduces the incentive to trade of both sellers and buyers.

Figure 6. Marginal effects of the buyers' most recent history and the sellers' most recent feedback on the buying probability*

Strangers

average buy = 37 percent

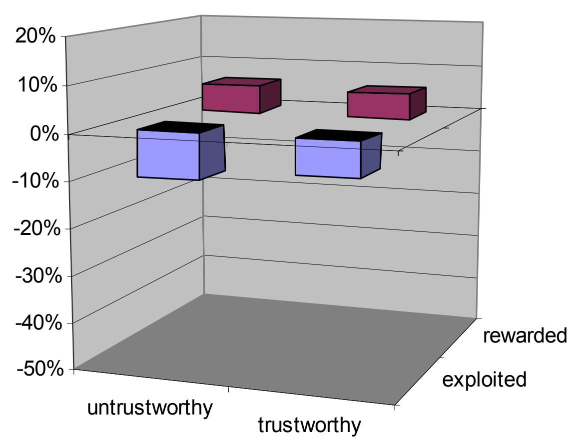

Reputation

average buy = 56 percent

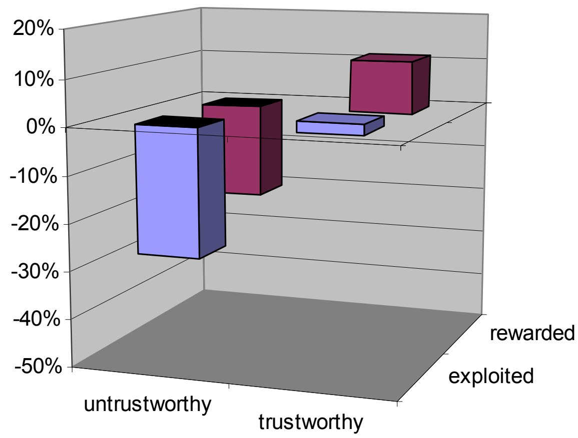

Partners

average buy = 83 percent

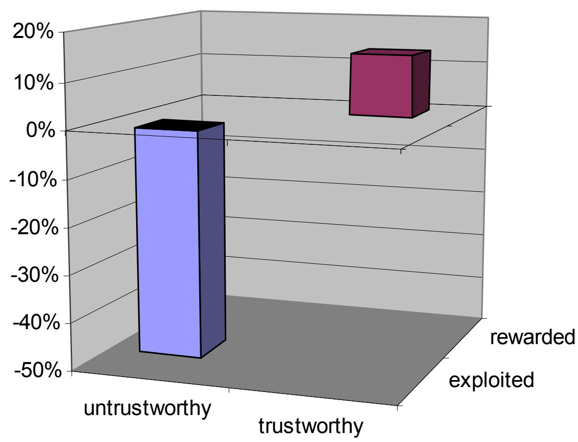

* As in Figure 5, in each graph, the zero level is normalized with respect to the respective "average buy" that includes all encounters.

3.4 A probit model of buyer behavior

Models 2 and 3 in Table 2 provide formal inferential support for the findings in the last section. They also provide some further insight into how buyers evaluate seller histories.

Model 2 replaces the round effect variables with variables reflecting the buyer's history with shipping: the cumulative number of times he has bought and had the item shipped (CBSH) and the cumulative number of buys that were not shipped (CBNH). Figure 6 suggests that buyer decisions are influenced by their experience with seller shipping. The alternative hypothesis is that the differing learning effects we see in Figure 3 are due to the differing rules of the three markets; so, for instance, there is less information for buyers to process in strangers and this is plausibly why the experience effect is greater in strangers than in the other treatments (see Section 3.1). If differing rules, not buyer experience, were responsible for the dynamic across rounds, then we would expect both C¨H coefficients in Model 2 to be about the same, and certainly Model 2 should fit no better than Model 1. In fact, the C¨H coefficients differ sharply and the likelihood of Model 2 is higher than that of Model 1, indicating that buyer experience is quite a good explanation for the differing dynamics across markets.[23] What we see in Model 2 is that negative experiences – making a purchase and having a seller fail to ship – erodes buyer trust in the market, while positive experiences have substantially less effect.

The influence of buyer history persists in Model 3, where we add variables reflecting the information buyers have about sellers at the time they decide whether to purchase. For the reputation market, we add variables to reflect the cumulative shipping history of the seller (TOTALSHIPS and TOTALNOSHIPS) as well as variables to reflect the most recent shipping history (SHIPLAST and NSHIPLAST). In the partners market, the seller's cumulative history is already reflected in the buyers C¨H variables. So for partners, we need add only most recent history variables. We also drop the REPUTATION variable since this has not been significant in the other two models.[24]

All of the information coefficients in Model 3 have the expected signs, save that for CBSH, but this estimate is very small and not significant. There are two main observations. First, both reputation and partner buyers weight recent observations more heavily than older ones (p = .0000 for both markets, Wald test). Second, in reputation markets, NSHIPLAST has a more reliable effect on buyer decision than SHIPLAST, whereas partner buyers weight recent positive and negative history about the same; in fact, SHIPLAST and NSHIPLAST are, in absolute terms, virtually identical.

[23] Models 1 and 2 are not nested and so not conducive to a formal inference test of the buyer experience hypothesis. If we re-estimate model 1 with three CNBH*x* variables, one for each treatment *x*, all three of the new variables are negative and significant (two-tail $p < 0.01$ in all cases). A full model, with history variables broken out by treatment, is given in Appendix B.

[24] Including it would make little difference to the estimates of Model 3 and the coefficient of REPUTATION would still not be significant. There are, however, some indications that REPUTATION is collinear with SHIPLASTreputation and NSHIPLASTreputation.

We also see the information wedge effects described in the last section: First, buyers are more likely to trust newbies in the partners market than in the reputation market as indicated by the significant coefficient for PARTNERS (while REPUTATION is insignificant). Second, a partner seller who has been trustworthy in the recent past is granted higher trust (SHIPLASTpartners is greater than SHIPLASTreputation, two-tail p = 0.0007, Wald test). Third, we can see the greater incentive sellers have to be trustworthy due to the internalization of buyer history. Note that a decision *not to ship* has very similar immediate total negative effect on buying in both Reputation and Partner markets: compare NSHIPLASTrep + CBNH to NSHIPLASTpar + CBNH. The difference is that the seller in the systems incurs only about 65 to 70 percent of this cost since the CBNH becomes part of the history of a buyer he is unlikely to meet again in the near future.

3.5 Payoffs

One would correctly conclude from the analysis of the efficiency reached by the different markets in Section 3.1 that average payoffs are smallest in strangers, larger in reputation, and larger still in partners. In fact, the strategic incentives created by reputation systems also translate into positive *correlations* between trust(worthiness) and payoffs. The Spearman rank correlation coefficients between subjects' total payoffs and the total number of ship and buy in the strangers markets is negative (−0.307, p = .034, for the frequency of buy and −0.038, p = .795, for ship), making all trade activities unprofitable. Sixteen out 48 subjects receive total payoffs that are smaller than in the Nash equilibrium in which nobody ever buys or ships. A reputation mechanism, on the other hand, makes trust and trustworthiness lucrative behavior. The corresponding coefficients in the reputation market are significantly positive (.288, p = .047, for buy and .504, p = .000, for ship) and not a single subject made payoffs smaller than equilibrium payoffs. Finally, the pecuniary incentives to trade are strongest in partners markets (.706, p = .000, for buy and .728, p = .000 for ship). While two subjects in this market received less than they could expect in equilibrium (both subjects never shipped to their partners but vainly tried to buy from them), 34 buyers and 31 sellers did their part of the exchange 14 or 15 times, numbers that are reached only once in strangers and reputation together.

3.6 Approaches to the public good aspect of trust and trustworthiness

Our data reveal that there are external effects of trust and trustworthiness that diminish the strategic incentives to cooperate in reputation markets compared to partners markets. The external effect of trust is due to the empirical fact that a seller's reputation profile has, starting in the first round, predictive value for his future behavior. This is inconsistent with the models outlined in Section 2.3. In particular, in standard sequential equilibrium approaches to our game, reputation information, though critical for the emergence of trust, is *not* valuable in deciding whether to trust or not: In the early phase of the finitely repeated game all sellers always cooperate regardless of their preference-type (trustworthy or strategic) so that in equilibrium early play cannot reveal valuable information. In all other periods, strategic sellers either mix such that buyers are indifferent between trusting and not trusting, or, once their reputation profile proves them as strategic players, never ship anymore. Thus, in this second phase, a buyer cannot make more than his outside option regardless of his seller's reputation profile. However, the dynamics of reputation effects are sensitive to the set of preferences that exists (see, e.g., Diamond, 1989). It is well conceivable that in a model in which some sellers are committed to behave trustworthy in all encounters, but some others are committed to behave always untrustworthy (and the rest behaving strategically), a seller's reputation is critical information for the buyer's behavior. Then, depending on the distribution of seller-types, it may be unprofitable to trust a newby but profitable to trust a 'reputable'seller.[25]

The external effect of trustworthiness is due to the empirical fact that buyers condition their behavior on their own histories. Again, this is not in line with sequential equilibrium approaches if one assumes that the seller-type distribution is commonly known. However, if

[25] Avery, Resnick and Zeckhauser (1999) developed a related model of evaluations for goods with fixed but initially unknown quality. The more feedback is available, the less the uncertainty about the true product's quality.

buyers update their beliefs about the distribution as they gain experience, they should be more willing to trust after a positive experience, and less willing to trust if trust was exploited.

There are also non-strategic modeling approaches that are in line with the role of the buyers' histories. First, adaptation theories predict that people tend to choose strategies that performed well in the past (such as Roth and Erev's, 1995, reinforcement learning theory or Selten's, 1988, learning direction theory). Hence, if trust was rewarded it has a higher probability of being chosen again. Second, the own experience effect may reflect a non-strategic (backward-looking) reciprocal motive. Market participants may not be willing to cooperate in a market in which they were exploited. This argument appears to have more bite in the partners market where *not buy* can be straightforwardly interpreted as a reciprocal punishment for unfair behavior against oneself. The fact that the own experience effect also occurs in the reputation market suggests, however, that such kind of unfairness aversion is also relevant among strangers (similar observations have been made by Blount, 1995, and Bolton et al., 2001, among others).

We neither attempt to advance a formal model, nor was our experimental study designed to separate such theoretical ideas.[26] Our data suggests, however, that a satisfactory model − be it an equilibrium model or a model of boundedly rational behavior − need capture the interaction of strategic behavior (as seen, e.g., in conditional trust, last round effect, etc.) with non-selfish or noisy patterns of behavior (such as non-shipping in the first encounter or shipping in the last encounter) that we observe.

4. Summary

Our market participants systematically respond to the strategic incentives created by a reputation mechanism. Buyers condition their trust on their sellers' feedback profile, similar to what has been observed in field studies, creating incentives for the sellers to be trustworthy − at

[26] A theoretical model would go far beyond the scope of this paper. But see Bolton and Ockenfels (2003) for a formal model of the informational dilemma in a related reputation game that was motivated by the present experimental study. Dellarocas (2003) surveys theories of reputation and relates them to online feedback mechanism.

least as long as reputation can be exploited in the future. As a consequence, reputation markets perform better than strangers markets with respect to straightforward measures of efficiency, trust, trustworthiness, and stability.

But at the same time, our experiments demonstrate that the strategic incentives to trade in reputation markets are weaker than in partner relations. Reputation markets drive a wedge between buyer and seller histories. Thus, in our experiment, both trust and trustworthiness create public benefits not internalized by a reputation mechanism.

Trustworthiness creates public benefits because buyers condition their behavior on their own experience. There is good reason to believe that one's own experience is also important to the decisions field traders make. Online reputation mechanisms typically create various incentives to manipulate feedbacks making them far less reliable than the feedbacks truthfully generated in our experiment.[27] Ones own history is informative of the severity of these problems. In this sense, our findings, if anything, underestimate the influence of buyer histories.

An important implication is that electronic reputation mechanisms may be able to multiply these public benefits by informing all market participants about the shipping probability in the whole market, and not only about the trustworthiness of individual traders. If buyers also take the experiences of other buyers into account (a hypothesis that can be tested within our framework), then a market with high levels of trustworthiness may be able to push trust to higher levels, while of course a market with already low shipping probabilities would deter even more buyers from trusting.

The same extenuating circumstances mentioned with regard to personal experience suggest that the influence of the information dilemma may be even larger in the field than in our experiment. Recall that the informational dilemma occurs when buyers are reluctant to bear the

[27] See e.g., Dellarocas (2001). Ebay distinguishes 4 forms of fraudulent feedbacks: defensive and offensive shill feedback (using secondary Ebay User IDs or other Ebay members to artificially raise the level of your own feedback or to leave negative comments for another user), feedback extortion (demanding any action of a fellow user that he or she is not required to do, at the threat of leaving negative feedback), and feedback solicitation (offering to sell feedback, trade feedback undeservedly, or buy feedback); see http://pages.ebay.com/help/community/investigates.html.

costs associated with verifying trustworthiness. In many online market platforms participants can change their identity at no costs. This creates a stronger incentive for newbies not to ship, because buyers typically cannot distinguish a 'real' newby, who trades the first time, from a 'deceptive' newby, an experienced seller who got rid from his bad reputation. So, all other things equal, the probability that a newby is not trustworthy is likely to be higher on such platforms, which aggravates the informational dilemma.[28] The important implication is that market platforms should try to gain control over the agents' identities (see Friedman and Resnick, 2001, for more theoretical reasons along these lines and how control can be realized).[29]

As we alluded in the Introduction, Granovetter (1985) argues that people have greater confidence in information when it comes from "a trusted informant" that has dealt with the individual in question than when the same information comes from a stranger or an institution; and that they have even greater confidence in information gained from first hand dealings. In particular, Granovetter argues that first hand information (a) is usually cheaper to gather, (b) is often more detailed, (c) offers the promise of future business that provides a great motivation to be trustworthy, and (d) often become overlaid with social content. As we explained earlier, none of these factors can rationalize our experimental data. But our finding on the differential value of the same information across repeated and one-shot relationships would seem to add a new factor to the list.

---

[28] Of course, the conflicting impacts of endogenous pricing and free identity changes can be examined with the help of new experiments within our framework.

[29] Other, more immediate design implications from our study come from the analysis of how buyers respond to feedback patterns. In particular, buyers in our experiment put more weight on negative than positive feedback, and more on recent than old feedback. The emphasis on recent negative feedback has also been reported in field studies (Lucking-Reiley et al., 1999, and Resnick and Zeckhauser, 2001, among others). Given this trader predilection, a cumulative measure may not be appropriate because it hides information critical to the buyers' decision to trust. Another problem with cumulative feedback measures comes from the fact that the seller's actions have diminishing impact in influencing the buyers' assessment of the trustworthiness. This typically leads to increasing incentives to exploit one's good reputation (see Holmstrom, 1999, for a model along these lines in a different context).

References

Andreoni, J., and J.H. Miller (1993). “Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma: Experimental Evidence.” *Economic Journal*, 103, 570-85.

Ariely, D., A. Ockenfels, and A.E. Roth (2002). “An Experimental Analysis of Ending Rules in Internet Auctions.” Working Paper, Harvard Business School.

Avery, C., P. Resnick, and R. Zeckhauser (1999). “The Market for Evaluations.” *American Economic Review*, 89, 564-584.

Blount, S. (1995). “When Social Outcomes Aren’t Fair: The Effect of Causal Attributions on Preferences.” *Organizational Behavior and Human Decision Processes*, 63, 131-44.

Bolton, G.E., A. E. Katok, and A. Ockenfels (2002). “Cooperation among Strangers with Limited Information about Reputation.” Working Paper, Smeal College of Business Administration, Penn State University.

Bolton, G.E and A. Ockenfels (2000). "ERC: A Theory of Equity, Reciprocity and Competition." *American Economic Review*, 90, 166-193.

Bolton, G.E and A. Ockenfels (2003). "Reputation as a Function of Information and Matching in Sequential Equilibria." Mimeo.

Ba, S. and P.A. Pavlou (2002). “Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior.” *MIS Quarterly*, 26, 243-268.

Brynjolfsson, E., and M.D. Smith (2000). “Frictionless Commerce? A Comparison of Internet and Conventional Retailers.” *Management Science*, 46, 563-585.

Dellarocas, C. (2001). “The Design of Reliable Trust Management Systems for Online Trading Communities.” Working paper, MIT Sloan School of Management.

Dellarocas, C. (2003). “The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms.” Working paper, MIT Sloan School of Management.

Diamond, D. (1989). “Reputation Acquisition in Debt Markets.” *Journal of Political Economy*, 97, 828-862.

Fehr, E., and Schmidt, K. (1999). “A Theory of Fairness, Competition, and Cooperation.” *Quarterly Journal of Economics*, 114, 817-868.

Friedman, E., and P. Resnick (2001). "The Social Cost of Cheap Pseudonyms." *Journal of Economics and Management Strategy*, 10(2), 173-199.

Granovetter, Mark (1985), "Economic Action and Social Structure: The Problem of Embeddedness," *American Journal of Sociology*, 91, p. 481-510.

Greif, A. (1989), “Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders,” *Journal of Economic History*, 49, 857-882.

Güth, W. and A. Ockenfels (forthcoming). “The Coevolution of Trust and Institutions in Anonymous and Non-anonymous Communities.” in: M.J. Holler, H. Kliemt, D. Schmidtchen and M. Streit (eds.), *Jahrbuch für Neue Politische Ökonomie*, 20, Tübingen: Mohr Siebeck.

Holmstrom, B. (1999). “Managerial Incentive Problems: A Dynamic Perspective.” *Review of Economic Studies*, 66, 169-182.

Houser, D. and J. Wooders (2001). “Reputation in Auctions: Theory and Evidence from eBay.” Working Paper, University of Arizona.

Keser, C. (2002). “Trust and Reputation Building in E-commerce.” Working paper, IBM, New York.

Kreps, D.M.; Milgrom, P.; Roberts, J. and Wilson, R. (1982). “Rational Cooperation in the Finitely Repeated Prisoners' Dilemma.” *Journal of Economic Theory*, *27*, 245-252.

Ledyard, J. (1995). “Public Goods: A Survey of Experimental Research.” In: J.H. Kagel and A.E. Roth (eds.), *Handbook of Experimental Economics*, Princeton: Princeton University Press, 111-94.

Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves (1999). “Pennies from eBay: the Determinants of Price in Online Auctions.” Working paper, Vanderbilt University.

McMillan, John (2002), *Reinventing the Bazaar: A Natural History of Markets*, New York: Norton.

Melnik, M.I., and J. Alm (2001). “Does a Seller’s e-commerce Reputation Matter?” Working paper, Georgia State University.

Milgrom, P., D. North, and B. Weingast (1990). “The Role of Institutions in the Revival of Trade: The Law Merchant, Private Judges, and the Champagne Fairs.” *Economics and Politics*, 2, 1-23.

Ockenfels (forthcoming). “New Institutional Structures on the Internet: The Economic Design of Online Auctions.” In: M.J. Holler, H. Kliemt, D. Schmidtchen and M. Streit (eds.), *Jahrbuch für Neue Politische Ökonomie*, 20, Tübingen: Mohr Siebeck.

Ockenfels, A. (2003), “Reputationsmechanismen auf Internet-Marktplattformen: Theorie und Empirie.“ *Zeitschrift für Betriebswirtschaft*.

Resnick, P. and Zeckhauser, R. (2001). “Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System.” Draft prepared for NBER workshop.

Resnick, P., R. Zeckhauser, E. Friedman, and K. Kuwabara (2001). “Reputation Systems: Facilitating Trust in Internet Interactions.” Working Paper, mimeo.

Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood (2002). "The Value of Reputation on Ebay: A Controlled Experiment." Working Paper, mimeo.

Roth, A. E. (2002). "The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics." Fisher Schultz lecture, *Econometrica*, 70, 1341-1378.

Roth, A.E., and I. Erev (1995). "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior*, 8, 164-212.

Roth, A.E., and A. Ockenfels (2002). "Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet." *American Economic Review*, 92(4), 1093-1103.

Selten, R. (1998). "Features of Experimentally Observed Bounded Rationality." Presidential Address, *European Economic Review*, 42, 413-36.

Selten, Reinhard, und Rolf Stöcker (1986). "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames." *Journal of Economic Behavior and Organization*, 7, 47-70.

U.S. White Collar Crime Center and the Federal Bureau of Investigation, "IFCC 2001 Fraud Report," http://www1.ifccfbi.gov/strategy/statistics.asp.

Appendix A. Subject instructions and buyer screen

BELOW ARE THE WRITTEN INSTRUCTIONS THAT WERE GIVEN TO SUBJECTS IN THE REPUTATION TREATMENT. INSTRUCTIONS FOR THE OTHER TREATMENTS WERE PARALLEL, THE ONLY DIFFERENCES BEING THE DESCRIPTION OF THE REPUTATION SYSTEM (REMOVED FOR STRANGERS), OR THE DESCRIPTION OF PARTNER ROTATION (PARTNERS).

General. The purpose of this session is to study how people make decisions. If at any time you have questions, raise your hand and a monitor will happily assist you. From now until the end of the session, unauthorized communication of any nature with other participants is prohibited.

During the session you will play a game that gives you an opportunity to earn cash. At the end of the session, you will be paid your earnings plus a $5 show-up fee. Decisions and payments are confidential: No one will be told your actions or the amount of money you make.

[The figure that appeared here is the same as Figure 1 in the text of the paper.]

Description of the game. You and the other participants in the room (but not the monitors) are the players in the game. The game proceeds in a series of rounds. Each round, each player is randomly matched with another player to trade a (fictional) commodity. First, one of the players, the "Buyer," chooses to either *buy* or *not buy*. If the Buyer chooses *not buy*, then the game ends and both players receive $.35. If the Buyer chooses to *buy*, then the game continues and the other player, the "Seller," makes a decision to *ship* or *not ship*. *Ship* pays each player $0.50 while *not ship* pays the Buyer nothing and the Seller $0.70.

The game will last for 30 rounds. You will be a Buyer for half of the rounds, and a Seller for the other half. When you switch between roles is a matter of random chance, so you may be in one role for more than one round in a row before switching to the other role, and the pattern of switching may be different for you than for other players in the game.

Seller's feedback history. For each game played, the computer will record whether the Seller chose *ship* or *not ship* (if the Seller did not get to move, the computer records nothing). This feedback will then be made available to all future Buyers that are matched with this Seller. The feedback will include a summary of the number of times the Seller shipped in the past, as well as a round-by-round history of their shipping decisions, beginning with the most recent decision. Buyers will see this feedback history prior to making their buy decision.

Pairings. All partner pairings are anonymous: Your identity will not be revealed to the person you are playing with either before, during or after the game. You will never be matched with the same player in the same role more than once.

Money earnings. You will be paid your earnings from all of the rounds of the game (plus a $5 show-up fee) in cash.

Practice games. When the monitor gives the OK, play some practice games. Your partner for the practice games will be the computer. It has been programmed to choose its moves at random. The practice games will allow you to experience the game from both the Buyer and Seller's perspective. Practice until you feel comfortable with the game and its rules.

Consent Forms. If you wish to participate in this study, please read and sign the accompanying consent form. The consent form explains your rights as a subject as well as the rules of confidentiality that will be adhered to regarding your participation.

Figure A1. Buyer screen

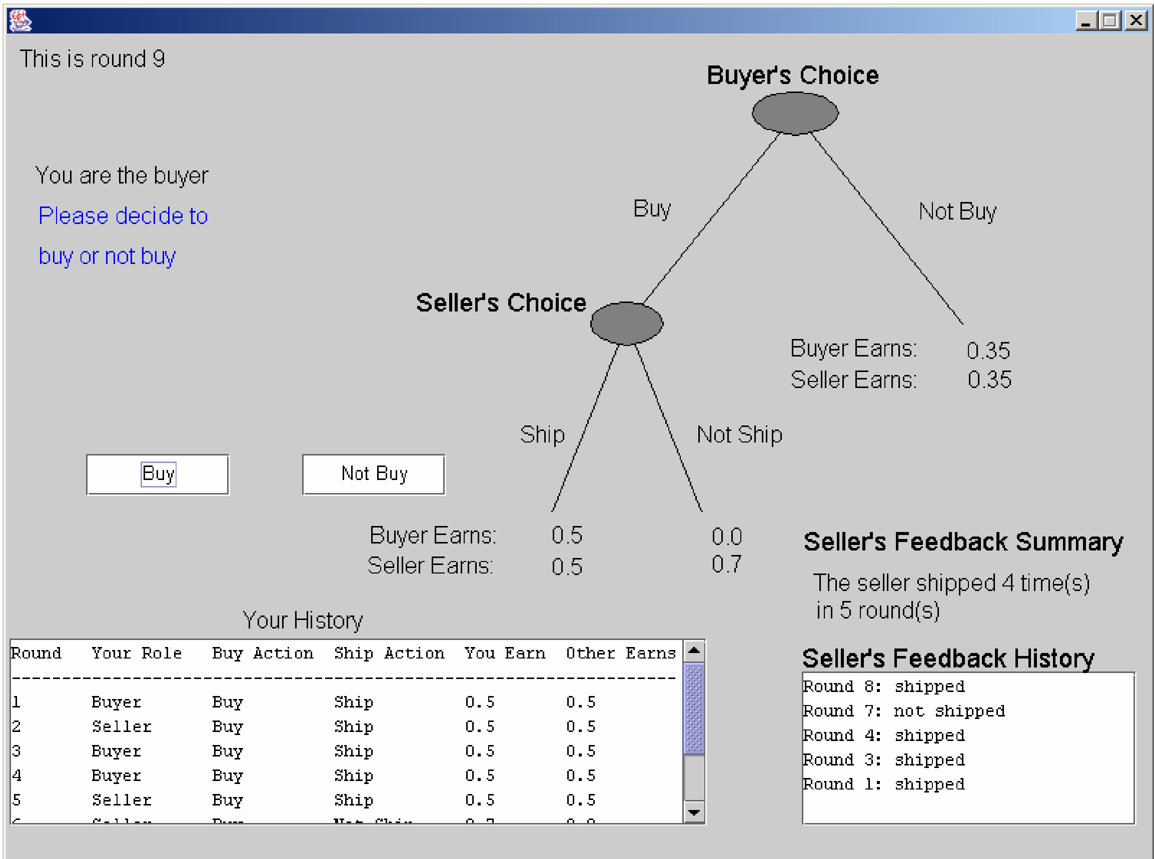

Appendix B.

Table A2. Fixed effects linear models, buyers[a]

OLS estimates (and two-sided *p*-values) for buyer behavior

Dependent variable = "1" for *buy*

Independent variable	Model 1	Model 2	Model 3
CONSTANT	---	---	---
REPUTATION	---	---	---
= 1 if buyer is from reputation treatment, and 0 else.			
PARTNERS	---	---	---
= 1 if buyer is from partners treatment, and 0 else.			
TOTALSHIPSreputation			0.021
= number of seller ships prior to last order.			(.0000)
TOTALNOSHIPreputation			-0.035
= number of seller no ships prior to last order.			(.0092)
SHIPLASTreputation			0.009
= 1 if reputation seller shipped last order, and 0 else.			(.8697)
NSHIPLASTreputation			-0.261
= 1 if reputation seller did not ship last order, and 0 else.			(.0000)
SHIPLASTpartners			0.167
= 1 if seller in partners shipped last order, and 0 else.			(.0023)
NSHIPLASTpartners			-0.224
= 1 if seller in partners did not ship last order, and 0 else.			(.0005)
CBSH		0.000	-0.006[b]
= number of past times item was shipped to buyer.		(.9103)	(.1428)
CBNH		-0.133	-0.129[b]
= number of past times buyer bought but not shipped.		(.0000)	(.0000)
ROUNDstrangers	-0.184		
= round in strangers treatment, and 0 else.	(.0000)		
ROUNDreputation	-0.007		
= round in reputation treatment, and 0 else.	(.0003)		
ROUNDpartners	0.001		
= round in partners treatment, and 0 else.	(.6651)		
LAST2ROUNDstrangers	-0.005	-0.069	-0.070
= 1 if round 29 or 30 in strangers treatment, and 0 else.	(.9434)	(.2393)	(.2227)
LAST2ROUNDreputation	-0.301	-0.263	-0.263
= 1 if round 29 or 30 in reputation treatment, and 0 else.	(.0000)	(.0000)	(.0000)
LAST2ROUNDpartners	-0.213	-0.136	-0.111
= 1 if round is 29 or 30 in partners treatment, and 0 else.	(.0010)	(.0225)	(.0610)
Number of observations	2160	2160	2160
Adjusted R-squared	0.368	0.401	0.442
F-test *p*-value	.0000	.0000	.0000

[a]These are analogous estimates for Table 2 in the text.

[b] History for Partner's buyers does not include last transaction.

Table A2. Random effects probit models, buyers[a]

Maximum likelihood estimates and two-sided *p*-values for buyer behavior

Dependent variable = "1" for *buy*

Independent variable	Coefficient	P-value
CONSTANT	0.512	0.0008
PARTNERS	0.835	0.0016
SHIPTOTALreputation	0.077	0.0009
NOSHIPTOTALreputation	-0.089	0.1194
LASTSHIPreputation	0.270	0.1520
LASTNOSHIPreputation	-0.566	0.0073
LASTSHIPpartners	1.184	0.0007
LASTNOSHIPpartners	-0.776	0.0127
CBSHstrangers	0.003	0.9691
CBSHreputation	-0.067	0.1171
CBSHpartners	0.055	0.2151
CBNHstrangers	-0.382	0.0000
CBNHreputation	-0.334	0.0000
CBNHpartners	-0.474	0.0000
LAST2ROUNDSstrangers	-0.420	0.1896
LAST2ROUNDSreputation	-0.926	0.0001
LAST2ROUNDSpartners	-1.649	0.0000
RHO	0.442	0.0000
Number of observations	2160	
Log-likelihood	-985.78	
F-test *p*-value	.0000	