

TRUST IN A NETWORKED WORLD:
EXPERIMENTAL GAMES FOR THE DESIGN OF
REPUTATION MANAGEMENT SYSTEMS

Claudia Keser

IBM T. J. Watson Research Center and CIRANO, Montreal

March 24, 2003

Abstract: The extent to which people trust each other influences the performance of organizations to which they belong and the growth rates of countries in which they reside. Trust also positively impacts Internet adoption. We argue that trust levels can be increased through the use of trust enhancing mechanisms. Reputation management systems, such as ebay's Feedback Forum, are examples of such mechanisms in the context of e-business. Reputation management systems used by various online markets differ substantially but there has been little research done on the design of optimal mechanisms. Experimental economics provides a framework to measure trust and trust enhancement in a controlled laboratory environment. We present an experimental laboratory study examining the functionality of two ebay-like reputation management systems and the factors that influence their effectiveness.

1) Trust and the Internet

Trust is at the root of almost any economic or personal interaction. We need to trust in our government, the credit card company, the car dealer, or our business partner. Francis Fukuyama (1995) argues that trust, loosely defined as the expectation of trustworthy and cooperative behavior of others, impacts the performance of all social institutions, including firms, and thus the overall economic performance of a country [1]. Individuals in higher-trust societies spend less to protect themselves from being exploited in economic transactions. Trust is an economical substitute for extensive contracts, litigation, and monitoring in transactions and thus economizes on transaction costs. Knack and Keefer (1997) provide empirical foundation for Fukuyama's hypothesis, showing that differences in trust across countries help explain differences in investment and economic growth [2]. They measure trust in a country based on an indicator from the World Values Surveys conducted in 1981 and 1990-1991. The indicator is the percentage of respondents from that country who answered positively when asked:

“Generally speaking, would you say that most people can be trusted, or that you can't be too careful in dealing with people.”

Based on the same trust indicator, Huang et al. (2002) show that more trusting countries tend to exhibit higher levels of Internet penetration [3]. Litan and Rivlin (2001) and Varian et al. (2002) predict that, due to reduced transaction costs associated with production and distribution of goods and services, the Internet will positively affect

productivity in the coming years [4,5]. To the extent that productivity gains increase with the level of internet adoption, low trust countries, the majority of which tend to be of low and middle income, will take a double hit in terms of economic growth, penalized for low trust first in terms of direct impact on investment and growth, and then again through lower adoption of a growth-enhancing technology. In other words, we will observe a developmental-cum-digital-divide (see Figure 1).

This is a bad message for the low-trust countries. However, there might be ways to mitigate the implications. In specific situations, trust can be enhanced or complemented by related factors. In Section 2 we identify some of these factors and illustrate a case of successful trust enhancement in the context of Internet transactions by reputation management—ebay’s Feedback Forum.¹ However successful, ebay’s Feedback Forum shows a number of shortcomings, which leads us to the question how ebay’s reputation management system could be improved upon. In section 3 we present a new methodology appropriate to answer this type of question in the experimental economics laboratory. Section 4 concludes this article.

2) Trust enhancement and the ebay example

Trust is typically assigned the role of solving problems caused by social uncertainty when we are incapable of correctly determining the intentions of other persons or organizations having (monetary) incentives to act against our own best interest. Following Yamagishi and Yamagishi (1994) we define trust as the expectation of other persons’ goodwill and benign intent, implying that in certain situations those persons will place the interests of others before their own [7].

Yamagishi and Yamagishi (1994) distinguish trust from assurance, which they define as the expectation of another person's goodwill and benign intent based on the knowledge of the incentive structure surrounding the relationship. They give the following example:

Suppose I have a special tie with the Mafia, and my trading partner knows this. I am certain that he will not cheat on me; he knows that if he does he will be quickly sent to a mortuary. My expectation of the partner's "honesty" is based on the fact that acting "honestly" is in his own interest, not on the belief that he is a benevolent person. Here, assurance exists but no trust. (Yamagishi and Yamagishi 1994, p. 132)

Assurance can thus complement or substitute for trust.

A related trust substitute is commitment. Maintaining long-term relationships with loyal partners rather than making deals with new partners is a kind of commitment where incentives for non-cooperative behavior are reduced.² In such relationships it is mutual assurance based on the nature of the relationship rather than trust that leads to cooperative behavior. This was demonstrated, for example, by Axelrod (1984), Selten, Mitzkewitz and Uhlich (1997) and Keser (2000) who examined human strategies in repeated "social dilemma situations" where the individually payoff maximizing non-cooperative behavior leads to socially inefficient outcomes [8,9,10]:³ people often actively attempt to establish and maintain mutual cooperation when they expect to repeatedly interact with each other.

In early interactions they signal their willingness to cooperate and then use reciprocity—cooperate if the other(s) cooperated and defect from cooperation if the other(s) defected in the previous interaction—as an instrument to establish cooperation. Following such a strategy typically pays for an individual that is involved in many repeated encounters with others. Keser and van Winden (2000) show that in a social dilemma situation people who interact with the same others (“partners”) cooperate significantly more than those who interact in each repetition with others who are randomly associated to them (“strangers”) [11].

Also familiarity can be seen as a complement to trust [12]. Familiarity deals with the understanding of the current action of another person while trust deals with the beliefs about future actions of another person. The latter may often be based on familiarity. Familiarity with Amazon.com, for example, would be the knowledge of how to search for books and information about them and how to order these books through the website interface. Trust in Amazon.com might entail willingness to provide credit card information based on the belief that the information will not be inappropriately used in the future. Though trust and familiarity are distinctly different, they are related.

Another trust complement or trust enhancing factor is reputation. Reputation may play two different roles in social interactions involving trust. The first role is informational. It makes the recipient of positive reputation information trust more. Trust has been defined above as the expectation that others will show goodwill in their dealings with us. Lacking perfect information about others’ intentions, we have to infer their intentions from other available information, like their reputation. In other words, reputation may enhance trust by the creation of assurance. The second role of reputation

is a kind of sanctioning. The attribution of a negative reputation may work as a sanctioning mechanism to punish dishonest behavior. This likely makes the owner of reputation act in a more trustworthy way.

In e-business we observe successful trust enhancement by reputation management systems. The most popular of those is ebay's Feedback Forum. At ebay, anonymous individuals spread over the globe may buy and sell almost everything from Pez dispensers to Ferraris or castles. With nearly 50 million registered users and 170 million transactions (for \$ 9.3 billion worth of goods) in 2001, it is the largest of the informal online markets [13]. These numbers are impressive given the risks involved in trading on such a market. Typically there is no opportunity for a buyer to inspect the item for which he has paid before delivery, and if he isn't satisfied with the quality, it may be impossible for him to track down the seller. Even worse, the buyer has no guarantee that the item will be delivered at all. The seller, on the other hand, if he chooses to deliver before receiving the payment faces similar risks with respect to the buyer. To put it differently, each of the parties involved in a trade might be tempted to cheat. Ebay has a fraud protection program that covers losses for up to \$200. However, beyond that, if users do not want to make use of costly escrow services that ebay also offers, they must bring a large portion of trust when they engage in transactions on this informal online market.

To enhance trust and trustworthiness of its users ebay created its Feedback Forum. In addition to textual comments, the participants in a transaction are asked to rate each other with a "+1" for a positive comment, a "-1" for a negative comment or a "0" for a neutral comment. All ratings that an ebay user has received from distinct other ebay users are summed up to build his or her Feedback Rating number. This number is

attached to each seller's or bidder's identity. A user who has 100 positive comments is characterized by a Feedback Rating number of 100. However, another user who has 150 positive and 50 negative comments would also be characterized by 100. Any user whose Feedback rating number reaches -4 is suspended from participation. The Feedback Rating number is part of the user's Feedback Profile, which can be obtained by clicking on the user's Feedback Rating number. A user's Feedback Profile provides the full list of textual comments on that user, the distribution of all previous ratings that he or she received from distinct other users, and the distribution of his or her recently received ratings.⁴

Recently, a number of empirical studies addressed the question of whether the prices that sellers obtain on ebay are correlated with their reputation. Kalyanam and McIntyre (2001) examined auctions of Palm Pilot personal digital assistants, Houser and Wooders (2000) examined auctions of Pentium III processors, Lucking-Reiley et al. (2000) examined collectible coin auctions [14,15,16]. All of these studies come to the conclusion that buyers are willing to pay more for a good coming from a seller with a good reputation. Resnick et al. (2001) conducted a field experiment in which they (1) sold matched pairs of items (batches of vintage postcards) either under the extremely high reputation of an established ebay auctioneer or under newcomer identities with little reputation, and (2) compared sales under newcomer identities with and without negative feedback [17]. They observed (1) that the established identity fared better than the newcomer identity, and (2) that among the newcomers, one or two negative feedbacks showed no price effects. These empirical and experimental field studies thus show that the reputation heterogeneity created by ebay's reputation management system leads to

price dispersion, which was assumed to be minimized on online markets due to reduced buyer search cost [18]. We are aware of no field study, however, that examines the function of specific aspects of reputation management systems and, in particular, their impact on trust and trustworthiness.

Although the reported fraud rate at ebay is as low as one percent of all listings, there are recurring incidents of fraud that may be due to shortcomings in the reputation management system. Dingleline, Freedman and Molnar (2001) for example observed the following non-unique incident [19]:

“In the eBay case, a group of people engaged in auctions and behaved well. As a result, their trust ratings went up. Once their trust ratings were sufficiently high to engage in high-value deals, the group suddenly ‘turned evil and cashed out.’ That is, they used their reputations to start auctions for high-priced items, received payment for those items, and then disappeared, leaving dozens of eBay users holding the bag.”

Such incidents indicate that there are drawbacks in the functioning of the current Feedback Forum. The Feedback Rating number sums up a user’s ratings received both in the role of a seller and a buyer. It appears, however, easier for a buyer than for a seller to collect positive evaluations. Thus, it might be worthwhile for ebay to consider segmenting the Feedback Rating number. The segmentation could not only be according to the role in the transaction (buyer or seller) but also by the product category. A seller of

PEZ dispensers may enjoy an excellent reputation among PEZ dispenser buyers. He might, however, be lousy at selling a car on ebay. This could be unintentional due, for example, to lack of experience or it could be a matter of intent: a car represents a much higher value than a PEZ dispenser. The strategy of building reputation on low-value deals only to misuse the reputation on high-value deals could also be discouraged by modifications in the reputation function: for example, high-value deals could be weighed more heavily than low-value deals in the construction of the Feedback Rating number. Fraud might also be discouraged by making it more difficult for traders to change their identity. A user with a negative reputation has, in the current version of the Feedback Forum, an incentive to reappear on the market with a new identity that comes along with a neutral reputation. Without going into further detail, we conclude that many potentially beneficial modifications could be made to ebay's Feedback Forum. One cannot always foresee the implications such modifications have on the behavior of users and the general performance of the reputation management system. Experimenting in the field can be costly and time consuming. Furthermore, it is unlikely that we can control the environment to the extent that we would like to. We thus propose to examine the effectiveness of different rating systems in controlled laboratory experiments. In the following we present a specific way to measure trust, trustworthiness, and the enhancement of both through reputation management in the experimental economics laboratory.

3) Designing reputation mechanisms^{3/4}fast prototyping using experimental game theory

In a model e-business environment, we want to consider the interaction of individuals in an abstract situation involving trust on one side and trustworthiness on the other. We do not necessarily want a one-to-one translation of the ebay auction situation, as the auction design involves other issues that, at this point, we want to keep isolated from the trust question.⁵ In the experimental economics literature, Berg, Dickhaut, and McCabe (1995) introduced the following “investment game”, often also called the “trust game” [22].

The two players in this game—let us denote them as buyer and seller—are endowed with \$10 each. Buyer and seller may interact according to the following rules, of which they both have full knowledge. The buyer may send (invest) part or all of his endowment to the seller, but he need not send anything. The amount invested by the buyer will be tripled, so that the seller will receive three times the amount invested by the buyer.⁶ The seller then has the opportunity to return part or all of the amount he received, but he need not return anything. Then the game is over.

Suppose that the two players in this game are, as typically assumed in economic theory, egoists who are striving to maximize their personal payoffs. Suppose further that each assumes the other’s objective is to maximize his or her own personal payoff. The game can then easily be solved by backward induction: a seller who egoistically maximizes his personal payoff will not return anything to the buyer, whatever amount the latter invested. The buyer, also striving to maximize his personal payoff and anticipating

the zero return by the seller, will not send anything to the seller. Thus, economic theory predicts zero flow of money in this game.

In the experimental laboratory, as in real life, we often observe people behaving less egoistically than postulated by economic theory. Berg, Dickhaut and McCabe (1995) conducted laboratory experiments on the trust game under strong anonymity conditions, guaranteeing that none of the other participants and not even the experiment monitors could identify their personal decisions. They gathered a number of participants, half of them in the role of buyers and the other half in the role of seller. Buyers and sellers were in different rooms and never saw each other. To conduct the experiment, an envelope-mailbox system was used. Each buyer had the opportunity to leave any amount between zero and ten dollars in an envelope, which he then deposited in a mailbox that he shared with an unknown seller. The experiment monitor would triple the amount in each of the envelopes before the sellers accessed their mailboxes to pick up the envelopes. Each seller then had the opportunity to return anything between zero and the received amount in the envelope which he had to put back into the mailbox to be picked up later again by the buyer with whom he shared the mailbox. Participants in these experiments were undergraduate students from the University of Minnesota who played this game for real money. More than 90 percent of those in the role of a buyer sent positive amounts of money to the seller (\$5.16 on average) and about 45 percent of those in the role of a seller who received some amount from the buyer returned positive amounts (\$4.66 on average). The amount invested by the seller provides a measure for his trust in the seller's goodwill while the relative return by the seller provides a measure for trustworthiness. Thus, the experiments provide evidence that people, to some extent, do trust in the goodwill of

others. Furthermore, they are not entirely untrustworthy: the participants in the role of a seller returned, in aggregate, almost as much as the participants in the role of a buyer invested. This means that in this game trust does not really pay for a buyer. The sellers tend to keep the entire surplus created by the buyers' investments for themselves, which implies an inequitable payoff distribution among buyers and sellers. Note also that the observed level of trust is still far below the socially efficient level: the sum of payoffs to both players would be maximal if the buyer invested the entire endowment. In other words, only full trust could lead to a socially efficient outcome. In the case of full trust, any return by the seller greater or equal to ten implies a Pareto efficient situation, in which none of the players could be made better off without reducing the payoff to the other player.

The experimental results by Berg, Dickhaut and McCabe have been replicated in a large number of experimental studies conducted in various countries. The results are qualitatively robust although the actual level of trust varies across countries. Interestingly, the observed differences are directionally in keeping with the trust measure used by Knack and Keefer (1997) and Huang et al. (2001) in the empirical studies discussed in the introduction. Willinger et al. (forthcoming), for example, find that Germans do trust significantly more than the French [23]. Buchan, Croson and Dawes (2002) observe a higher trust level among Americans and Chinese than among Koreans and Japanese [24]. Ensminger (2002) reports on a study where norms of altruism, trust, and cooperation of people have been studied in places like New Guinea, the Amazon Rain Forest, Kenya, and the urban and rural Missouri [25]:

“...[t]he results [of the trust game] indicate that in the small-scale societies of the developing world, there is less trust, and in the United States there is more (in both rural and urban Missouri).” (Ensminger 2002, p.13)

The situation described by the trust game relates, the auctioning aspect being abstracted away, to trading on online markets. The focus is on the issues of trust and trustworthiness. If I buy PEZ dispensers at ebay I face risks similar to those of the buyer in the trust game: is the quality as described, will the packaging be appropriate, will the delivery be in a timely fashion, will the item be delivered at all? The seller on ebay may be considered in the role of the seller in the trust game in as far as his or her trustworthiness is at stake. Although the trust game, due to the tripling of the buyer's investment, does not directly translate into the ebay transaction, it captures in a simple way its major issues of trust and trustworthiness. The tripling of the investment corresponds to the existence of gains from trade in the buyer-seller context. As discussed in Section 2 above, it is to be assumed that at ebay the buyer's risks with respect to the seller's trustworthiness are somewhat mitigated owing to the Feedback Forum.⁷ On the one hand, the seller's reputation provides a signal that is likely to influence the buyer's expectation of trustworthiness; on the other hand, the fear of being negatively rated might make sellers behave in a more trustworthy way.

In Keser (2002), we have designed, based on the trust game, computerized laboratory experiments that examine the functionality of such reputation management systems [26]. The big advantage of using the laboratory experimental method, compared to field studies or experiments, is that we can directly compare a situation without a

reputation management system to situations with specific reputation management systems. We can thus measure to what extent each reputation management system impacts trust and trustworthiness. We have conducted three different experiments. All of the experiments are based on a twenty-fold repetition (periods) of the trust game. The participants remain in the same role of either a buyer or a seller over all twenty periods. The pairing of a buyer with a seller is random in each period, but it is guaranteed that the same two players never meet in two consecutive periods. The first experiment is called the *baseline experiment*; it is based on the trust game such as described above. The other two experiments are based on a modified trust game with a reputation management system in which the buyer, after having been informed of how much the seller returned to him, is asked to evaluate the seller's cooperation as *positive*, *neutral* or *negative*. The seller is informed about the evaluation that the buyer gave him. From the second period on, in the beginning of each period, each buyer is informed of the previous rating(s) attributed to the seller with whom he is going to interact. In the second experiment, called the *short-run reputation experiment*, the buyer is informed of the seller's most recent rating. In the third experiment, called the *long-run reputation experiment*, the buyer is informed of the most recent rating attributed to his seller as well as of the distribution of all previous positive, neutral and negative ratings to that seller. This information is similar to the information given in a user's Feedback Profile on ebay.

The method: 320 students from different disciplines at various universities in Montreal voluntarily participated in the experiments. Thirty-two experimental sessions were organized, eight for the baseline experiment, and twelve of each reputation

experiment. In each session, ten students participated—five in the role of a buyer and five in the role of a seller. Buyer-seller pairs were randomly created in each of the 20 periods, under the constraint that no two participants should meet in two consecutive periods. Note that the random matching implies that the aggregate behavior in a session represents an independent observation on which non-parametric statistics (SPSS 10.0 for windows) will be based. The rules of the game were identical in each of the 20 periods. The payoffs of a participant over the 20 periods were added up to determine his or her individual earnings in the experiment. An experimental session lasted about 1.5 hours and the average earnings were \$ CN 30 (including a show-up fee of \$ CN 5). The experiments were conducted in French in the computerized experimental economics laboratory at CIRANO/LUB (Centre Interuniversitaire de Recherche en ANalyse des Organizations/Laboratoire Universitaire Bell) in Montreal. Instructions were distributed and then read aloud. Before the experiment could start, all participants had to correctly answer a number of questions testing their understanding of the instructions. Payment was in private at the end of the experiment.

The results of the baseline experiment are in keeping with previous experimental results. In our baseline experiment the buyers' average investment (the *measure of trust*) is 3.91 while the sellers' average return is 3.3, which amounts to a relative return with respect to the received amount (the *measure of trustworthiness*) of 32 percent. In other words, in aggregate, the buyers receive a little less than the invested amount back, which implies that the sellers keep not only the entire surplus but also a small part of the buyers' investment for themselves. While trustworthiness is relatively stable over time, we

observe a statistically significant decrease in the trust level from the first ten periods to the last ten periods of the game (two-sided Wilcoxon signed ranks test, 5 percent level).

The experiments show significant effects of the introduction of the reputation management system on trust and trustworthiness. When compared to the baseline, the short-run reputation management system increases trust and trustworthiness by more than 30 percent (two-sided Mann-Whitney U-tests, 10 percent and 5 percent level, respectively): it increases trust to 5.15 and trustworthiness to 46 percent. The long-run reputation management system is even more efficient increasing both trust and trustworthiness by more than 50 percent (two-sided Mann-Whitney U-tests, 1 percent level): it increases trust to 6.05 and trustworthiness to 49 percent.

The development of trust over time in the three experiments is shown in Figure 2, while Figure 3 shows how trustworthiness develops. These figures show that, in aggregate, with a reputation management system in place both trust and trustworthiness are always higher than in the baseline experiment except for the final period. The introduction of a reputation management system increases the levels of trust and trustworthiness from the first period but, toward the end of the interaction, there is a dramatic drop in trust and trustworthiness. This is a typical end game effect, as observed in the vast experimental literature on finitely repeated prisoner's dilemma type of games where cooperation tends to break down toward the end [27].

In terms of payoff, the major winners from the introduction of a reputation management system are the buyers, which is due to the fact that, with a reputation management system in place, sellers are forced to return more than what was invested by the buyers in order to get a positive rating. The buyers' average per period payoff

increases from 9.80 in the baseline to 11.95 in the short-run and to 12.83 in the long-run reputation experiment. Both of these increases are significant at the 1 percent level (two-sided Mann-Whitney U-tests).

With a reputation management system in place also sellers earn higher payoffs than in the baseline experiment: the sellers' average per period payoff increases from 17.93 in the baseline to 18.30 in the short-run and to 19.27 in the long-run reputation experiment. These increases are, however, statistically not significant, requiring a 10 percent significance level of the two-sided Mann-Whitney U-test.

Sellers care very much for their long-run reputation when it is at stake. They have reason to do so as buyers tend to trust significantly more the better the seller's long-run reputation. Giving a value of +1 to each positive, a value of 0 to each neutral and a value of -1 to each negative rating, we define a seller's long-run reputation as the sum of this seller's previous ratings. Based on this, we observe that buyers' investments tend to be higher in the case of a positive than in the case of a neutral long-run reputation and higher in the case of a neutral than in the case of a negative long-run reputation (two-sided Wilcoxon signed ranks tests, 1 percent level). Short-run reputation also matters, but in the experiment where information on the long-run reputation is available additional short-run reputation significantly affects the buyers' trust only when it is positive (two-sided Wilcoxon signed ranks test, 1 percent level). In the experiment where short-run reputation is the only information available, a positive reputation significantly increases the buyers' trust while a negative reputation significantly decreases it (two-sided Wilcoxon signed ranks tests, 1 percent level in the comparison of positive and neutral reputation and 5 percent level in the comparison of neutral and negative reputation).

To summarize the results, the introduction of a reputation management system significantly increases both the transaction volume (trust) and the level of trustworthiness. Defining efficiency as the percent of the maximal possible sum of payoffs to buyer and seller, we observe that the overall efficiency increases from 70 percent in the baseline to 76 percent in the short-run and to 80 percent in the long-run reputation experiment owing to the increased trust level. The introduction of a reputation management system leads to a Pareto improvement as both player types earn higher payoffs. At the same time, the payoff distribution between buyers and sellers becomes more equitable. This is due to the fact that in an environment with a reputation management system in place engaging in a transaction that requires trust will, in aggregate, pay, which is not the case in an environment without any reputation management system. Comparing the two reputation experiments, we observe that the management of long-run reputation leads, in the intermediate phase of the interaction, to more trust and trustworthiness than the management of purely short-run reputation.

4) Conclusion

We started our discussion with the observation that trust plays an important role at the level of individual transactions, organizations and the economic performance of a country. Differences in trust across countries exist. In specific contexts, there are, however ways for low-trust countries to substitute or enhance trust by other factors. In e-business we observe successful trust enhancement through reputation management systems. However, there is no standard reputation management system (each online market uses its own design) nor even a set of guidelines for designing effective systems

[28]. We suggest using the methodology of experimental game theory for the design of efficient reputation management systems.

The experiments presented in Keser (2002) examine the effectiveness of two reputation management systems, of which one, the long-run reputation management system, is very similar to ebay's Feedback Forum. The results suggest that if ebay had not introduced any reputation management system, they would have experienced significantly lower transaction volume and significantly more fraud. The management of long-run reputation appears preferable to the management of short-run reputation. It remains an open question whether ebay could do even better at enhancing trust and trustworthiness by using a reputation management system that is different from the one they currently use. Should they use, for example, a system, in which each rating that a user receives is weighted based on the value of the transaction? Should a rater's own reputation be taken into account to weight his or her rating? How important is it that ebay users have no doubts about the technical reliability of the feedback forum. How do communities of interest on ebay influence trust and reputation formation. Laboratory experiments provide us with a controllable, fast and relatively inexpensive means of addressing all these questions. As such the laboratory is an important complement to theoretical analyses and field studies.

Acknowledgements

I thank CIRANO/LUB in Montreal for their hospitality and in particular Claude Montmarquette and Jacinthe Majeau for their help in running the experiments. Thanks are also due to Robert Baseman, Bill Grey, Jonathan Leland, Mark Nixon, and Carl Singer for their input and many discussions. I am grateful to anonymous referees who made valuable comments. Prince Stanley developed the experimental software, which is based on z-Tree by Urs Fischbacher at the Institute for Empirical Research in Economics at the University of Zurich.

References

- [1] Fukuyama, F., 1995, *Trust: The Social Virtues and the Creation of Prosperity*, New York: Free Press.
- [2] Knack, S., and P. Keefer, 1997, Does social capital have an economic payoff? a cross-country investigation, *Quarterly Journal of Economics*, 1251-1288.
- [3] Huang, H., C. Keser, J. Leland, and J. Shachat, 2002, Trust, the Internet, and the digital divide, IBM Research Report RC22511.
- [4] Litan R. E., and A. Rivlin, 2001, *Beyond the Dot.Coms: The Economic Promise of the Internet*, Internet Policy Institute, Brookings Institution Press.
- [5] Varian, H., R. E. Litan, A. Elder, and J. Sutter, 2002, *The Net Impact Study*, <http://www.netimpactstudy.com/>
- [6] IBM Global Services, 1999, IBM Multi-National Consumer Privacy Survey, www.ibm.com/services/files/privacy_survey_oct991.pdf.
- [7] Yamagishi, T., and M. Yamagishi, 1994, Trust and commitment in the United States and Japan, *Motivation and Emotion* 18, 129-166.
- [8] Axelrod, R., 1984, *The Evolution of Cooperation*, New York: Basic Books.
- [9] Selten, R., M. Mitzkewitz, and G. Uhlich, 1997, Duopoly strategies programmed by experienced players, *Econometrica* 65, 517-555.
- [10] Keser, C., 2000, Strategically planned behavior in public goods experiments, Working Paper, CIRANO, Scientific Series 2000s-35.

- [11] Keser, C., and F. van Winden, 2000, Conditional cooperation and voluntary contributions to public goods, *Scandinavian Journal of Economics* 102, 23-39.
- [12] Gefen, D., 2000, E-commerce: the role of familiarity and trust, *Omega* 28, 725-737.
- [13] Jerry Adler, The eBay way of life, *Newsweek*, June 17, 2002.
- [14] Kalyanam, K., and S. McIntyre, 2001, Returns to reputation in online auction markets, mimeo, presented to IBM Consulting Group.
- [15] Houser, D., and J. Wooders , 2000, Reputation in auctions: theory and evidence from eBay, Working Paper, University of Arizona.
- [16] Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves, 2000, Pennies fom ebay: the determinants of price in online auctions, Working Paper, University of Arizona.
- [17] Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood, 2002, The value of reputation on eBay: a controlled experiment, Working Paper, University of Michigan.
- [18] Bakos, J. Y., 1997, Reducing buyer search costs: implications for electronic marketplaces, *Management Science* 43, 1676-1692.
- [19] Dingledine, R., M. J. Freedman, and D. Molnar, 2001, Accountability, in: A. Oram (ed.), *Peer-to-Peer: Harnessing the Benefits of a Disruptive Technology*, O'Reilley & Associates.
- [20] Roth, A. E., and A. Ockenfels, 2002, Last-minute bidding and the rules for ending second-price auctions: evidence fom ebay and amazon auctions on the Internet, *American Economic Review* 92, 1093-1103.

- [21] Ockenfels A., and A. E. Roth, 2002, Late and multiple bidding in second price Internet auctions: theory and evidence concerning different rules for ending an auction, Working Paper, Harvard Business School.
- [22] Berg J., J. Dickhaut, and K. McCabe, 1995, Trust, reciprocity, and social history, *Games and Economic Behavior* 10, 122-142.
- [23] Willinger, M., C. Keser, C. Lohmann, and J.-C. Usunier, *forthcoming*, A comparison of trust and reciprocity between France and Germany: experimental investigation based on the investment game, *Journal of Economic Psychology*.
- [24] Buchan, N. R., R. T. A. Croson and R. M. Dawes, 2002, Swift neighbors and persistent strangers: a cross-cultural investigation of trust and reciprocity in social exchange, *American Journal of Sociology*.
- [25] Ensminger, J, 2002, From Nomads in Kenya to small-Town America: experimental economics in the bush, *Engineering and Science* 2, 6-16.
- [26] Keser, C., 2002, Managing brands in e-business: an experimental study in trust and reputation management, IBM Research Report RC22654.
- [27] Selten, R., and R. Stoecker, 1986, End behavior in sequences of finite prisoner's dilemma supergames, *Journal of Economic Behavior and Organization* 7, 47-70.
- [28] Kollock, P., 1999, The production of trust in online markets, in: *Advances in Group Processes* (Vol. 16), edited by E. J. Lawler, M. Macy, S. Thyne, and H. A. Walker, Greenwich, CT: JAI Press

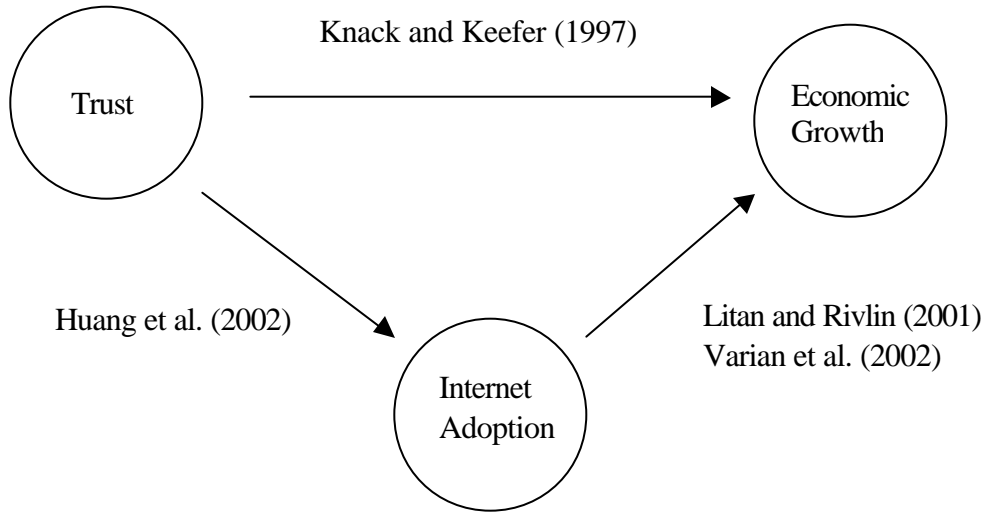


Figure 1: The developmental-cum-digital divide




Figure 2: Buyers' investments (trust measure) over time




Figure 3: Sellers' relative returns (trustworthiness) over time

Endnotes

¹ We have to acknowledge that the relationship between the general level of trust in a country and trust in the specific context of an Internet transaction has not yet been firmly established. Survey studies on privacy, an issue very relevant to online transactions, have shown, for example, that even among relatively high trust countries there may be significant differences in privacy concerns [6].

² Yamagishi and Yamagishi (1994) argue that this is typical for Japanese firms.

³ Axelrod (1984) considered the so-called prisoner's dilemma game, Selten, Mitzkewitz and Uhlich (1997) a Cournot oligopoly, and Keser (2000) a public goods game.

⁴ Amazon.com uses a different feedback system. Any time a buyer makes a purchase from a seller he or she is encouraged to rate the seller's performance and leave a short comment. The average ratings accompany a seller's name in every reference and appear as one to five stars, with five stars being the best. Also Half.com uses 5 rating categories, numbers from 1 ("poor") to 5 ("excellent") and characterizes each user by the average rating.

⁵ The auction design on online markets per se is a highly interesting question that has been addressed by, for example, Roth and Ockenfels (2002) and Ockenfels and Roth (2002) [20,21]. The role of trust in the auction framework is part of our future research agenda.

⁶ The tripling of the amount invested might be considered as gains from trade.

⁷ The Feedback Forum also mitigates the seller's risks with respect to the buyer's payment.