The Role of Reputation in Open and Closed Groups:

An Experimental Study of Online Trading*

May, 2003

Toshio Yamagishi
Graduate School of Letters, Hokkaido University

Masafumi Matsuda
NTT R & D

Running head: Reputation in open and closed groups

Word count: 16,100

*Address all correspondences to Toshio Yamagishi, Graduate School of Letters, Hokkaido University, N10 W7 Kita-ku, Sapporo, Japan 060-0810, email: Toshio@let.hokudai.ac.jp. The experiments presented in this paper were supported by a grant to the first author from Japan Foundation for the Promotion of Science, 10410029. The first author thanks support from Center for Advanced Study in the Behavioral Sciences.

The Role of Reputation in Open and Closed Groups:

An Experimental Study of Online Trading

Abstract

Greif (1989, 1993) argues that reputations provide a solution to the agent problem deriving from information asymmetry when principal-agent relations are confined within a closed group. This raises an interesting question of whether or not reputations provide a solution to the lemon's market problem (Akerlof, 1970) when the market has open boundaries. In order to address this question, we create an online trade market in the laboratory. In three experiments, we examine how reputation helps alleviate the lemon's problem in open and closed markets. We manipulate openness and closeness of the market by allowing and not allowing traders to change their identities and cancel reputations. When they are allowed to change their identities, they can re-enter the same market under a different identity and a fresh reputation score. We find that a reputation system helps alleviate the lemon's problem in both types of markets, but for different reasons. In *closed* market, having a reputation system greatly improves the market efficiency while individual traders who have accumulated good reputations benefit only little from their reputations. In contrast, individual traders who have acquired good reputations in *open* markets enjoy substantial financial advantage. We argue that the effect of reputation in the closed market is mostly based on traders' untested beliefs, whereas the effect in the open market is based on individual traders' experience to a much greater degree than in the closed market.

The Role of Reputation in Open and Closed Groups:

An Experimental Study of Online Trading

Introduction

Increasing interest in the role that reputation plays in promoting cooperation among egoistic agents characterizes the recent development of game theoretic accounts of social order, both in biology and behavioral ecology and in social sciences. The fundamental question here is: "how is mutual cooperation possible among egoists when cooperation implies self-sacrifice for the welfare of others?" Early research on this issue focused on the dyadic form of interdependence called the prisoner's dilemma game, culminating in the conclusion that mutual cooperation between two egoists is not only possible, but also very likely to occur, when the two interact with each other indefinitely (Axelrod, 1984). However, mutual cooperation among three or more egoistic agents—a situation often referred to as a social dilemma—is more difficult to explain even when they play interacting in the group for a long time. The most obvious solution to the problem of mutual cooperation among multiple egoists is the administration of selective incentives or sanctions (Olson, 1965). However, administration of selective sanctions is insufficient as an endogenous explanation of mutual cooperation unless the problem of the second-order social dilemma (Oliver, 1980; Yamagishi, 1986) is resolved. Like the Aesop story variously known as "Belling the Cat" or "the Mice in Council," administration of selective sanctions as a solution to the problem of social dilemmas faces the problem of who ties the bell to the cat. Since the benefit of sanctions is shared by all members, no egoist is willing to volunteer to bear the cost of administering sanctions against non-cooperators.

Currently, an impressive array of solutions is proposed in various fields of biological and social sciences. Some researchers propose that some forms of sanctions such as ostracism and avoidance is relatively costless, and thus the personal benefit of behaving cooperatively can outweigh personal cost of doing so. A good example of this type of solution is a "tolerated theft hypothesis" of meat sharing among some hunter-gatherers (Blurton Jones 1984). Since large games are beyond one family's consumption, the cost of fighting against other members who demand a share is more than the benefit of keeping every piece of the meat to one's own family. Furthermore, meat sharing constitutes a system of insurance among hunter-gatherers (Kaplan and Hill 1985). The one who claims an insurance premium without having paid dues faces the majority who refuses his claim. Sanctioning non-cooperators is not costly when the sanctioners are the majority, and the cost of sanctions is widely shared. And, this is more true when joining the sanctioning activities is considered a part of the dues—i.e., when only the ones who shared their meat with others *and* joined the collective effort of sanctioning are welcome to the meat sharing system.

The above example of meat sharing assumes that everyone's behavior is more or less transparent to other members of a community. As the group size increases, this transparency is lost and finding out who are non-cooperators and who do not join sanctioning activities is more difficult. One member's direct observation is limited to a few others around them. Reputation comes to play an important role in such a situation. For ostracism and avoidance to function as effective mechanisms to promote cooperation among egoistic agent, access to information through indirect means becomes essential. In other words, reputation—indirectly obtained information about other members' behavior—plays a critical role in providing and maintaining ostracism- and avoidance-based solution to the problem of cooperation.

The first formal model that demonstrates the role of reputation for promoting cooperation

among egoists was proposed by Nowak and Sigmund (1998a, 1998b). A member of a community who helps others earns an “image score” every time she does so, and loses a score every time she hurts others. When a sufficient number of others adopt the strategy of helping only those who have earned a high image score, helping others is a gainful strategy. The one who helps others is helped by similar-minded others, and the community of high image-scorers will thrive. Milinski, Semmann and Krambeck (2002) further elaborate on the nature of reputation that makes its role more effective. Their “standing scores” discriminate between defection on cooperators with high standing in the community, and defection on defectors with low standing in the community. A standing score increases when a member helps another member of high standing, but not when she helps another member of low standing.

In those models of reputation, reputation itself is, in a sense, extraneously provided. Each agent’s behavior is ***automatically recorded*** and reported to every other member. The potential problem of free-riding in the provision of personally gained information—sharing valuable information with others—is mostly ignored. Sharing information based on personal experience with others is often costly. This is especially the case when relative standing in a community is at stake since the advantage of having extra information gives one a competitive advantage. Furthermore, most of the formal models of reputation implicitly assume what we call below the closed social relationship. The “community” is the whole world, and agents cannot escape from it. This assumption may be tolerable for models of the role of reputation in traditional communities to which members belonged all their lives. However, such closed communities are rare in modern societies. Does this imply that reputation-based solution to the problem of mutual cooperation is not a viable solution in societies more open than traditional, closed communities? If reputation-based solution is also viable in such societies, can we apply the same logic to explain why reputation works in promoting mutual cooperation among egoists in

the two kinds of societies? These are the questions we address in this study. Specifically, we have two general goals in this study. One is to empirically examine if reputation provides a solution to the problem of mutual cooperation among contemporary humans—specifically, contemporary Japanese students who participate in an experimental study to earn money. The other is to investigate the role of reputation in closed and open social relations..

Open and Closed Social Relations

As examples of two extreme forms of social relations described above, we refer to the trading society among Maghribi merchants of the 11th century and trading relations among contemporary Internet traders. Both Maghribi traders and online traders successfully apply reputations for controlling the problem of mutual cooperation among egoists. More specifically, the problem these traders face are commonly referred to as the problem of lemons and the principal-agent problem. Results of our study suggest that reputation provide solutions to the problem of lemons on different bases in the two forms of social relations.

The *problem of lemons* is a potential threat to traders who conduct trades without institutional mechanisms for enforcing contracts. As discussed by Greif (1989, 1993), Maghribi traders of the 11th century who conducted business across the Mediterranean faced this problem in dealing with their agents. They faced highly uncertain environments in which they cannot directly observe the quality of the commodities they traded. They thus faced a high level of information asymmetry in dealing with their agents, that is known to produce a "market for lemons" when left with no institutional arrangement to prevent its occurrence. In a classic paper on the market for lemons, Akerlof (1970) argues that asymmetry of information drives honest traders and high quality goods *out* of the market, resulting in a market where only lemons or fraudulent commodities pervade. Akerlof (1970) examines the used car market as an example of a market for lemons. In the used car market, only sellers have information about problems with

the cars they are selling, and most consumers are incapable of discerning these problems. This is also the situation Maghribi traders faced. Because the merchants cannot find if their agents are behaving honestly or cheating them, it is in the agents' best interest to cheat their patrons. Knowing this, it is in the merchants' best interest not to hire "rational" agents. This would have prevented extended trading activities across the Mediterranean. Contrary to this pessimistic theoretical prediction, however, Maghribi merchants were successful in maintaining trades across the Mediterranean Sea.

Contemporary online traders such as users of Internet auction sites face the same problem Maghribi traders or consumers in the used car market face. Individuals who purchase a good online can learn about the quality (or condition) of the good *only after* they have paid for it. In the worst scenario, the good never arrives and the buyer cannot track down the seller. Thus, online traders face greater risk than traders in the "brick and mortar" business (Friedman and Resnick 2001; Houser and Wooders 2001; Kollock 1999; Resnick and Zeckhauser in press; Resnick, Zeckhauser, Friedman, and Kuwabara, 2000; Standifird 2001). The high level of information asymmetry in online trades thus creates a situation ripe for frauds, embezzlements, and the lemons problem. However, students of Internet auction report that relatively few frauds are observed in Internet trades (Kollock 1999; Resnick and Zeckhauser in press; Resnick et al., 2000)[1]. It is interesting to note that both scholars studying online trades and those studying the Maghribi traders have found that the reputation system which is created endogenously keeps the level of fraud low and keeps the problem of lemons from developing. Although the online

[1] While the Federal Trade Commission reports that the number of consumer complaints about Internet auctions have exploded from 107 in 1997 to 10,700 in 1999, the number is relatively small compared to the sheet volume of trades conducted online (Federal Trade Commission press release, February 14, 2000).

trading market and the Maghribi trading market are separated by a millennium, the same solution—the reputation system—operates in both markets as the central mechanism addressing problems arising from information asymmetry.

Before discussing the role of reputation as a solution to the problem of lemons, however, let us briefly present the distinction between shared and non-shared reputation. An individual can learn about her trading partner through direct experience, or she can learn about her trading partner from someone else who has had direct experience with him. We may call the former "private reputation" or "direct reputation" and the latter "public reputation" or "shared reputation." To make the terminology clear, we use "experience-based information" to refer to the former and restrict the use of the term "reputation" to the latter. According to a definition of reputation by Wilson (1985: 27-8), a reputation is a "characteristic or attribute ascribed to one person … by another." Similarly, according to Standifird (2001: 281), "[r]eputation is defined as the current assessment of an entity's desirability as established by some external person or group of persons." According to our terminology, a reputation is a characteristic or attribute ascribed to one person (and is believed to be useful in predicting that person's future behavior) ***by a third-party***. When a person ascribes a characteristic or attribute to another based on her own experience with that person, we call it "experience-based information." Sharing experience-based information with someone else produces a reputation.

The Roles of Reputation among Maghribi and Online Traders

Through a thorough reading of archival records of trades made by Maghribi merchants, Greif (1989) concluded that a unique institution of the *coalition* provided a solution to the problems arising from information asymmetry. Grief (1989) argues that the key to the success of the *coalition* in curtailing opportunistic behavior and promoting "trust" among traders was that reputations were shared within the *coalition*. A reputation system to share experienced-based

information among a group of traders, making that information a public reputation, provides a solution to the problem of lemons in a market characterized with information asymmetry. However, this solution works insofar as the reduction in the expected future profit caused by a reputation of a dishonest trader exceeds the immediate profit derived from dishonest behavior. Greif (1989) argues that the ***closed*** nature of the coalition among Maghribi traders was essential to the success of the reputation system. An agent who behaves dishonestly to a member acquires a reputation for dishonesty, and is excluded from the coalition and thus from future opportunities to trade with any member of the coalition. Exclusion from the coalition and future opportunities to trade with other coalition members works as a deterrent to dishonest trades for the agent only to the degree that he has no alternative trading partners outside the coalition. When the coalition members are the only possible trading partners, exclusion from the coalition means no future trading opportunities. This is a situation in which exclusion from the coalition works as a powerful deterrent to dishonesty.

In contrast, when the agent who cheats a patron merchant and shunned from trading with any coalition member can trade with non-coalition members and make an equivalent level of profit there, exclusion based on reputation information will have no disciplinary power. Okazaki (1999) applies the same logic in his analysis of *kabunakama* in pre-modernization Japan to make this point. During two and a half centuries under the Tokugawa regime before modernization, *kabunakama* or merchants' guild was banned by the Edo (Tokyo) government. Immediately following each time *kabunakama* was banned was a sharp decline in trading activities, and the government was forced to reinstate *kabunakama* to restore trading activities. In the absence of civil intervention by the central government, pre-modernization Japanese merchants had to resort to the closed *kabunakama* coalition to control the agent problem.

The requirement of closed boundaries for reputation to function as a powerful deterrent to

dishonesty is not met in online trade. The effectiveness of a negative reputation system to control the lemons problem is compromised to the degree that dishonest traders have alternative markets to move into without paying exit and entrance costs. Contemporary online traders face a market that greatly differs from the one faced by Maghribi traders of the 11th century. Most importantly, the online market has no closed boundaries. This means that exclusion from the market is difficult. Online traders can assume many identities and/or create many pseudonyms. They can change their identities and pseudonyms when necessary and come back to the same market under a different identity. Furthermore, students of Internet auction sites find another potential problem for a negative reputation system. That is, fear of retaliation prevents online traders from providing a negative evaluation to their trading partners (Resnick and Zeckhauser in press). Thus, using a reputation system as *the* solution to the lemons problem in the online market seems to be a hopeless endeavor.

The reality, however, is much brighter than it is theoretically predicted (Kollock 1999; Resnick and Zeckhauser in press). Reported frauds on the Internet are relatively infrequent (at least compared to the theoretically predicted level) and the reason for the low rate of dishonest behavior has been attributed to the reputation system used in online markets. Many researchers (e.g., Kollock 1999; Resnick and Zeckhauser, in press; Resnick, Zeckhauser, Swanson, and Lockwood 2002; Standifird 2001) claim that the disadvantages of online markets for reputation are well offset by the sheer quantity of quickly and cheaply accumulated and disseminated reputation information. Specifically, observers of online trades (Avery, Resnick and Zeckhauser 1999; Kollock, 1999; Resnick and Zackhauser in press) point out that the rapid development of information technology makes the cost for sharing information very low. They claim that the vast quantity of cheaply available reputation information in online trades offsets the lack of quality and the lack of reliability of information. The studies that examined the effects of

reputation on auctioning prices in eBay, one of the largest Internet auctioning sites, report that those who have acquired a highly positive rating score (good reputation) enjoy an advantage in the price for which they sell their goods. Further, those with a negative reputation score are "punished" when they receive a low bid in the auction (Eaton 2002; Houser and Wooders 2001; Lucking-Reiley, Bryan, Prasad, and Reeves 2000; Resnick and Zeckhauser in press; Standifird 2002; See Resnick et al. 2002, for a review of 12 studies examining the effect of positive and negative reputations on price premium).

Kollock (1999), however, proposes another reason for the theoretically unexpected strong effect of reputation in online trades. He makes a distinction between two types of reputation system—the positive reputation system and the negative reputation system. A positive reputation system evaluates traders only in the positive direction. A new entrant to the system receives a neutral reputation of zero, which is in fact the lowest level of evaluation since all information-based evaluation or reputation is only in the positive range. A negative reputation system, in contrast, evaluates traders only in the negative direction. The neutral reputation of zero that a new entrant (or one who re-enters the market under a new identity) receives is the highest in this system since all information-based evaluation is only in the negative direction. Kollock asserts that traders who have acquired a negative reputation have an incentive to change their identity and re-enter the market, and thus reputation does not accumulate under such a system. On the other hand, traders under the positive reputation system have incentives to maintain a positive reputation because a positive reputation is a valuable asset. The freedom to change identities will not affect the effectiveness of the positive reputation system since traders voluntarily keep their "brand name" once they have acquired a positive reputation. Dellarocas (2003) provides a formal model that requiring newcomers to start with the lowest level of reputation score is helpful to prevent negative effects of "cheap pseudonyms" (Friedman and

Resnick 2001).

The goal of the study reported in this paper is to examine if reputation provides a solution to the lemons problem in an open market as well as in a closed market, and if so, whether or not the directionality of the reputation system makes a difference in the two types of market. Specifically, we address the following set of research questions in a series of laboratory experiments to be presented below. Except one ingeniously designed experiment (Resnick *et al*. 2002), existing studies of reputation in Internet auctioning use actual trading logs. Examining the effect of reputation with actual trading logs may enable researchers to assess the effect of accumulating either positive or negative reputation on the success rate of selling a good and on the price of the good achieved at the auction. However, it is difficult to assess the effect of installing a reputation system on the global performance of the market. For the researchers who study reputation, one of the most central issue is whether or not reputation helps solve the problem of social dilemmas and makes mutual cooperation possible, whether reputation can solve the problem of lemons. This critical problem cannot be empirically assessed with the analysis of existing trading logs, unless data are extracted from two different auctioning sites, one with a reputation system and one without such a system. Perhaps as testimony to its effectiveness, most existing online auction sites use a reputation system. Laboratory experiment, or quasi-laboratory experiment using a web-based laboratory (Yamagishi, Yoshikai and Takahashi 2003), is the ***only*** methodology currently available to empirically assess the system-wide effectiveness of reputation. Another advantage of the experimental methodology is that the sensitive private information in the analysis of reputation—i.e., how honestly traders behave and how honestly they report their personal experience—is accessible. In the existing trading logs including evaluation scores provided by traders, it is impossible to measure how accurate the evaluation scores represent what actually happened. The one who receives a

positive evaluation score may be conducting an honest business, or may be colluding with his band, giving each other positive evaluations to entice naïve users to their sites. Or the buyer may have given a positive evaluation to the seller before or without realizing the true quality of the purchased commodity. In the laboratory experiment, we can record the actual quality of the goods traded and compare it to the evaluations accorded to the seller of the quality. We decided that these advantages of controlled laboratory experiments outweigh the often-cited disadvantage of experiment, or the lack of external validity. This, of course, does not imply that analysis of actual behaviors observable on Internet trading sites is useless. Not at all. What we propose is that carefully designed and conducted laboratory experiments complements studies of actual behaviors. Each has its advantages and disadvantages, and thus each complements the other. Given the paucity of laboratory studies, findings from the experiments will provide useful information to anyone who are interested in the role of reputation.

Research Questions

First, we ask: Do people voluntarily share their own experience-based information with others? And, if they do, does the voluntarily shared reputation provide a solution to the lemons problem? These questions can be addressed only in experimental research, since the true quality of the purchased good is impossible to assess in analyzing the actual evaluation scores provided by participants in on-going Internet auction sites. This is a question concerning free-riding that might occur when individuals are faced with the decision to share experience-based information with others in order to create a reputation (Avery et al. 1999; Resnick et al. 2000; Resnick and Zeckhauser in press). There is no incentive for an online trader to share his own experience-based information with others. Shared reputation is a public good which one can freely access; contributing to its provision is not a condition for access. According to the "logic of collective action" (Olson 1965), a public good that individuals can freely access without

contributing to the provision of it will eventually fail to be provided. Or, if it is provided at all, it will be underprovided. Furthermore, there is an incentive *not* to share one's own experience with others. A person who has found an honest trader may try to keep other people from trading with her in order to monopolize her trading. For this purpose, he will not provide a positive reputation for her and he may even provide a false negative reputation for her. Such a situation occurs, for example, when one finds a nice little restaurant and does not want to see it overcrowded.

The provision of sanctioning is logically identical with the provision of a public good, which involves a second-order social dilemma problem as discussed in the introduction. In short, people can benefit from a sanctioning system without paying the cost for providing and maintaining the system. One solution to this second-order social dilemma problem is the higher-order sanctioning (Axelrod 1986). It is often argued that the higher-order sanctioning (e.g., sanctioning against non-sanctioners) is less costly than lower-order sanctioning (i.e., sanctioning against non-cooperators) (Gintis 2003). Applying this logic to the provision of a reputation system, the second-order dilemma problem can be resolved when those who do not share their personal experiences with others receive a negative reputation and are then treated just as dishonest traders are treated. It is not too difficult to imagine that those who do not join sanctioning activities also become the target of the group's hatred in closely-knit groups. On the other hand, the current reputation systems used in most Internet auction sites such as eBay do not have an option to give a negative evaluation to those who failed to give a negative evaluation to a fraudulent trader. The free-riding problem concerning the provision of reputation is thus expected to be more serious among online traders than it is among Maghribi traders. Whether or not people voluntarily share their own experience-based information with others is the first research question addressed in this paper.

Before examining this issue, however, we address the more fundamental issue of whether or not the predicted lemon's market will occur in our laboratory. We predict that the market characterized with information asymmetry will result in a lemon's market in which only the lowest quality goods are traded. Despite this prediction, previous studies of labor contracts that deals with the same issue found no strong evidence of this prediction. Experimental economists (DeJong, Forsythe, and Lundholm 1985; Fehr, Gachter and Kirchsteiger 1997) examining principal-agent relations have found that agents exhibit the predicted opportunistic behavior in the absence of a continued relationship between specific principals and agents. However, the effort level of the agent in the relationship is not reduced to zero as expected by the rational choice model of the agent. For example, Fehr et al. (1997) examine an experimental labor market in which principals (firms) hire agents (workers) at a wage level *w* set by the principal. The principal, in turn, expects an effort level *e* from his agent. The agent who receives the wage *w* can then freely choose the level of *e* that he actually performs. That is, the contract between the principal and the agent about the level of *e* for *w* cannot be enforced. In this situation, the agent who cares only about his own welfare should exert *zero* level of effort. In contrast to this zero effort prediction, Fehr et al.'s (1997) experiment did *not* find that the agents exerted a zero level of effort in such a situation. Similar experiments by DeJong and his colleagues (DeJong et al. 1985, DeJong, Forsythe, Lundholm, and Uecker 1985) also fail to confirm the *zero* effort prediction by rational choice models of behavior.

According to Fehr and his colleagues (Fehr and Gachter 2000; Fehr et al. 1997), reciprocity explains the failure of the zero effort prediction. This may be true in the context of an exchange of wage and effort. In the experimental online market we create in our laboratory, however, reciprocity is expected to play a minimal role. The economic transactions of selling and buying are less likely to activate the norm of reciprocity than the exchange of labor for

wages. Thus, we expect that the pure market for lemons will result in our experimental market in which players are not identifiable to each other.

Hypothesis 1: In an experimental market in which players cannot be identified and contracts cannot be enforced, the quality of the goods sold will be reduced to the lowest level.

At the same time, we predict that if the players *can* be identified, the lemons problem will be alleviated. Players can accumulate experience-based information when they can be clearly identified, so that players can avoid interacting with known "cheaters."

Hypothesis 2: The ability to identify other players and to remember their past behavior will improve the quality of the goods sold in the market.

Furthermore, reputation—indirect information obtained from others—will improve the quality of the goods sold in the market if, and only if, players share their experience-based personal information with other players. Because we have no theoretical basis for the claim that players will voluntarily provide their experience-based information *honestly*, we test the following hypothesis as a tentative hypothesis.

Hypothesis 3: The information the buyer provides to other buyers concerning the seller's behavior will reflect the seller's *actual behavior*.

If Hypothesis 3 is confirmed, we will test the following hypothesis:

Hypothesis 4: Reputation—i.e., information obtained from other players—will improve the quality of the goods sold in the market.

In a study of principal-agent (firm-worker) relations, DeJong, Forsythe, Lundholm and Uecker (1985) examined the effect of sharing experience-based information on the effort level exerted by the workers in the experimental labor market they created in their laboratory, and found a positive effect of information sharing. However, in their experiment, information-sharing and

experience-based information were manipulated simultaneously such that the effects of the two could not be separated. We will examine the role of reputation independent of the role of experience-based information in our experiment.

The second research question concerns the contrast between a closed market—i.e., trades among Maghribi members—versus an open market—i.e., online trades. Online traders can have as many identities or "handles" as they wish, *and* they can change these identities as often as they wish. An online trader who has accumulated negative reputations can shake them off by assuming a new identity with a new email address and thus pretending to be another person. As a result, the effectiveness of reputation as a solution to the problem of lemons can be greatly undermined in online trading. Thus, we address in this study whether stability of players' personal identities is required for reputation to be effective in solving the lemons problem. Specifically, we examine the following hypothesis in the second experiment:

> *Hypothesis 5*: When players are permitted to freely change their identities, the quality of the goods sold in the market will be reduced *even if* both experience-based information and reputation about other players are available.

The problem of unstable identities discussed above is especially serious with respect to negative reputations, which, according to research on trades on eBay (Houser and Wooders 2001; Lucking-Reiley et al. 2000; Resnick and Zeckhauser in press; Standifird 2001) have a more powerful effect on price than do positive reputations. As suggested by Kollock (1999), however, positive reputations are free from this problem since online traders have every incentive to maintain the positive reputation that accompanies their identity. Having a positive reputation is a valuable asset, and traders who have acquired a positive reputation will not voluntarily change their identities. The third research question we address in this study is whether the effectiveness of reputation as a solution to the lemons problem varies with the

direction—positive or negative—of the reputation. To examine this question, we introduced a positive reputation system and a negative reputation system in the third experiment.

Hypothesis 6: The quality of the goods in the market will be higher in the positive reputation system than it will be in the negative reputation system.

Basic Features of the Experiment

The experiment was conducted in Japan with Japanese participants using Japanese yen as incentives. In this experimental market, all players (i.e., voluntary student participants) simultaneously take on two roles: the buyer and the seller. As a seller, each player produces a "commodity" (an abstract entity without any substance) by expending some cost, ranging from 10 yen to 100 yen per commodity, in increments of 10 yen.[2] The quality of the produced commodity is determined by the amount of money a seller has expended for its production. For simplicity, the quality is expressed in terms of the amount of money (yen) the seller has spent for its production, ranging from 10 to 100 as an integer. For example, a commodity with the quality of 70 is a commodity that the seller spent 70 yen to produce. The commodity produced by a seller is then put on the market for sale with the price set by the seller. At the same time, the seller announces the quality of the commodity—i.e., how much she spent for its production. The "advertised" quality may or may not reflect its true quality. For example, a seller who has spent 30 yen for the production of a commodity may put it on the market for sale at 100 yen with an announcement that its quality is 80 (whereas its true quality is 30). Any other player, acting as a buyer, can purchase the commodity for that price, 100 yen. Once a buyer purchases the commodity, it disappears from the market. That is, each commodity can be purchased only by one buyer. A commodity placed on the market remains on the market for 5 minutes, after which

[2] The exchange rate at the time the experiment was roughly 100 yen to the dollar.

unsold commodities perish. They disappear from the market and the producer/sellers cannot recover the cost of production. This option was introduced to the current experiment to avoid overcrowding the market with too many commodities. Once a commodity is placed on the market, the seller of the commodity cannot change its advertised quality or its price.

As a buyer, each player may purchase any of the commodities placed on the market. The information available to a buyer prior to the purchase is the price and the advertised quality of the commodity. Each commodity is worth 1.5 times its true quality to the buyer; the experimenter pays the buyer 1.5 times the true quality of a commodity. Through this "reimbursement" deal that immediately takes place after a buyer purchases a commodity on the market, the buyer finds out the true quality of the commodity he has just purchased. The difference between the value to the buyer (1.5 times the true quality of the commodity) and the purchase price becomes the profit for the buyer. In the above example of a trade in which a buyer pays 100 yen for a commodity with true quality of 30, the buyer loses 55 yen since the commodity is only worth 45 yen (1.5 times its true quality of 30). As another example, s seller spends 80 yen for the production of a commodity, advertises that its quality is 90, and sets the price at 110. The commodity is worth 120 yen to a buyer. If someone purchases it, the seller makes 30 yen (110 – 80), and the buyer makes 10 yen.

As mentioned before, each player takes on the role of a seller as well as that of a buyer. He can act as a buyer at any time. On the other hand, he can act as a producer/seller only sporadically. Specifically, each player is provided with a "production opportunity" every 50 seconds with a random variation in the range of plus and minus 15 seconds. Once a production opportunity is provided, a player is given 100 yen from the experimenter as an endowment, and decides how much of it to invest in the production of a commodity (in increments of 10 yen). A player is required to produce one and only one commodity per production opportunity. She

cannot act as a buyer until she finishes producing a commodity. As explained above, how much she spends for the production of a commodity determines the quality of the commodity she produces. The remaining amount of the endowment (100 yen minus the amount she invests in the production of a commodity) is hers to keep. In addition to the price and the advertised quality of the commodity, each commodity for sale on the market may come with other information, depending on the experimental conditions, i.e., the identity or reputation of the seller.

The experiment was conducted in a fully computerized laboratory at a major national university in Japan. Upon arrival at the laboratory, each participant was escorted to her individual room. Each room was equipped with a computer that was connected via LAN to a host computer controlled by the experimenter. All instructions were presented on the computer screen. The production and purchasing decisions made by the participant were entered to the participant's computer with a click of a mouse. In each session, participants were allowed to produce commodities as many times as possible for the first two minutes. That is, a new production opportunity was given to the participant as soon as she finished producing a commodity during the first two minutes of the session. The experiment was conducted in this way to provide the market with a sufficient number of commodities with which to start. After the first two minutes, a production opportunity was provided to the participant every 50 seconds with a random variation within the range of plus or minus 15 seconds.

Participant's Computer Screen

Figure 1 is an example of how the market looks on a participant's computer screen. The screen is divided into upper and lower sections. The upper section of the screen displays commodities placed for sale by other players, and the lower section of the screen displays commodities the player herself has produced and placed for sale in the market. Information

related to the player's personal "account" is also displayed on the right side of the lower section of the screen. In the example of the display shown in Figure 1, the other players have placed ten commodities on the market. Each box shown in the upper section represents a commodity. In the box located on the lower-left corner of the screen, for example, the seller/producer *to-me-ko* sells the commodity for 40 yen, while advertising its quality to be 30. Please note that the identity of the producer/seller is provided in this example of the identity condition of Experiment 1, but this is not always the case in other conditions. Note also that the original screen is written in Japanese, not in English. The player's identity is expressed in three Japanese characters. In the example of *to-me-ko*, "to," "me," and "ko" stand for three distinct Japanese characters. Each participant's screen is unique, since the commodities the player has produced does not appear in the upper section of the screen for that player. Furthermore, the boxes in the upper section are periodically shifted so that it is difficult to keep track of all of the commodities on the market in the condition in which identities are not stable.

Insert Figure 1 About Here

The lower section of the display in Figure 1 indicates that this player *se-se-i* is currently selling two commodities that are represented by the two boxes shown in that section. To produce the first commodity displayed as the box on the far left side of the lower section, this player spent 10 yen. She advertises its quality as 100, and is selling it for 140 yen. If someone purchases it, the seller, *se-se-i*, makes a profit of 130 yen, since she spent only 10 yen to produce the commodity that she sells for 140 yen. The buyer who purchases this commodity loses 135 yen since it is worth only 15 yen (1.5 times its true quality of 10). The second box to the right represents another commodity this player is currently selling on the market. The right corner of the lower section displays the identity of this player, *se-se-i*, and her current financial situation. According to the information displayed there, this player has thus far made a profit of 170 yen,

including revenue from the sales of her commodities and the part of the "endowment" of 100 yen that she did not use for the production of her commodity. If the experiment ends at this moment, this player will be paid 170 yen for her participation in the experiment.

Experiment 1: The Effect of the Player's Identity and Reputation

We conducted the first experiment to examine the first four hypotheses outlined previously. The research design to test these hypotheses includes the following three conditions:

Anonymity condition. In the anonymity condition, the commodities placed on the market have no identity markers. Buyers are unable to find out who the seller of a particular commodity is. In this purely anonymous market with no possibility of intentionally repeating transactions between a particular seller and a particular buyer, we predict that sellers will come to sell only the lowest quality goods (Hypothesis 1). Even honest sellers are forced to sell the lowest quality goods since, given the possibility of lemons, buyers will come to pay only the price for the lowest quality goods. The comparison of the quality of the commodities produced in this market with the theoretically predicted level of 10 (i.e., the lowest possible quality) serves as a test of the first hypothesis. Since only the lowest quality of commodities are produced and traded, participants in this condition fail to attain the benefit of trading commodities of high quality. When a seller produces a commodity with the quality of 100, he and the buyer of that commodity share the extra value of 50 yen provided by the experimenter. When a seller produces a commodity with the quality of 10, the extra value they share is only 5 yen. Sellers and buyers together forgo opportunities to enjoy a larger value to share when only the lowest quality commodities are produced and sold on this anonymous market. Thus, the average level of qualities of the commodities produced in the market is a direct indicator of the market efficiency. When only the lowest quality goods are produced and sold on the market,

only 10% of the maximally possible profits is achieved.

Identity condition. The second condition is one in which commodities placed on the market come with identity markers (one of 48 Japanese characters[3]). In this condition, buyers can use "experience-based information" about sellers' past behavior. They would avoid purchasing commodities from the seller with whom they have had sour experiences, and would prefer purchasing commodities from the sellers who have sold them profitable commodities. Thus, we predict that giving players a permanent identity will improve the market and reduce the likelihood of lemons to a certain degree (Hypothesis 2). Comparing the average quality of the goods sold in this condition with that in the anonymity condition serves as the test for the second hypothesis.

Reputation condition. The last condition used in the first experiment is the reputation condition. As mentioned in the introduction, we use the term "reputation" as indirectly acquired information. To assess the effect of reputation independent of the effect of experience-based information (available in the identity condition), only reputation information is made available in the reputation condition. Unique identities are not attached to the commodities placed on the market. Each buyer evaluates the seller from whom she buys a commodity. Evaluation is made by assigning a number between -2 and +2 (-2 = very bad, -1 = bad, 0 = neutral, +1 = good, +2 = very good) to the seller from whom the buyer has just purchased a commodity. The cumulative sum of evaluation scores becomes the seller's overall score. Each commodity placed on the market for sale by a seller comes with a signal indicating the overall evaluation score of the seller at the time she made the product. Since providing an exact evaluation score for each seller makes it possible for the buyers to identify a seller by his evaluation score, an evaluation score

[3] The number of letters was increased to three in the second and the third experiments.

for the seller is translated into a color and the shade of the color. Specifically, each box representing a commodity produced by a particular producer/seller is colored either blue or red. Blue indicates the total evaluation score of that producer/seller is positive, and red indicates that the total evaluation score is negative. In addition, the shade of blue or red reflects the absolute value of the total evaluation score. That is, dark blue indicates a highly negative score while dark red indicates a highly positive score. Light blue indicates a slightly negative score and light red indicates a slightly positive score. The screen is flushed periodically and the locations of the commodities on the display change places randomly each time. Thus, it is impossible for the players to identify the producer/seller of each commodity, and yet it is possible for them to determine how other buyers evaluate the producer/seller of each product on the market.

If buyers honestly reveal their evaluations of the producer/sellers from whom they have purchased a commodity (Hypothesis 4), the quality label expressed as the color of the box will provide at least a partial solution to the lemons problem (Hypothesis 5). Buyers will avoid commodities with negative reputation (in red color) and will seek commodities with positive reputation (in blue color). The buyer's response to the reputation information will exert discipline on the seller, since those who receive negative reputation will have hard time successfully selling their commodities. The sellers who want to earn good reputation for successfully selling their commodities have to behave in a way that does not offend the buyer. Examining how honestly or accurately buyers report their evaluations of the seller from whom they have bought commodities will provide a test of the third hypothesis. The comparison of this condition with the anonymity condition serves as a test of the fourth hypothesis.

Procedure

Fourteen students (10 males and 4 females) were recruited from a participant pool of about 1,500 students enrolled in a major university in Japan to participate in the experiment.

Monetary incentives were emphasized, and no class credit was involved in the recruitment of the participants. Participants were randomly assigned to two groups, and each group consisted of seven participants. Each group participated in six experimental sessions. The first three sessions lasted 45 minutes each, and the last three sessions lasted 20 minutes each. In total, the experiment lasted an entire afternoon (approximately 5 hours). Participants were paid the amount they earned in the experiment. The most any one participant earned was 8,965 yen (about $80) and the least was 910 yen (about $8). The average pay was 5,597 yen. Each of the three 45-minute sessions was assigned to one of the three conditions. We added the 20-minute sessions to learn how the experience of the first three sessions would affect the participants' behavior. Each of the three 20-minute sessions were assigned to one of the three conditions as well. Participants were not told how long each session would last. The order of presentation of the three conditions was randomized for each group. A rest period of 10 minutes was provided between sessions.

Findings

Figure 2 presents the average level of quality of the commodities produced and placed on the market during the 45-minute session and the 20-minute session. Each session is broken into time blocks lasting five minutes each. The main effect of the condition in a condition (a repeated factor) × time block (a repeated factor) ANOVA on the 45-minute session indicates a significant main effect of the conditions, $F(2, 39) = 16.73$, $p < .0001$. The main effect of the time block, $F(8, 312) = 2.05$, $p < .05$, and the conditions × time block interaction, $F(16, 312) = 4.34$, $p < .0001$, were also significant. As shown in the figure, both the main effect of the time block and the interaction effect reflect the downward trend of the average quality of the commodity that occurred only in the anonymity condition. A post-hoc analysis of mean quality in the three conditions (using a Tukey's studentized range test) indicates that the mean quality of the

anonymity condition (29.65) is significantly different from the mean quality of the other two conditions (70.76 in the identity condition and 71.07 in the reputation condition). The mean quality of the identity condition and the mean quality of the reputation condition were *not* significantly different from one another. A similar analysis was conducted on the 20-minute sessions, and similar results were obtained. As in the 45-minute session, the main effect of the conditions was significant, $F(2, 39) = 41.71$, $p < .0001$. The main effect of the time blocks was also significant, $F(3, 117) = 12.74$, $p < .0001$. The conditions × time blocks interaction effect, however, was not significant, $F(6, 117) = 1.60$, *ns*. As in the 45-minute session, there was a significant difference between the mean quality in the control condition (16.52) and the mean quality in the other two conditions (67.51 in the identity condition and 76.89 in the reputation condition). These results support hypotheses 2 and 4.

The first hypothesis is also supported by the data. The average quality of the commodities in the control condition approached the theoretically predicted lowest quality of 10 in either the first 45-minute session or in the second 20-minute session. The average quality during the last time block (11.96) of the 45-minute session is only very slightly above 10, and is not significantly different from 10, $t(13) = 1.69$, *ns*. In the 20-minute session, this equilibrium was reached more quickly. The average quality was not significantly different from 10 in the 3rd time block (10.94, $t(13) = 1.46$, *ns*.) and the 4th time block (10.29, $t(13) = 1.00$, *ns*.) These results confirm the first hypothesis.

Insert Figure 2 About Here

Finally, we tested the third hypothesis by calculating the within-participant correlation between how fraudulent the commodities were that a participant purchased and the evaluation scores she gave to the sellers of those commodities. To calculate this correlation, we first gave each commodity purchased by a buyer a “fraudulence score,” which is the advertised quality

minus the true quality. For example, a commodity that is advertised for the quality of 80 while its true quality is 30 receives a fraudulence score of 50 (80 – 30). We then calculated a correlation for each participant between the fraudulence scores of the commodities she purchased and the evaluation scores she gave to the sellers of those commodities. A negative correlation means that the participant evaluated the seller accurately and honestly. A correlation around zero means the evaluation was random or inaccurate. And, a positive correlation between the fraudulence scores and the evaluation scores means that the buyer gave an intentionally false evaluation of a seller. To avoid confusion, we reversed the sign of the correlation, so that high scores reflect honest evaluations and low scores reflect dishonest evaluations. We call this inverse of the correlation the *index of honest evaluation*. The average index score of honest evaluation in the reputation condition was .67, $p < .0001$, in the 45-minute session and .73, $p < .0001$, in the 20-minute session. These results indicate that the participants in this experiment fairly honestly and accurately reported their personal experience-based information to be shared with others. This provides a support for the third hypothesis.

In addition to confirming all four hypotheses stated earlier, the data also provide further interesting findings. Figure 3 shows the "fraudulence level"— the difference between the advertised quality and the true quality—in the three conditions. As shown in the figure, the fraudulence level was much higher in the anonymity condition than in the other two conditions. In the ANOVA of the level of dishonesty for the 45-minute session, the main effect of the conditions was significant, $F(2, 39) = 5.36$, $p < .01$. The main effect of the time block, $F(8, 312) = 1.93$, $p < .06$, and the conditions × time blocks interaction effect, $F(16, 312) = 1.56$, $p < .08$, were marginally significant. For the 20-minute session, the main effect of the conditions was significant, $F(2, 39) = 3.61$, $p < .01$, and so were the main effect of the time blocks, $F(3, 117) = 2.93$, $p < .05$, and the interaction effect, $F(6, 117) = 4.74$, $p < .001$.

Insert Figure 3 About Here

Another interesting finding concerns the relationship between dishonesty and profit. To analyze this relationship, we performed a regression analysis in which the independent variable was the participant's *fraudulence level* (the average fraudulence score of the commodities the participants produced) and the dependent variable was the *profit* participants earned by producing and selling commodities (excluding the profit they earned from purchasing deals). To control for between-groups differences, a dummy variable for the groups (0=Group 1 and 1=Group 2) was added to the list of independent variables. We examined the correlation between the fraudulence level and the profit level in each condition. We conducted the same analysis separately for the 45-minute session and the 20-minute session. In the anonymity condition, the fraudulence level was positively related to profit in both the 45-minute session, $b = 30.03$, $t(11) = 4.36$, $p < .01$, and the 20-minute session, $b = 4.38$, $t(11) = 1.23$, *ns*. In the 45-minute session, an average increase of one point in fraud (i.e., overstatement of the quality) led to an increase of about 30 yen in seller's profit. The average fraudulence level in this condition was 26.53, and thus the average participant made 803.86 yen more than an honest participant with the fraudulence level of zero whose expected profit was –79.39 yen. Dishonestly was clearly a better strategy for making profits than honestly revealing the true quality of the commodity. The average fraudulence level in the second session was 26.72, and the average seller thus made 117.03 yen more than the honest seller. The honest seller with the fraudulence level of zero was expected to make –115.64 yen. In both sessions, the honest seller would have lost substantial amount of money. This is because substantial proportion (32% in the first session and 29% in the second session) of the commodities placed on the market did not sell.

In the identity condition, the result of the same regression analysis indicates a weaker positive effect of fraudulence level on profit. In the 45-minute session, the effect was positive

but not significant, $b = 11.87$, $t(11) = 1.23$, *ns*. The average sales profit was 1975.71 yen, and the expected sales profit of the honest seller was 1807.02 yen. In the 20-minute session, the effect was positive and statistically significant, $b = 14.44$, $t(11) = 3.02$, $p < .05$. The average seller's profit was 810.00 yen and the expected profit of the honest seller was 607.88 yen. Given this finding that dishonesty increased profits in the identity condition, the ability of identity information to curtail the lemons problem seems precarious. At the same time, sellers in the identity condition were not punished with financial loss for conducting honest deals as in the case of the anonymity condition. Finally, in the reputation condition, the regression coefficient for fraudulence in the 45-minute session was negative though not significant, $b = -8.49$, $t(11) = .46$, *ns*. On average, sellers made a profit of 1704.29 yen, and the honest seller was expected to make a profit of 1779.41 yen. In the 20-minute session, the regression coefficient was positive and not significant, $b = 11.16$, $t(11) = 1.29$, *ns*. The average seller's profit was 1006.43 yen, and the expected profit of the honest seller was 880.46 yen. These results suggest that reputation information seems more promising than identity information or experience-based information for resolving the lemons problem.

In the reputation condition, the overall level of reputation of a seller (defined as the average reputation score of the commodities she produced during the last time block of the experimental session) ranged from the minimum of –14.60 to the maximum of 44.40, with a mean of 11.62, in the 45-minute session, and from –10.80 to 19.57 with a mean of 3.97 in the 20-minute session. As expected, the average fraudulence level of a seller has a significant negative effect on the overall reputation score, controlling for group differences, $b = -1.67$, $t(11) = 2.25$, $p < .05$ in the 45-minute session and in the 20-minute session, $b = -0.30$, $t(11) = 1.82$, $p < .05$, one-tailed. We then examined the effect of the seller's overall reputation on her profit as a seller with a similar regression analysis controlling for group differences by adding a dummy

variable representing the group difference in the list of the independent variables. The regression coefficient is positive, b = 8,88, $t(11) = 1.58$, *ns*., in the 45-minute session and negative, b = -6.83, $t(11) = 0.41$, *ns*., in the 20-minute session. Neither of these effects is statistically significant. When the level of fraudulence is controlled, both become positive, though not significant (b = 10.62, $t(10) = 1.50$, *ns*., in the 45-minute session, and b = 2.18, $t(10) = .13$, *ns*., in the 20-minute session). This result poses an interesting question. On the one hand, the introduction of reputation greatly improved the overall level of the commodities produced in the experimental market, compared to the quality in the anonymity condition. This resulted in a tremendous difference in the overall profit participants in these two markets made. Participants in the anonymity condition made an average of 228.21 yen from selling and buying in the 45-minute session and lost 67.50 yen in the 20-minute session[4], whereas those in the reputation condition made an average profit of 1770.00 yen in the 45-minute session and 986.07 yen in the 20-minute session. The difference was highly significant, $t(13) = 3.89$, $p < .01$ in the 45-minute session and $t(13) = 5.90$, $p < .0001$ in the 20-minute session. On the other hand, the above result indicates that those who have received good reputations enjoyed only minor advantage in achieving better profit. This contrast between the group-level advantage of the reputation system and the individual-level advantage of having good reputations suggests that *the effect of a reputation system is a group-level phenomenon rather than an aggregation of individual-level effects*. We will come back to this issue in the general discussion.

The results of the first experiment supported all four hypotheses. Clearly, the market for lemons that was predicted to result in the anonymity condition (where neither identity nor reputation was provided) did emerge. Providing identity or reputation of the sellers alleviated

[4] These figures are the total profits of the participants, including both profits from sales and purchase. Buyers lost money in purchasing deals in the anonymity condition.

the problem of lemons. However, the existence of either identity or reputation provided only limited success in solving the lemons problem despite the fact that participants provided fairly accurate evaluation of the seller. The findings from these experiments provide a benchmark for evaluating the results of the second and the third experiments.

Experiment 2: Stability of Reputation

The second experiment was conducted to test the fifth hypothesis, which predicted that when players are permitted to change their identities freely, the quality of the goods sold in the market would be reduced. In the first experiment, identity and reputation each provided at least a partial solution to the problem of lemons. However, such results may critically depend on a unique feature of the experimental market used in the first experiment. That is, identity and reputation are permanent in the sense that the players themselves can not cancel them. The question we address in the second experiment is whether or not identity and reputation have a positive effect on resolving the problem of lemons when players can freely change their identities and thus erase their reputation at will.

Procedure

The purpose of the second experiment is to test the fifth hypothesis about the effect of personal identity on the quality of commodities sold on the market. To achieve this purpose, we used only one condition in the second experiment. To test the fifth hypothesis, we compare the results of this condition with the three conditions of the first experiment. This condition is one in which *both* the identity and the reputation of the seller are provided. That is, the commodity placed on the market by a seller has an identity label of the seller and her reputation color. The seller chose her identity as a combination of three Japanese characters at the start of the experiment, but she could change it at any time during the experiment. Each time the participant

chose to change her identity, a new set of three Japanese characters randomly generated by the computer was assigned to her. Since both identity and reputation had positive effects on the quality of the commodities in the first experiment, we predict that the average quality in this condition will be at least as high as the average quality in the identity or the reputation condition in the first experiment *if* the freedom of changing identities has no effect. We examine this prediction by comparing the average quality in the second experiment with those in the identity and the reputation conditions in the first experiment.

As in the first session of the first experiment, the experiment lasted for 45 minutes. Thirty-four participants participated in five groups. The experimental session was not repeated in a 20-minute session as it was in the first experiment. Each participant experienced only one 45-minute session and was paid the amount she earned in that session. The average payment was 881 yen with the minimum of 500 yen and the maximum of 3,195 yen.

Findings

The average quality of commodities produced and sold during the 45-minute experimental session (broken to 9 time blocks) is presented in Figure 4. The overall average quality of the commodities was 44.09 (sd = 19.94), which was much higher than that in the anonymity condition of the first experiment (29.64, sd = 21.95, $t(46) = 2.22$, $p < .05$). However, it was less than the average quality in the identity condition (70.76, sd = 25.27, $t(46) = 3.89$, $p < .001$) or in the reputation condition (71.07, sd = 17.44, $t(46) = 4.41$ $p < .0001$) of the first experiment.[5] These results clearly confirm the fifth hypothesis.

Insert Figure 4 About Here

The average index of honest evaluation (i.e., the inverse of the within-participant

[5] Only the 45-minute sessions in the first experiment were used in these comparisons.

correlation between fraudulence score of the purchased commodities and the evaluations the participant gave them) was .57, which is significantly greater than zero, $p < .0001$. This suggests that the weaker effect of reputation observed in the second experiment is not due to the participant's unwillingness to honestly and accurately report their evaluation of the seller.

The results of the second experiment indicate that dishonesty pays in this market in which traders can freely change their identities and cancel their reputations. The regression coefficient for the fraudulence level on the profit from producing and selling commodities was positive and significant, $b = 23.20$, $t(28) = 3.71$, $p < .001$. The average seller's profit was 544.41 yen, and the honest seller's profit was 116.88 yen. Overall, the limited positive effect of identity and reputation observed in the first experiment was weakened by the option to freely change identities and cancel reputations in the second experiment, despite the fact that both identity and reputation were available there.

The overall level of reputation of the seller ranged from the minimum of –16.00 to the maximum of 47.20, with a mean of 6.61. The effect of the average fraudulence level of a seller on her overall reputation, controlling for group differences with a set of dummy variables, is negative and significant, $b = -0.38$, $t(27) = 3.40$ $p < .01$. We then examined the effect of the seller's overall reputation on her sales profit with a similar regression analysis controlling for group differences. The regression coefficient is negative, though not significant, $b = -10.01$, $t(27) = 1.58$, *ns*. When the level of fraudulence is controlled, the coefficient becomes positive, though not significant ($b = 10.98$, $t(26) = 1.10$, *ns*). As in the first experiment, a reputation system substantially improved the market efficiency, but the individual advantage of acquiring good reputations was not strong.

Finally, we analyzed the frequency of identity changes. On average, participants changed their identities 7.00 times. While 38% of the participants did not change their identities at all,

23.5% of the participants changed their identities more than 10 times. As expected, the correlation between the frequency of identity changes and reputation was negative, $r = -.41$, $p < .05$. The participants whose overall average reputation scores were zero or less (n = 16) changed their identities 12.06 times on average, whereas those with a positive average reputation score (n = 18) changed their identities only 2.50 times on average. Among nine participants whose overall (i.e., cumulative) reputation score was 5 or more changed their identities only 0.33 times on average. The correlation between the fraudulence level of the participants and the frequency of identity changes was more pronounced, $r = .88$, $p < .0001$. These figures confirm our expectation that it is mostly dishonest sellers with negative reputations who change their identities. Finally, the correlation between the frequency of identity changes and profits from producing and selling commodities was found to be positive, $r = .41$, $p < .05$. In conclusion, selling fraudulent lemons and frequently changing identities constituted a "smart" strategy in the market where players could freely change identities.

Experiment 3: Positive and Negative Reputation Systems

The results of the second experiment show that the freedom of adopting new identities and canceling reputations does not completely nullify the positive effect of identity and reputation. Specifically, the average quality of the commodities in the condition in the second experiment was still higher than the average quality of commodities in the anonymity condition in the first experiment. This positive "residual" effect of the unstable identity and reputation may be based on the fact that those who have received a positive reputation did not want to change their identities. Participants want to cancel negative reputations, yet want to maintain positive reputations. In fact, the participants who have earned positive reputations rarely change their identities, whereas those who have earned negative reputations change their identities many

times. Thus, the positive “residual” effect of reputation found in the second experiment might be the result of the existence of sellers with positive reputations who do not change their identities.

The third experiment was conducted to test the sixth hypothesis concerning the effectiveness of the two types of reputations: positive and negative. Suppose that the evaluations of sellers that buyers provide are only in the negative direction. Specifically, suppose that evaluation scores used in the above experiments have only three levels, 0 for neutral, -1 for bad, and -2 for very bad. Such a scale would make reputations completely meaningless, since those who earn negative reputation scores soon change their identities and enter the market under new identities with a reputation score of zero. If this happens, practically everyone will have a reputation score of zero, which is equivalent to the situation in which no reputation is provided.

Next, suppose that the evaluations of sellers that buyers provide are only in the positive direction. Specifically, suppose that evaluation scores used have only three levels: 0 for neutral, 1 for good, and 2 for very good. We predict that the positive effect of reputation will be stronger here. Since even a slightly positive reputation is better than having none, anyone who has earned a positive evaluation score should be motivated to maintain it. The incentive to maintain a positive reputation will discourage a seller from changing identities. The players who have earned high reputation scores will be strongly motivated to maintain their current identity. Thus, the detrimental effect found in the condition where players can freely change their identities will decrease in a positive reputation system. We conducted the third experiment to examine this prediction.

Procedure

The experimental procedure used in the first two experiments was employed in the third and final experiment. Thirty-two participants were recruited from the same participant pool used in the first two experiments. Twenty participants were assigned to the positive reputation

condition in three groups, and the remaining 12 were assigned to the negative reputation condition in two groups. As in the second experiment, participants were allowed to assume a new identity, a randomly generated new combination of three Japanese characters. When they assumed a new identity, they started over with a reputation score of zero. In the positive reputation condition, participants assigned a positive evaluation score (0 for neutral, 1 for good, 2 for very good) to the seller every time they purchased a commodity. In the negative reputation condition, participants assigned a negative evaluation score (0 for neutral, -1 for bad, -2 for very bad) to the seller every time they purchased a commodity. The experimental session was not repeated. Each participant experienced only one 45-minute session and was paid the amount she earned in that session. The average payment was 675 yen with the maximum of 1,540 yen and the minimum of 500 yen.

Findings

In hypothesis six, we predict that there will be a significant difference between the average quality of the commodities produced in the positive reputation system and that in the negative reputation system. This difference is apparent in Figure 5, especially toward the end of the experiment. The overall quality is 48.46 (*sd* = 22.02) in the positive reputation system and 38.94 (*sd* = 20.60) in the negative reputation system. While the main effect of the condition is not statistically significant, $F(1, 29) = 1.33$, *ns*., the condition × time block interaction is statistically significant, $F(8, 232) = 7.48$, $p < .0001$. The main effect of time blocks was not significant, $F(8, 232) = 1.58$, *ns*. The difference between the two conditions in the average quality reached the significance level in the last time block, $F(1, 29) = 9.06$, $p < .01$. These results provide support to the sixth hypothesis concerning the differential effects of positive and negative reputations on the market.

Interestingly, the average quality was higher in the negative reputation condition (55.21)

than in the positive reputation condition (40.57) in the first time block, and the difference was significant, $F(1, 29) = 4.84$, $p < .05$. That is, negative reputations initially had a strong impact, but that effect quickly waned as sellers who had acquired negative reputations started to change their identities.

Overall, the negative reputation system with variable identities provided only a minor improvement in the average quality of produced commodities compared to the anonymity condition in the first experiment (where no identity or reputation information was provided (29.65 vs. 38.94, $t(24) = 1.11$, *ns*.). The positive reputation system was more promising. The overall average quality of the commodities was significantly higher in the positive reputation system than in the anonymity condition of the first experiment (29.65 vs. 48.46, $t(32) = 2.45$, $p < .05$). And, as shown in Figure 5, the difference increased toward the end of the 45-minute session. In the last time block, the difference between the overall average quality in the anonymity condition in the first experiment and that in the positive reputation system was fairly large (11.96 vs. 60.28, $t(31) = 4.95$, $p < .0001$).[6] The average quality in the positive reputation condition was lower than that in the reputation condition in the first experiment (48.46 vs. 71.06, $t(32) = 3.20$, $p < .01$). However, this difference was smaller and not significant in the last time block (60.28 vs. 77.31, $t(31) = 1.60$, *ns*.) These findings seem to suggest that the positive effect of the positive reputation system takes time to be realized by the participants.

Insert Figure 5 About Here

The average index of honest evaluation was significantly greater than zero both in the negative reputation condition (.66, $p < .0001$) and the positive reputation condition (.41, $p < .0001$). The fraudulence level had a negative but non-significant effect on profit from

[6] The degrees of freedom (*df*.) in this t-test is 31, instead of 32, because one of the participants in the positive reputation system did not produce any commodities during the last time block.

producing and selling commodities in the positive reputation condition, b = -8.82, $t(16)$ = 1.14, *ns*. Sellers in this condition made an average profit of 357.00 yen, and the honest seller was expected to make 452.31 yen. In the negative reputation condition, dishonesty had a positive but insignificant effect, b = 3.36, $t(9)$ = .49, *ns*. The average seller's profit was 190.00, and the expected profit of the honest seller was 125.21 yen.

The overall level of reputation of the seller ranged from the minimum of 0 to the maximum of 50.00, with a mean of 25.51 in the positive reputation condition. The range was between –7.00 to 0, with the mean of –0.75 in the negative reputation condition. The effect of the average fraudulence level of a seller on her overall reputation, controlling for group differences, is negative and significant, b = -1.06, $t(15)$ = 4.34 $p < .001$ in the positive reputation condition, while positive and non-significant, b = 0.06, $t(9)$ = 1.56 *ns* in the negative reputation condition. The seller's overall reputation has a positive and significant effect on her profit in the positive reputation condition, b = 14.55, $t(15)$ = 3.11, $p < .001$, but not in the negative reputation condition, b = -2.73, $t(11)$ = 0.05, *ns*. When the level of fraudulence is controlled, the positive effect of overall reputation becomes stronger in the positive reputation condition, b = 26.87, $t(14)$ = 4.64, $p < .001$, while the effect stayed negative and non-significant, b = -18.05, $t(8)$ = 0.30, *ns*., in the negative reputation condition. The positive reputation condition is the only condition examined in this study in which acquiring good reputations strongly help a seller to achieve better profit.

The expected asymmetry in the directionality of reputation clearly emerged, as shown in the differences in the frequencies of identity changes in the positive reputation system and the negative reputation system. In the positive reputation condition, participants changed their identities 1.75 times on average, and in the negative reputation condition, participants changed their identities 10.17. This difference was statistically significant, $t(12.3)$ = 3.05 $p < .01$. Interestingly, in the negative reputation condition, the correlation between reputation score and

frequency of identity changes was positive, although not significant, $r = .41$, *ns*. This result seems to suggest that frequent identity changers' reputation score stays close to zero, whereas infrequent identity changers accumulated negative reputations. The same correlation was negative in the positive reputation condition, $r = -.56$, $p < .05$. This result seems to suggest that those who have acquired high reputation scores maintained their good reputations.

The correlation between the fraudulence level and frequency of identity changes was very high in the negative reputation condition, $r = .88$, $p < .001$, whereas the same correlation was positive and modest in the positive reputation system, $r = .40$, $p < .08$. This suggests that, as important as "brand names" are in the positive reputation condition, it might even be more important in the negative reputation condition. In the latter, where the highest reputation that a seller can acquire is 0, a buyer could not distinguish between sellers who have a 0 reputation score because they, through their honesty, have been able to avoid any negative evaluations, and sellers who have a 0 reputation score because they have recently acquired a new identity and cancelled (very likely) highly negative reputations. In such cases, name recognition can serve as an indicator of seller's positive reputation, because buyers are likely to have seen the handles of the former type of sellers with 0 reputation scores, but not the latter type.

Finally, the correlation between frequency of identity changes and profit as a seller in the positive reputation condition was negative and non-significant, $r = -.34$, *ns*., confirming our expectation that participants have no incentives to change their identities. The correlation between frequency of identity changes and profit from selling commodities in the negative reputation condition was positive, though not significant, $r = .41$, *ns*.

Honesty as a Seller and Honesty as an Evaluator

Two types of honesty are involved in the experiments reported above. The first type of honesty is the seller's honesty when advertising the quality of the commodity, captured by the

seller's fraudulence level. The second type of honesty is the buyer's honesty in giving an evaluation score of the seller, captured by the index of honest evaluation. Interestingly, the correlation between these two types of honesty was not very high (r = -.27, *ns*., in the 45 minute session of the reputation condition in the first experiment; -.20, *ns*., in the 20 minute session of the reputation condition in the first experiment; -.42, $p < .05$, in the second experiment; -.36, $p < .10$, in the negative reputation condition in the third experiment; -.20, *ns*., in the positive reputation condition in the third experiment). While each participant acted both as a seller and a buyer, how honestly she behaved as a seller was not strongly related to how honestly or diligently she behaved as a buyer. Although not directly related to the issues addressed in this study, the weak correlation between the two types of honesty raises a very interesting question concerning the validity of the critical assumption underlying some of the solutions to the second-order social dilemma problem discussed in the introduction. As Yamagishi and Takahashi (1994) point out, the behavioral consistency between a higher-level and a lower-level social dilemma is critical for voluntarily provided sanctions to solve the problem of social dilemmas, while no currently available theoretical model explains the adaptive advantage of the consistency. The weak correlation in our study suggest that this across-level consistency may not be as common as assumed by some theorists (e.g., Axelrod, 1986).

General Discussion

All of the six hypotheses addressed in this study received support from the experimental results. The messages we receive from the three experiments presented in this paper are clear. First, information asymmetry drives the experimental market among anonymous traders into a lemons market in which only the lowest quality goods are traded. In this market, opportunities to achieve better profits from trading high quality goods are unrealized. Second, either

experience-based information or reputation about other traders alleviates the lemons problem when traders' identities are permanent. Third, the power of experience-based information or reputation as a solution to the problem of lemons is substantially reduced when traders can freely change their identities and cancel their reputations. Fourth, the negative reputation system that is designed to illuminate dishonest traders is particularly vulnerable to identity changes, whereas the positive reputation system designed to illuminate honest traders is not so vulnerable to identity changes.

Despite the theoretical problem of free riding in the provision of experience-based personal information to be shared with others, the overall picture emerging from the experimental findings is that reputation systems provide decent success in curtailing the lemons problem. So, how was the free rider problem resolved in our experiment? One answer to this question points to a feature of the experimental design. All participants were *required* to submit an evaluation score each time they made a purchase. Free riding was made impossible by the design of the experiment. However, participants did not need to report accurate evaluations. Because of this, free riding concerning submission of an ***accurate*** evaluation score was in fact possible. Furthermore, there were incentives to submit a false evaluation score. Lowering other traders' reputation by attaching them negative or low evaluation scores increased a trader's own relative standing as a seller in the market. The results of the experiments indicate that most participants honestly and accurately report their evaluations of the sellers from whom they purchased commodities.

In an effort to design better reputation systems to resolve the lemons problem, our findings concerning the positive and the negative reputation systems provide valuable insights. Previous research on Internet auction sites have pointed out stronger effects of negative reputation than positive reputation (e.g., Eaton 2002; Houser and Wooders 2001; Lucking-Reiley et al. 2000;

Resnick and Zeckhauser in press; Standifird 2001). Standifird (2001) attributes the stronger effect of negative reputation to the gain-loss asymmetry in subjective utility (Kahneman and Tversky 1979). At the same time, negative reputations are used rarely. For example, Resnick and Zeckhauser (in press) conclude, based on 36,233 randomly chosen transactions in eBay, that less than one percent of ratings submitted by buyers are negative. According to Resnick and Zeckhauser, fear of retaliation as well as a courtesy helps explain why Internet traders use more-effective negative reputations less often than less-effective positive reputations. Internet traders provide positive evaluations to their trading partners to elicit similar positive reputations from them. Internet traders also refrain from providing negative evaluations to their trading partners to avoid retaliatory negative evaluations from them. Dellarocas (2003) suggests that the paucity of negative ratings in eBay reflects actual lack of dissatisfying experience since the reputation system help eBay traders to achieve an equilibrium in which cheating rarely occurs. The results of the third experiment, however, suggest a different reason why negative reputations are used less often. ***Negative reputations may be more effective than positive reputations in the short run, but can be less effective in the long run***. In the third experiment, the average quality of commodities was higher in the negative reputation condition than in the positive reputation condition *in the first time block*. This initial strong effect of negative reputation may reflect the strong aversion of negative outcomes suggested by Standifird (2001). However, the advantage of the negative reputation system did not last long. Soon, the average quality in the negative reputation condition plummeted whereas the average quality in the positive reputation condition steadily improved. The positive effect of the negative reputation system was short-lived because it was neutralized by frequent identity changes. The positive effect of positive reputations, on the other hand, may take time to be realized, but it works in a cumulative manner. In part, such a situation occurs because identity changes are rare in the

positive reputation system. The differential long-term effects of positive and negative reputation systems, in addition to the psychological factors discussed by Resnick and Zeckhauser (in press), may be behind the differential frequencies of the two types of reputations.

These findings have important implications on in the roles of positive and negative reputations in the two types of social relations we discussed in the introduction. The success of the Maghribi coalition, according to Greif's (1989) analysis, lies in the fear of exclusion. A Maghribi coalition member balances the immediate profit of behaving dishonestly with the risk of being excluded from the coalition (and from future profit) once his dishonesty is detected and his reputation of dishonesty spreads among coalition members. Thus, it was predominantly ***negative reputation*** that was critical to the success of the coalition. Greif (1989) further argues that the success of negative reputation requires closure of the coalition. Exclusion from a market matters only when it is closed to non-members. This issue speaks to why negative reputation was not very effective in curtailing the lemons problem in the third experiment. The individuals who have acquired bad reputations and been "excluded" or shunned by the other members can freely "re-enter" the same market under a new identity. The central characteristic of online trading—that is, its openness—thus prevents negative reputation from exerting its power. The openness of online trading, on the other hand, promotes positive reputation as an effective means for curtailing the lemons problem. To understand why positive reputation is useful in online trading, we need to realize that there are two functions in reputation: exclusion and inclusion. The power of negative reputations is based on the ***principle of exclusion***. Negative reputations exclude dishonest traders from the market. In contrast, the power of positive reputations is based on the ***principle of inclusion***. Positive reputations are not effective for excluding dishonest traders from the market. However, positive reputations are a useful means to attract potential trading partners. Attracting new partners is not of central importance in a

closed market since the membership of the market is limited. Establishing a good reputation in the Maghribi coalition does not help a Maghribi trader to expand his deals beyond the boundaries of the coalition. In sharp contrast to this, the number of potential trading partners is unlimited in an open market or in online trading. Therefore, the merit of obtaining good reputations in an open market is also unlimited. Established brand names are much more valuable, in this sense, in open markets than in closed markets.

The distinction between exclusion based on negative reputation and inclusion based on positive reputation may explain the differential effects of the seller's reputation on her profit in the reputation condition of the first experiment and in the positive reputation condition of the third experiment,[7] the two conditions in which strong group-level effect of having a reputation system was observed. Despite the strong group-level effect of the reputation system, sellers with good reputations did not enjoy much advantage in making profit in the first experiment. That is, in the first experiment in which identities are fixed, the positive effect of the reputation system at the group level is not an aggregation of the positive effect of having good reputation at the individual level. This result challenges the conventional economist's logic that "in equilibrium, a good reputation must command a price premium" (Resnick et al., 2002: 4). This assertion is based on the truism that aggregation of individuals' rational behavior sustains an equilibrium. Upon careful examination, however, this kind of equilibrium is found to be untenable in the online market. This is because individual sellers have no means to find out about the price premium of maintaining good reputations, as discussed by Macy (2003). One way the sellers find out about the price premium of reputation is to read journal articles and conference papers.

[7] Individual differences in reputation did not have an effect on the seller's profit in the second experiment. We exclude second experiment from the following discussion since the group-level effect of the reputation system was not strong there.

The findings concerning price premium have been scarce so far, and these findings are not consistent (see Resnick et al. 2002, for a review). Another way for the sellers to find out the price premium of reputation is by experimenting for themselves, behaving honestly and dishonestly under different pseudonyms. After experimenting for themselves, they may choose the optimum level of honesty and reputation. The participants of the first experiment in the reputation condition certainly could not do this, however, because experiment on bad reputation is too costly. This strongly suggests that the strong group-level effect of reputation in that condition is not an aggregation of individual sellers' rational decisions.

If the group-level effect of reputation observed in the reputation condition of the first experiment in the absence of individual-level advantage is not an aggregation of individuals' rational decisions, what does explain it? Macy (2003) suggests that it is the belief shared by the market participants that reputations do work. The key point here is that, as discussed above, those who believe in the sanctioning power of reputation have no opportunity to personally experience the validity of the belief. The belief thus can be maintained among market participants, and whether having good or bad reputations actually affects their profit is irrelevant. The trend over time of the fraudulence level in the reputation condition shown in Figure 3 further adds support to the claim that the group-level effect of reputation in this condition was based on pre-existing belief, rather than experience. The average level of fraudulence as well as the average quality of commodities produced for sale in the market stayed at about the same level throughout the experimental session as in the first 10 minutes. The effect of having a reputation system existed from the very beginning of the first experiment.

The pattern is quite different under the positive reputation system in the third experiment. First, individual differences in the level of reputation are positively correlated with the profit they earned as sellers. Clearly, sellers were rewarded by maintaining good reputations. Second,

the average level of fraudulence decreased over time as shown in Figure 3, and the average level of the qualities of the commodities increased over time as shown in Figure 5. These differences in how a reputation system improves the market in closed and open markets present an interesting paradox. In the closed market in which the possibility of being excluded functions as a strong sanction against dishonest behavior, traders do not have personal means to test the power of reputation as a sanctioning mechanism. In an open market in which those who are excluded can re-enter with a new identity, it is not as costly to personally examine the benefit (or cost) of having good (or bad) reputations as in the closed market. This may explain why the reputation system has a positive effect at the group level but not at the individual level in the first experiment in which identities of the traders are fixed, on the one hand, and the positive reputation system in the third experiment had a positive effect at both levels. The analysis presented above is thus suggestive of the mechanism through which positive and negative reputation improve the quality of the commodities produced and sold in open and closed markets—through simple belief or through personal experience. However, identifying the exact mechanisms in the two types of social relations definitely require further investigation.

In conclusion, we started with the assertion that the closed nature of the coalition was a critical condition for the negative reputation system to control the lemons problem among Maghribi traders (Greif 1989). Based on the findings from our experimental study, we propose that the lack of a closed market among online traders, which appears, at first glance, to be a formidable problem, may actually be a blessing. To realize why positive reputations, rather than negative reputations, are effective in resolving the lemons problem in our experimental market as efficiently as in the real Internet auction sites (Kollock 1999), we must turn to the principle of inclusion rather than the principle of exclusion. Put differently, although the principle of exclusion has been the central principle adhered to when examining the effect of reputations in

markets, we find that the effect of reputations in *open* markets is best understood using the principle of inclusion.

References

Akerlof, George. 1970. “The Market for Lemons: Quality Uncertainty and the Market Mechanisms.” *The Quarterly Journal of Economics* 84: 488-500.

Avery, Christopher, Paul Resnick, and Richard Zeckhauser. 1999. “The Market for Evaluation.” *American Economic Review* 89: 564-584.

Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.

Axelrod, Robert. 1986. “An Evolutionary Approach to Norms.” *American Political Science Review* 80: 1095-1111.

Blurton Jones, N. G. 1984. “A Selfish Origin for Human Food Sharing: Tolerated Theft.” *Ethology and Sociobiology* 5: 1-3.

Blurton Jones, N. G. 1987. “Tolerated Theft, Suggestions About the Ecology and Evolution of Sharing, Hoarding, and Scrounging.” *Social Science Information* 26: 31-54.

DeJong, Douglas V., Robert Forsythe, and Russell J. Lundholm. 1985. “Ripoffs, Lemons, and Reputation Formation in Agency Relationships: A Laboratory Market Study.” *The Journal of Finance* 40: 809-820.

DeJong, Douglas V., Robert Forsythe, Russell J. Lundholm, and Wilfred C. Uecker. 1985. “A Laboratory Investigation of the Moral Hazard Problem in an Agency Relationship.” *Journal of Accounting Research* 23: 81-120.

Dellarocas, Chrysanthos. 2003. “Efficiency and Robustness of Binary Feedback Mechanisms in Trading Environments with Moral Hazard.” MIT Sloan School of Management Working Paper 4297-03.

Eaton, David H. 2002. “Valuing Information: Evidence from Guitar Auctions in eBay.” Working paper, Department of Economics and Finance, Murray State University.

Fehr, Earnst and Simon Gachter. 2000. “Fairness and Retaliation: The Economics of Reciprocity.” *Journal of Economic Perspective* 14: 159-181.

Fehr, Earnst, Simon Gachter, and Georg Kirchsteiger. 1997. Reciprocity as a Contract Enforcement Device: Experimental Evidence. *Econometrica* 65: 833-860.

Friedman, Eric, and Paul Resnick. 2001. “The Social Cost of Cheap Pseudonyms.” *Journal of Economic and Management Strategies* 10: 173-199.

Gintis, Herbert. 2003. “Solving the Puzzle of Prosociality.” *Rationality and Society*, in press.

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

Greif, Avner. 1993. “Contract Enforceability and Economic Institutions in Early Trade: The Maghribi Traders Coalition.” *American Economic Review* 83: 525-548.

Houser, Daniel and John Wooders. 2001. “Reputation in Auctions: Theory, and Evidence from eBay.” Working paper, Department of Economics, University of Arizona.

Kaplan, Hillard and Kim Hill. 1985. “Food Sharing Among Ache Foragers: Test of Explanatory Hypotheses.” *Current Anthropology* 26: 223-246.

Kahneman, Daniel and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decisions under Risk.” *Econometrica* 47: 263-291.

Kollock, Peter 1994. “The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust. *American Journal of Sociology* 100: 313-345.

Kollock, Peter 1999. “The Production of Trust in Online Markets.” Pp. 99-123 in *Advances in Group Processes*, Vol. 16, edited by E. J. Lawler, M. Macy, S. Thyne, and H. A. Walker. Greenwich, CT: JAI Press.

Lucking-Reiley, David, Dong Bryan, Naghi Prasad, and Daniel Reeves. 2000. “Pennies from eBay: The Determinants of Price in Online Auctions.” Working paper, Department of

Economic and Business Administration, Vanderbilt University.

Macy, Michael W. 2003. “Reputations as Signals and Sanctions.” Presented at the First Interdisciplinary Symposium on Online Reputation Mechanisms, MIT, April 26-27.

Milinski, Manfred, Dirk Semmann and Hans-JuÈrgen Krambeck. 2002. “Reputation Helps Solve the ‘Tragedy of the Commons’.” *Nature* 415: 424-426

Nowak, Martin A. and Karl Sigmund. 1998a. “Evolution of Indirect Reciprocity by Image Scoring.” *Nature* 393: 573-577.

Nowak, Martin A. and Karl Sigmund. 1998b. “The Dynamics of Indirect Reciprocity.” *Journal of Theoretical Biology* 194: 561-574.

Okazaki, Tetsuji. 1999. *Market Economy during Edo Period: Kabunakama in the Perspective of Comparative Institutional History*. Tokyo: Kodansha. (In Japanese)

Oliver, Pamela. 1980. Rewards and punishments as selective incentives for collective action: Theoretical investigations. *American Journal of Sociology* 85: 1356-1375.

Olson, Mancur. 1965. *The Logic of Collective Action*: *Public Goods and the Theory of Groups.* Cambridge, MA: Harvard University Press.

Resnick, Paul, and Richard Zeckhauser. In press. “Trust among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System.” In *Advances in Applied Microeconomics,* Volume 11*: The economics of Internet and E-commerce,* edited by M. R. Bay. Amsterdam: Elsevier Science.

Resnick, Paul, Richard Zeckhauser, Eric Friedman, and Ko Kuwabara. 2000. “Reputation Systems: Facilitating Trust in Internet Interactions.” *Communications of ACM* 43: 45-48.

Resnick, Paul, Richard Zeckhauser, John Swanson, and Kate Lockwood. 2002. “The Value of Reputation on eBay: A Controlled Experiment.” Working Paper for the June 2002 ESA Conference, Boston, MA

Standifird, Stephen S. 2001. “Reputation and E-commerce: eBay Auctions and the Asymmetrical Impact of Positive and Negative Ratings.” *Journal of Management* 27:

279-295.

Wilson, R. 1985. “Reputations in Games and Markets.” Pp. 27-62 in *Game-Theoretic Models of Bargaining*, edited by A. Roth, Cambridge: Cambridge University Press.

Yamagishi, Toshio. 1986. “The Provision of a Sanctioning System as a Public Good.” *Journal of Personality and Social Psychology* 51: 110-116.

Yamagishi, Toshio and Nobuyuki Takahashi. 1994. “Evolution of Norms Without Metanorms.” Pp. 311-326 in *Social Dilemmas and Cooperation*, edited by U. Schulz, W. Albers, and U. Mueller. Berlin: Springer-Verlag.

Yamagishi, Toshio, Noriaki Yoshikai, and Hiroyuki Takahashi. 2003. “A Web-based Experiment of Internet Auctioning: How Reputation Works.” Presented at the First Interdisciplinary Symposium on Online Reputation Mechanisms, MIT, April 26-27.

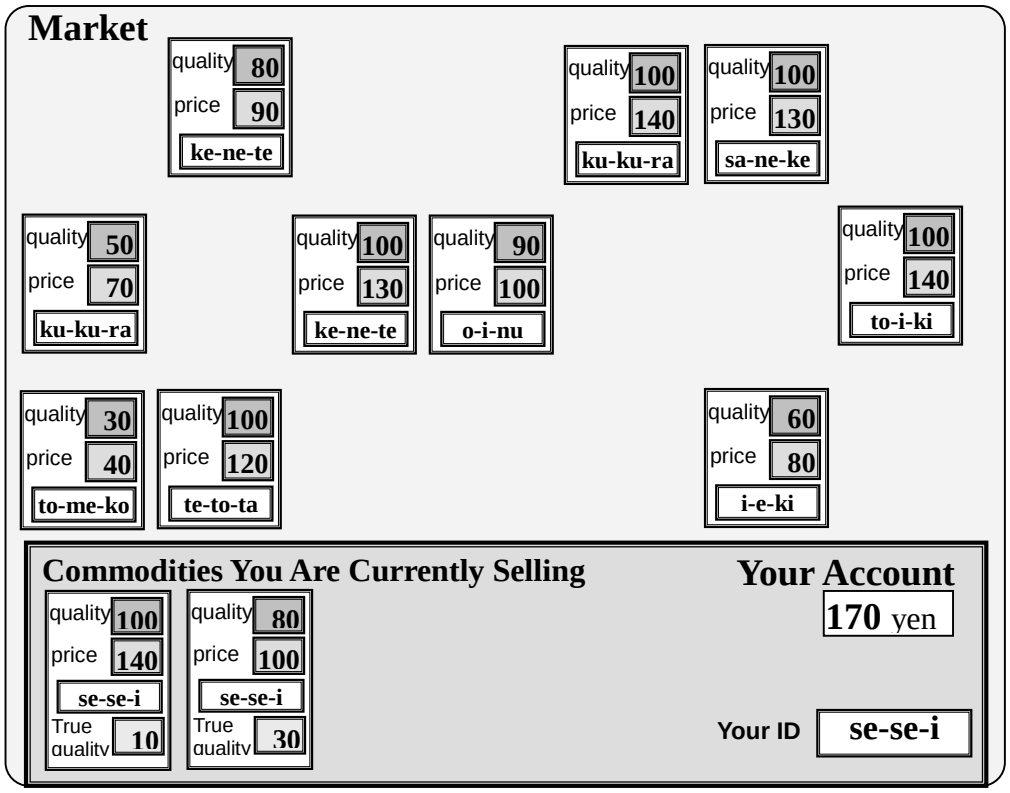

Figure 1 An image of the computer monitor of the identity condition of the first experiment. All characters are translated from Japanese.

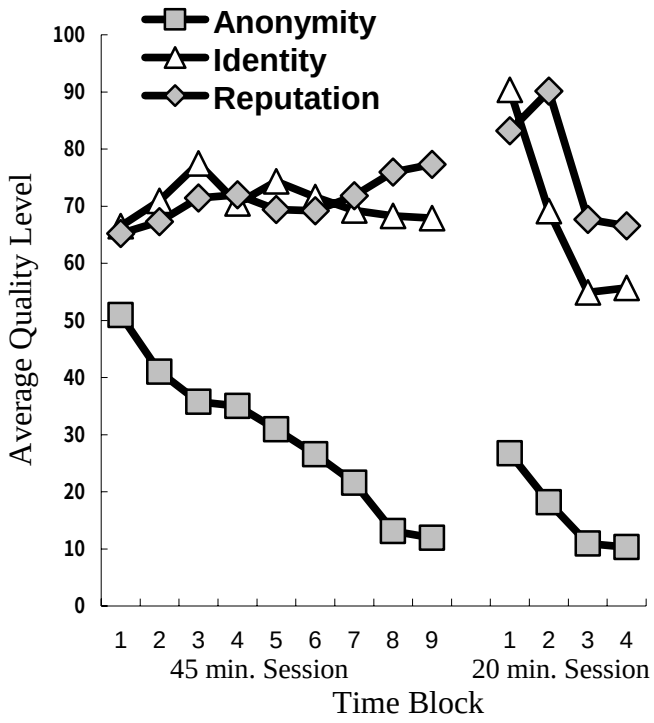

Figure 2 Average quality of the produced commodities in Experiment 1.

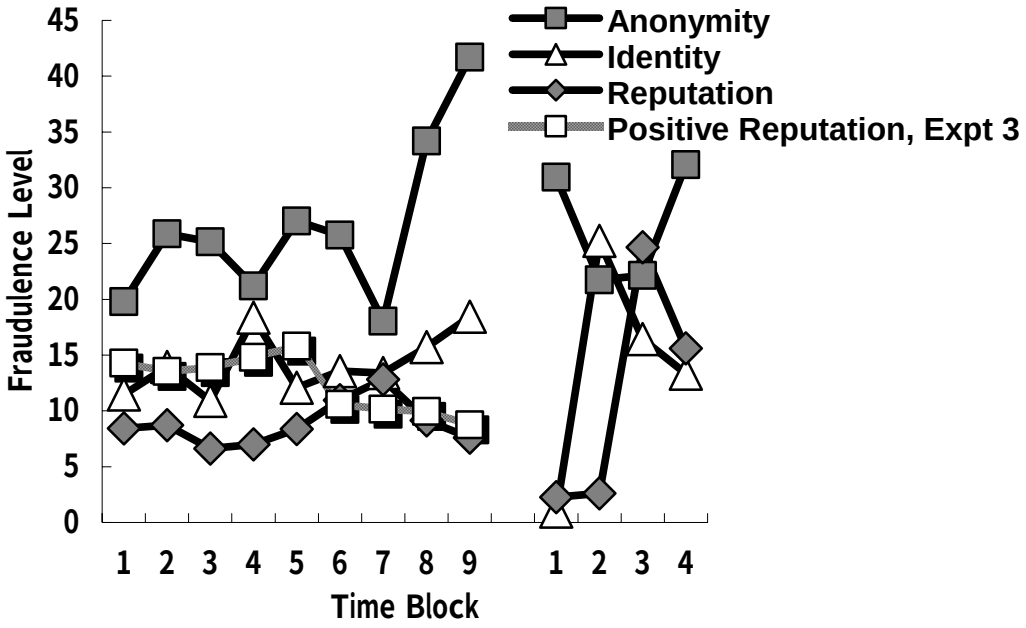

Figure 3 Level of dishonesty (the difference between the advertised quality and the true quality) in the first experiment

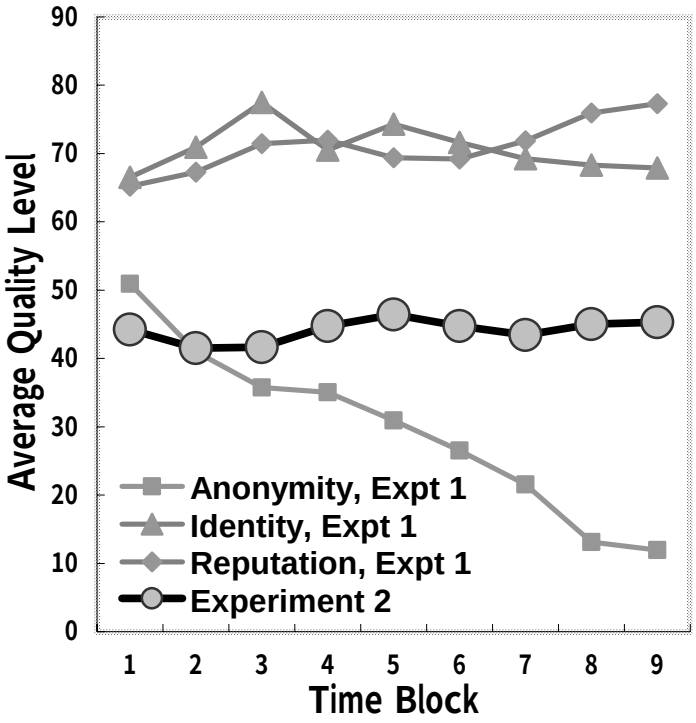

Figure 4 Average quality in Experiment 2, compared to the levels in the three conditions in Experiment 1

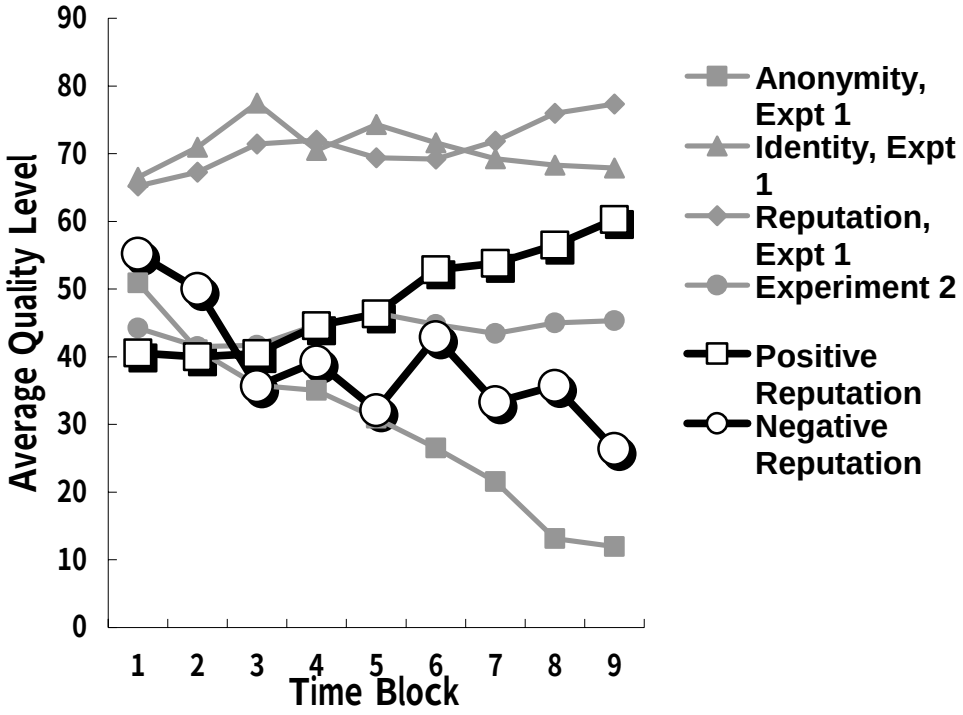

Figure 5 Average quality in the positive reputation condition and the negative reputation condition in Experiment 2, compared to the levels in Experiments 1 and 2.