Provision, Interpretation and Effects of Feedback in Reputation Systems

by

Tapan A. Khopkar

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
in The University of Michigan
2008

Doctoral Committee:

Professor Paul J. Resnick, Chair
Professor Yan Chen
Professor Judith S. Olson
Assistant Research Scientist Mary L. Rigdon
Associate Professor Chrysanthos Dellarocas, University of Maryland

Acknowledgements

This dissertation bears my name, and the primary responsibility for any shortcomings is mine. But this dissertation would not have been possible without the support from many individuals. If I were to thank everyone individually, it would take about as much time as writing this dissertation took. For the sake of brevity I can thank only a few people here by name, but there are many others that I am grateful to.

First, I thank all of my dissertation committee members and in particular the chair and my advisor, Paul Resnick. His support, guidance and feedback have been invaluable in carrying out the work described in this dissertation and also in learning the craft of research. He has helped keep me on track throughout this process, and helped clear obstacles that at times seemed insurmountable. Yan Chen and Judy Olson have been sources of mentoring and advice throughout my graduate student years. I have learnt a great deal from them, in particular about designing and conducting laboratory experiments. They have also been my role models and sources of inspiration for their dedication and efficiency. I would also like to thank Mary Rigdon and Chris Dellarocas for their time and valuable feedback on my work.

I also thank my coauthors- Robert Gazzale in Chapter 2, Paul Resnick in Chapters 3 & 4, and Xin Li in Chapter 3. Their help and feedback have been instrumental in bringing this work to fruition. Special mention must be made of Kan Takeuchi, who helped with z-Tree programming, and Ben Taylor, who helped with recruitment and administration for the experiments. NSF and Williams College provided financial support for the experiments, while eBay provided the dataset used for analysis in Chapter 3. I would like to thank Jeff Mackie-Mason for his advice and support, and the ICD lab group for the invaluable feedback on my work from time to time. I also wish to thank Sue Schuon who has been a constant and invaluable source of support and assistance.

Special thanks are in order for my friends and fellow graduate students for their time, patience and help in innumerable forms- ranging from pain relief patches to valuable feedback on my work. This journey would not have been possible without them. Mihir Mahajan, Lian Jian and Anya Osepayshvili deserve special mention, but I also wish to thank Amit Salvi, Anil Nayak, Derek Hansen, Devayani Bhave, Munish-Shilpa and Rhea Gupta, Nagaraj, Nandita Madhavan, Neha Jain, Rekha Santhanam, Rucha Vaidya and AJ, Shrihari Sathe, Swati Raul and Vaibhav Donde.

Finally, this long journey would not have been possible without the unflinching support from my family. There were times when I would have given up, but they egged me to push on, and provided me the wherewithal to do so. Arun, Nandu, Abha-Shreyas-Chochi, Pallavi, Shilpa, Shonil, Anil-Shobha, Vasanti, Shanta, Neeta have all played important roles in shaping my dreams, and giving me the strength to pursue them. I must make special mention of my grandfather, R.S. Bhagwat, who instilled in me the desire for learning. Amidst a sea of family, friends and well-wishers, three people stand out- my grandmother- Kamal Bhagwat (aaji), my mother- Neela Bhagwat, and my fiancée-Preeti Pansare. The three of them have provided the ballast in my life, whatever I have achieved so far is thanks to them. Dedicating this dissertation to Aaji, Neela and Preeti is necessary, but hardly sufficient.

Table of Contents

List of Figures

List of Tables

List of Appendices

Chapter 1

Introduction

1.1 Trust in Electronic Marketplaces

Every day, millions of people engage in commercial transactions over the Internet, using portals such as Amazon and eBay. The revolutionary feature of these electronic marketplaces is that they enable traders to transcend the limits of geographic distance by providing them with a platform to exchange goods and services without ever meeting their interaction partners (Ba and Pavlou 2002). However, there is a trade-off.

The increased opportunity for trade comes with a significant increase in the risk involved. Buyers and sellers in an electronic market are likely to be complete strangers and are separated by physical distance. This leads to an information asymmetry between them. The seller often has a much better idea about the true value of the product than the buyer does. The buyer pays first and has no way to enforce that the seller fulfills the transaction as agreed. The seller can act in a number of ways that are detrimental to the buyer's welfare. The most extreme scenario is one where the seller receives payment and does not ship the product at all. There are several other possibilities such as intentional misrepresentation of product quality, or use of inferior shipping material, that can lead to an unsatisfactory transaction for the buyer. The common feature in these scenarios is that the seller's and the buyer's incentives are misaligned. The seller can make a choice that benefits himself[1], while imposing a cost on the buyer. Moreover the buyer cannot exert any control over the seller's choice.

Social uncertainty is a useful framework to understand the situation faced by the buyer in an electronic marketplace. Social uncertainty is said to exist for an actor, when her interaction partner has an incentive to act in a manner that imposes a cost on her, and she has neither the means to control, nor the information to predict, the partner's behavior (Yamagishi, Jin et al. 1998). A buyer in an electronic marketplace as characterized above faces such social uncertainty.

Trust plays a key role in facilitating interactions in situations of social uncertainty. 'Trust' is defined as the confident expectation that the other party will perform a transaction according to the trustor's expectations (Gambetta 1988). A buyer's decision to buy in an electronic marketplace is an act of trust. The buyer is placing her trust in the seller, because she expects the seller to fulfill the responsibilities in the transaction as agreed upon. A seller who reciprocates the buyer's trust by completing the transaction as agreed is being trustworthy. A seller who chooses an action that benefits himself, at a cost to the buyer can be considered as being untrustworthy. Having thus classified the seller and buyer's actions, we can use the trust game (Berg, Dickhaut et al. 1995) to understand the contradictory incentives faced by sellers and buyers in an electronic market.

---

[1] For the sake of clarity, I will use a female pronoun to refer to the buyer, and a male pronoun to refer to the seller

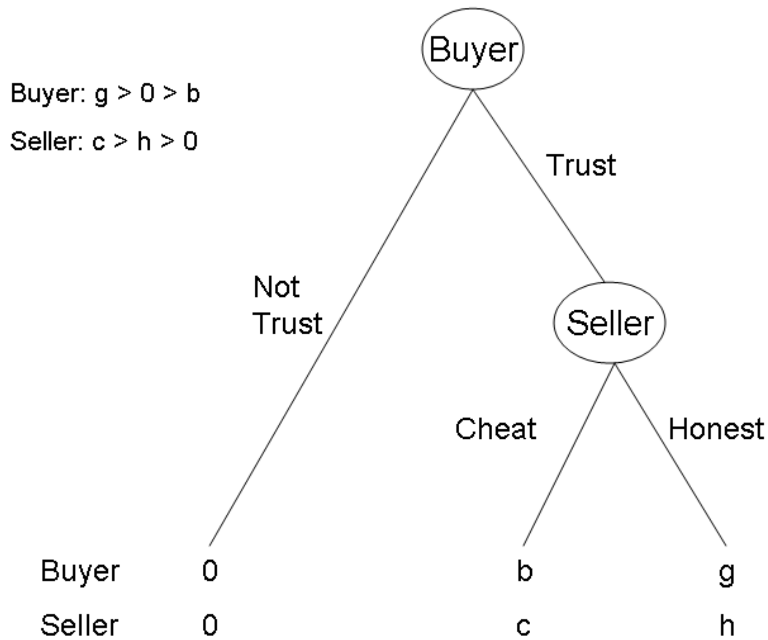

Figure 1 Trust Game Representation of a Transaction in an Electronic Marketplace

When a buyer encounters a seller in an electronic marketplace, she can choose to 'Trust' or 'Not Trust'. If the buyer chooses to trust, the seller can choose between 'Honest' and 'Cheat'. The players' actions and consequent payoffs are depicted in Figure 1.If the seller is honest the buyer gets a good outcome, while if the seller cheats the buyer gets a bad outcome. If the seller is trustworthy, the payoffs to the (buyer, seller) are (g,h) respectively. If the seller cheats, the respective payoffs for (buyer, seller) are (b,c). On the other hand, if the buyer decides to not trust, the payoffs for both parties are zero, since no transaction takes place. The following inequalities characterize the misaligned incentives for the buyer and the seller.

For the buyer: $g > 0 > b$

For the seller: $c > h > 0$

If the buyer trusts and the seller is trustworthy, a mutually beneficial transaction transpires between the buyer and the seller. However, in a one-shot transaction, once the buyer trusts the seller, the seller has no incentive to be trustworthy. Buyers who are aware of this have no reason to trust the seller. Thus, a mutually beneficial transaction would not take place if the buyer and the seller don't expect to interact with one another in the future.

The situation is quite different if the buyer and the seller expect interact with each other frequently. In such a scenario, the 'shadow of future' (Axelrod 1984) provides the seller with an incentive to be trustworthy in the present. If the seller cheats today, the buyer can punish him in the future by choosing to not trust him. The Folk theorem suggests (Fudenberg and Maskin 1986) that if the seller is sufficiently patient , it will be an equilibrium for the buyer to trust, and the seller to be trustworthy, as long as they expect to interact in the future.

Unfortunately, this condition is not satisfied in an electronic marketplace. Most transactions are between users who have never interacted with each other earlier, and will likely never interact in future[2]. Also, in most electronic marketplaces, users are identified only by means of a freely available pseudonym. As a result users can easily assume a new identity. (Resnick, Zeckhauser et al. 2000)

It is clear that the folk theorem does not apply in electronic markets due to the lack of repeated interactions and persistent identities. With the 'shadow of the future'

---

[2] Resnick and Zeckhauser (2002) found that, during a five month period from February 1999 to June 1999, 89% of the transactions were between users who interacted only once.

absent, we would be relegated to a single-shot world, where no transaction would take place in equilibrium, and the market would collapse.

In reality, the situation is not as grim as predicted by game theory. Bolton et al. (Bolton, Katok et al. 2004) found positive (although low) levels of trust and trustworthiness in a simulated electronic market implemented using a trust game. Such outcomes are usually attributed to factors other than rational self interest, such as altruism. While such outcomes are of interest to researchers, low levels of trustworthiness stemming purely from altruism would not be acceptable in commercial marketplaces. Many electronic marketplaces use a reputation system to foster trust and trustworthiness among their participants and to protect the market.

1.2 Reputation Systems in Electronic Marketplaces

A reputation system collects information about the participants' past behavior, aggregates this information, and makes it available to other participants. We have seen earlier that electronic markets don't achieve folk theorem like efficiency due to the lack of repeated interactions and persistent identities. A reputation system attempts to redress both these issues. Even though a user can be identified only by a pseudonym, the reputation system reveals the history of his interactions with other users. Now, even if a buyer does not know the identity of the seller she is dealing with, she knows how he has behaved in his past transactions. Similarly, the seller knows that even if he does not encounter this particular buyer again, potential buyers in the future can base their purchase decisions on the outcome of this transaction (Falk and Fischbacher 2006). A reputation system performs the following functions: (Resnick, Zeckhauser et al. 2000)

Signaling: A reputation system enables the buyer to distinguish between trustworthy and untrustworthy sellers. A seller's reputation acts as a signal of his skill and integrity. Buyers can use this information to distinguish dishonest and incompetent sellers from good ones, thereby overcoming adverse selection.

Sanctioning: A reputation system reveals the seller's past actions to future buyers. Sellers, aware that their actions are being reported and that future buyers can discriminate against sellers with a bad reputation, have an incentive to act in a trustworthy manner. The reputation system thus helps the sellers overcome moral hazard.

Self-selection: Another consequence of the reputation system is that sellers, who are incompetent or dishonest, acquire a bad reputation and are not able to attract buyers. Eventually, such sellers will be driven out of the market. Thus, the reputation system helps limit adverse selection.

Thus a reputation system fosters trust and trustworthiness by restoring the 'shadow of the future' in a transaction. Using the framework of trust, a boundedly rational individual should extend her trust, only if she thinks there is a sufficiently high chance of reciprocation (in this context, trustworthy behavior) based on the trustee's reputation (Ostrom and Walker 2003). Coleman (Coleman 1990) defines reputation as the information associated with a particular potential interactant concerning the interactant's likelihood of carrying out a certain behavior in future. In an electronic market with a reputation system, the history of the seller's interactions acts as his reputation.

Figure 2 presents a schematic diagram for a transaction in an electronic market with a reputation system. The reputation system stores the reputation of every participant in the market. When deciding to trust (or not trust) a seller, the buyer uses the seller's reputation available from the reputation system. If the buyer decides to trust the seller, the seller then chooses to be trustworthy or not; this decision affects the transaction outcome for the buyer and the seller. Upon completion of the transaction, the seller's reputation is updated to include a report of the transaction outcome, and the updated reputation is stored by the reputation system.

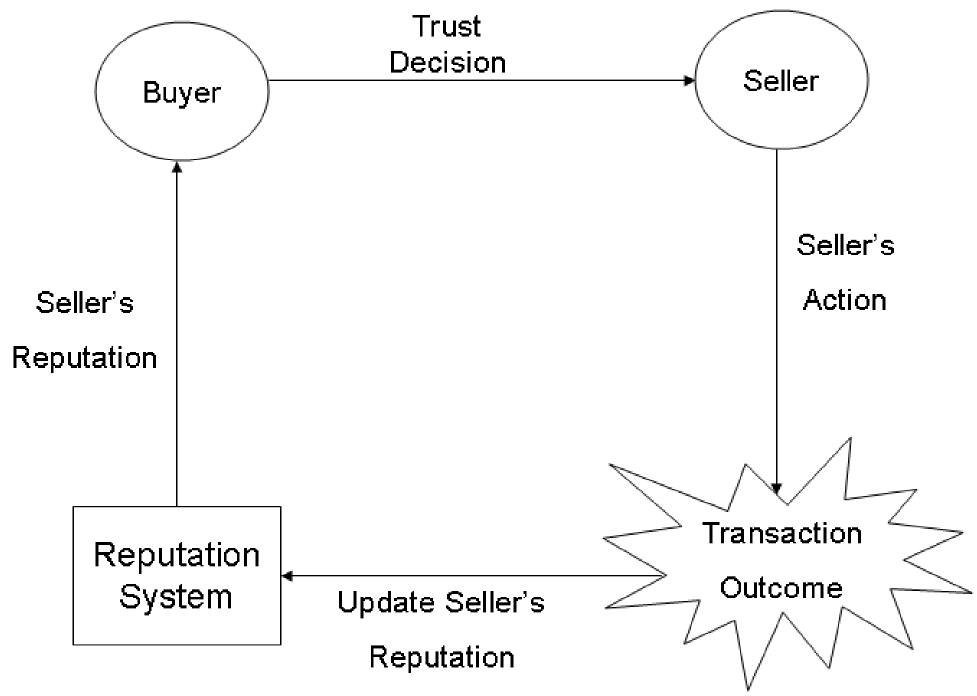

Figure 2 Reputation System in an Electronic Market

With aid from the reputation system, the buyer can distinguish between sellers based on their reputations, and trust only those sellers who have a reputation for being trustworthy. Sellers, aware that their prospects for trade depend on their reputation, will be careful to maintain a good reputation. This, in turn, will provide them with an incentive to be trustworthy in the current transaction. Thus, a reputation system can help traders achieve efficient folk-theorem like outcomes, where it is equilibrium behavior for the buyer to trust and a sufficiently patient seller to be trustworthy as long as he expects to remain in the market.

1.3 Feedback in Reputation Systems

Most reputation systems rely on voluntarily provided reports of transaction outcomes as indicators of the seller's choice of action. In the rest of my dissertation, I will use the term 'feedback' to denote a voluntarily provided report of a transaction outcome. Typically, a feedback in an electronic market consists of a numerical rating and a textual comment that describes the user's experience. On eBay's feedback forum, one of the most popular reputation systems, a feedback can be positive (+1), neutral (0) or negative (-1). A positive feedback indicates that the user was satisfied with the transaction outcome, a negative feedback indicates that the user was dissatisfied with the transaction outcome, while a neutral feedback indicates that the user was neither happy, nor unhappy with the transaction. Amazon.com uses an all positive star rating system, with a somewhat higher granularity. The feedback rating on Amazon varies from a single star to five stars, with higher number representing higher customer satisfaction.

Regulars in electronic marketplaces engage in thousands of transactions and get thousands of feedbacks. It would be an arduous and practically impossible task to browse all the comments and ferret out important information. To aid this process, reputation systems provide a few summary statistics of the user's history in the marketplace. Figure 3 is a screenshot of a feedback profile on eBay.

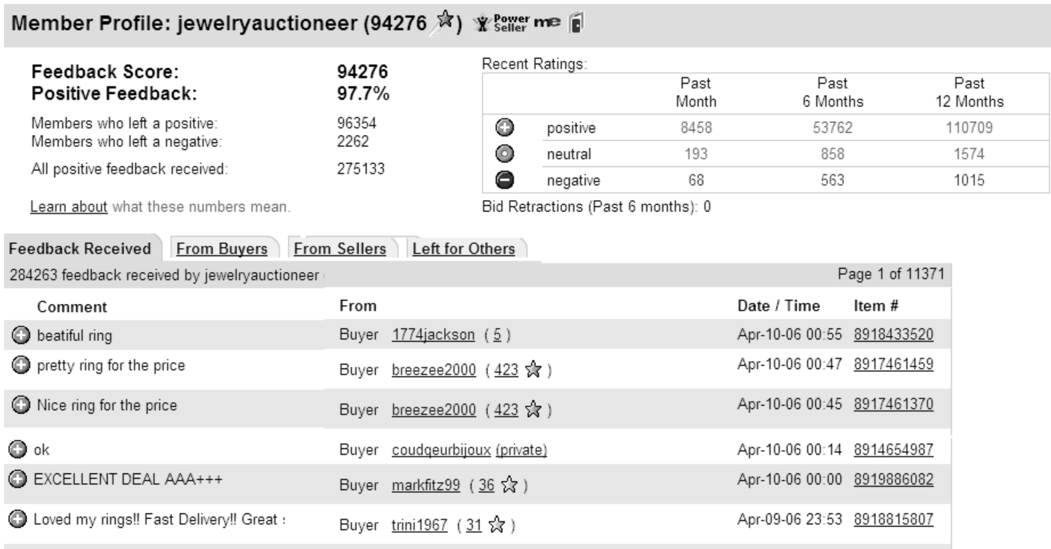

Figure 3 Feedback Profile on eBay

A user's feedback profile on eBay contains information such as her eBay ID, date of registration on eBay and a number of summary statistics such as feedback score, the percentage of positive feedback and a table summarizing recent feedback ratings. In addition to these, a list of all feedback received by a user is also presented on the user's feedback profile page.

A buyer who is interested in purchasing an item from a seller can access information about the seller's past dealings by looking at his feedback profile. The buyer can use this information to make an assessment of the seller's trustworthiness. The seller's trustworthiness, thus perceived, plays a role in the buyer's decision of whether or not to trust the seller (other factors involved are the buyer's risk attitude, the price of the item, buyer's value for the item etc.). If the buyer decides to trust the seller, she buyer can report the transaction outcome in the form of a feedback, which is used to update the seller's feedback profile. Figure 4 illustrates the various processes involving feedback in a reputation system.

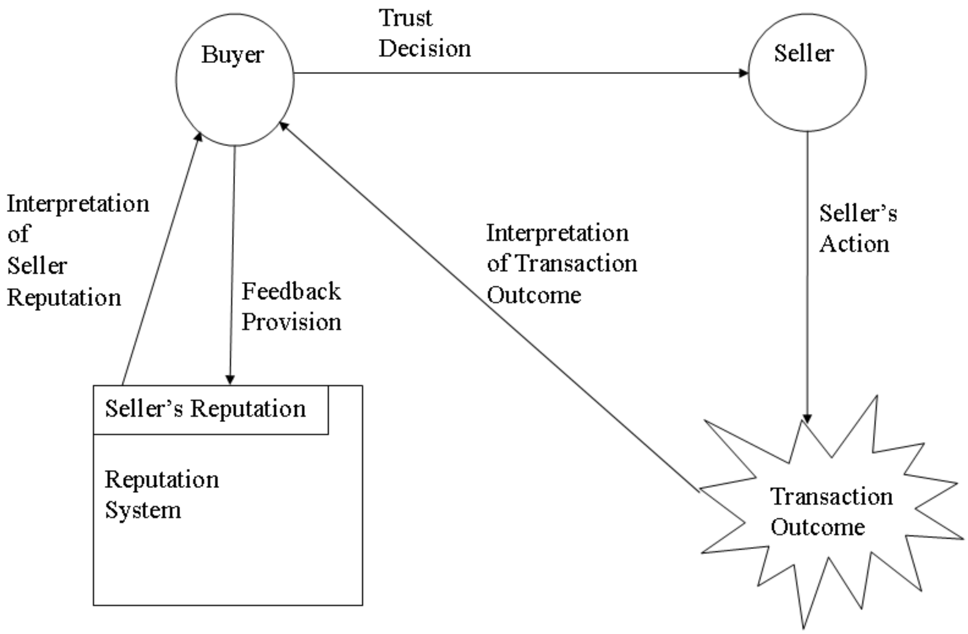

Figure 4 Feedback in a Reputation System

Feedback, thus, plays a crucial role in the functioning of a reputation system, by removing the disconnect between the present and the future. If feedback is complete and objective (i.e. outcome of every transaction is reported truthfully), then in theory, an electronic market with a reputation system can achieve the same high efficiency outcomes as a buyer and seller interacting repeatedly (Bolton, Katok et al. 2004). However, in commercial markets like eBay, Amazon, feedback is neither complete, nor objective. Feedback is voluntarily provided by users, and the processes that determine the provision and the content of feedback are often highly subjective.

A consequence of the voluntary nature of feedback is that not all transactions get a feedback. Resnick and Zeckhauser found that in a six month trading period, buyers on eBay provided feedback to sellers in about 50.2% of the transactions (Resnick and Zeckhauser 2002). Amazon marketplace reports even lower rates of feedback provision, about 10-20%. For transactions that don't get any feedback, the buyers (or anyone else) have no way of knowing whether the buyers were satisfied with the transaction outcome or not. Another serious problem is the lack of knowledge about whether the transactions that get a feedback are a representative sample of all transactions or whether they are biased in some way. Empirical research by Dellarocas and Wood (Dellarocas and Wood 2007) estimates that eBay buyers leave a transaction satisfied 78.9% of the time, mildly dissatisfied 20.4% of the time and very dissatisfied 0.7% of the time. These numbers are different from the 99% positive and 1% neutral/negative among the reported feedback, and suggest a bias towards reporting favorable outcomes.

The implications of the subjective nature of feedback are manifold. The provision of feedback by a buyer involves interpretation of the transaction outcome by the buyer, a process which is subjective in nature. Further, a buyer needs to decide whether a particular transaction outcome merits a positive, neutral or negative feedback, if any feedback at all. These processes that determine the content and frequency of feedback are subjective in nature. The buyer’s trust decision involves making an assessment of the seller’s trustworthiness. This, again, is a subjective decision, as there are no objective measures or processes that translate the information in the feedback profile into an accurate risk assessment for the next transaction. Overall, then, the provision of feedback is voluntary and subjective, and the interpretation is also subjective. In my dissertation, I present three papers that consider implications of the voluntary and subjective nature of feedback and its interpretation in reputation systems.

1.4 Motivation

In this section I highlight some specific consequences of the voluntary and subjective nature of the provision and interpretation of feedback in order to motivate the specific questions I address. To illustrate the consequences, I begin with a base model of a market with adverse selection and moral hazard, similar to the debt market considered by Diamond (Diamond 1989), but with automatically provided feedback. There are three types of sellers, who trade in the market indefinitely: good type (G), bad type (B) and opportunist type (BG). In each period, each seller is matched up randomly with one buyer. The lifetime of each buyer is one period, in which they are trying to maximize their payoff. A buyer knows that there are three types of sellers in the market, and has a subjective belief about the distribution of the three types. The buyer does not know the type of seller she gets matched up with.

In each period, randomly matched pairs buyer and seller play the trust game as depicted in Figure 1. Type G and type B sellers play Stackelberg strategies 'always honest and 'always cheat' respectively. Type BG seller is a rational profit maximizer, who can choose his strategy every period, and is trying to maximize his profit over the k periods. The type BG sellers discount the future at a discount rate δ.

Suppose first that the market has a reputation system with public monitoring, so that the outcomes of all the transactions a seller was involved in, can be viewed by all subsequent buyers. As shown in Figure 1, the payoff for the seller if he cheats is c, and the payoff if he is honest is h. If the seller cheats in any period, then this fact is automatically revealed to potential future buyers in the form of a negative feedback.

A buyer will buy from a seller only if she believed that the probability that the seller is type G is at least λ, where $\lambda g + (1- \lambda)b = 0$

In equilibrium no buyer will buy from a seller who has ever received a negative feedback (Diamond 1989).

Now consider the decision faced by a rational type seller:

If he chooses to cheat, he can never sell again, and has to be satisfied with c, the payoff he gets in the single period.

Payoff(cheat) = c

If he chooses to be honest, in equilibrium buyers will buy from the seller as long as he has no negative feedback.

$\text{Payoff(honest)} = h\,(1 + \delta + \delta^2 + \delta^3 + \ldots) = h/(1-\delta)$

The type BG seller will choose to be honest only if

$h/(1-\delta) > c$ Equation I

Using this framework, I now demonstrate how the subjective and voluntary nature of feedback and its subjective interpretation pose several challenges to the functioning of a reputation system.

1. Consequences of Voluntary Feedback

When feedback is provided voluntarily, not all transactions get a feedback. To restrict our focus on the voluntary nature of feedback, let us consider a reputation system where feedback provision is incomplete, but free from any noise. In such a system some transactions may not get any feedback, but trustworthy behavior will never get a negative feedback, and untrustworthy behavior will never get a positive feedback.

The knowledge that some transactions will not get a feedback can affect the seller's incentives greatly. When feedback is voluntary, the sellers know that some transactions will not get a feedback, but they cannot predict beforehand whether a particular transaction will get a feedback or not. Sellers' beliefs about buyers' propensity to give a positive and a negative feedback will govern their behavior. If buyers are reluctant to provide positive feedback, it becomes difficult for sellers to acquire a good reputation. On the other hand, if sellers think that buyers refrain from giving negative feedback, they might be tempted to cheat. One of the main challenges for a reputation system is to ensure that sufficient feedback is provided to deter such behavior.

Reputation systems often rely on buyers' having non-monetary motives such as altruism, and reciprocity for providing feedback. Reciprocity is particularly important, since many reputation systems allow buyers and sellers to rate each other. Dellarocas and Wood (Dellarocas and Wood 2007) find evidence that a trader's propensity to provide feedback increases after she received a feedback from her partner. However, this raises fears that buyers may be reluctant to post negative feedback for the fear of retaliation. Bolton et al (Bolton, Greiner et al. 2008) corroborate Dellarocas and Wood's finding about reciprocity, and also provide evidence for retaliation.

Use of explicit incentives for feedback provision has also been considered by many researchers. The simplest of such mechanisms would be direct payment for feedback by the platform operator, or by future buyers who benefit from feedback. Li (Li 2006) proposes a system in which the seller provides rebates to those who provide feedback. In a market with adverse selection and moral hazard, Li demonstrates the existence of a pooling equilibrium where both good and bad type sellers offer a rebate, leading to an increase in feedback provision. Alternatively, Gazzale (Gazzale 2004) proposes a mechanism where he tries to use the voluntary nature of feedback to overcome the problem of under-provision. In Gazzale's mechanism, buyers are allowed to acquire a reputation for feedback provision. Sellers can discriminate among buyers based on their feedback giving history, which can act as an incentive for buyers to provide feedback. Under such a mechanism, Gazzale demonstrates the existence of equilibria with high levels of trust, trustworthiness and efficiency, even with incomplete feedback provision.

Eliciting sufficient feedback is a challenge faced by all practical reputation systems. A number of mechanisms have been proposed to address this problem, and in theory they achieve attractive outcomes in equilibrium. In most cases these desirable equilibria require using complex information structures, coordination schemes or sophisticated strategic reasoning on part of the participants. For example, in a mechanism where buyers' feedback giving history is made publicly available, the equilibrium depends on users having access to second order reputation information, i.e., information about the partner's past actions, and also those of his past transaction partners, in order to assess whether the partner's bad actions were legitimate punishments of their past transaction partners. Such complex information structures and coordination schemes may be difficult to implement in practical reputation systems, and all participants may not have the ability to comprehend and play the equilibrium strategy. All of these factors may result in a mechanism not having its intended effect. In chapter 2, I consider a mechanism designed to give buyers incentives for feedback provision, and examine its effects on buyer and seller behavior in a controlled laboratory setting.

2. Consequences of Subjective Provision of Feedback

So far, we are assuming that feedback provision is voluntary (and hence incomplete), but free of errors. In reality, the process which determines the content of feedback (positive or negative) is neither objective nor error free. As seen in Figure 4, feedback provision involves two highly subjective processes:

1. Buyer interpreting the transaction outcome
2. Buyer reporting this outcome in the form of a positive or negative feedback

There are no objective measures for distinguishing between high and low quality transaction outcomes. The perceived transaction outcome depends on a number of factors, such as price, price relative to buyer's value for the item, buyer's expectation from the seller, buyer's prior experience in the electronic market and in other places, among others. There are no objective guidelines that specify what transaction outcome warrants a positive v/s a negative feedback. Some buyers are likely to be harsher than others, and as a result a transaction outcome that gets a negative feedback from one buyer might not get a negative feedback from another. There are also additional complications stemming from factors that are beyond the seller's control and can affect transaction outcome (loss or damage in shipping).

The strategies and incentives of both the seller and the buyer change radically as a result of this. If the buyer gives a negative feedback for every transaction she deems unsatisfactory, and subsequent buyers employ a trigger strategy (never purchase from a seller who has a negative feedback), some good sellers might be driven out of the market. Also, an opportunist type seller aware of this is possibility might be tempted to cheat and 'make a kill' instead of being driven out of the market because of one buyer's harshness or a quirk of fate.

On the other hand, buyers may be forgiving. Unsure of the transaction quality, a buyer might decide to not give a negative feedback. A buyer can also use the seller's current reputation to decide what type of feedback (if any) a particular transaction deserves. In theory, the equilibrium where the opportunist sellers cooperate critically hinges on the sellers expecting any transgression to be punished. If the buyers are forgiving for individual bad outcomes, the sellers may be tempted to cheat once in a while. Individual transgressions may be pardoned, but Cripps et al (Cripps, Mailath et al. 2004) show that if the seller cheats often enough, his true type will be revealed eventually. This will result in the buyers never buying from him. The availability of cheap pseudonyms in online markets means that sellers with a bad reputation might try to reenter the market with a new id. However, buyers aware of this fact will be reluctant to trust newcomers, thereby further reducing the market efficiency (Friedman and Resnick 2001).

When buyers are unsure of the transaction quality or of the seller's action that resulted in a particular transaction outcome for them, they may use additional information, such as the seller's current feedback profile, to determine what feedback a transaction merits. In particular, a recent negative feedback in the seller's history may bias the buyer into judging him in a negative light, and give rise to the possibility of what Resnick and Zeckhauser (Resnick and Zeckhauser 2002) refer to as 'stoning'. If such a phenomenon exists, it would indicate that feedback is not just a reflection of the perceived outcome of current transaction, but also depends on past profile, thus making the objective interpretation of feedback difficult. On the other hand it may also act as an additional safeguard against opportunistic cheating, because it makes negative feedback even more unattractive. In Chapter 3, I conduct an empirical analysis of transaction and feedback data from eBay and examine the effects of a negative feedback in the seller's profile, in particular whether there is evidence of stoning.

3. Consequences of Subjective Interpretation of Feedback

Reputation solves the seller's moral hazard problem only if the long term gains from having a good reputation are sufficient to offset the foregone profits in the present (Klein and Leffler 1981). In the context of an electronic market, having a good reputation should translate to higher prices and probability of sale for sellers (Dellarocas 2005). All else held constant, the amount a buyer is willing to bid would depend on the buyer's assessment of the seller's trustworthiness (Ba and Pavlou 2002). This is where the subjective interpretation of feedback comes into the picture. The buyer would be interested in predicting the probability of getting a good transaction from the seller. In the context of the framework presented earlier, the buyer would be interested in predicting the probability that the seller is type G, or a type BG who will act honorably in this transaction. A buyer's beliefs about the distribution of different seller types in the population, and beliefs about the proportion of good and bad transactions among those that did not get a feedback, will affect her assessment of the seller's trustworthiness.

When feedback is voluntary and hence incomplete, the buyers may also differ in their interpretations of missing feedback. The buyers have no way of knowing whether the transactions that got a feedback were a representative sample of all transactions, or whether they were biased in some way. If a buyer believes that feedback provision is biased in such a way that the probability of reporting an unsatisfactory transaction is low, she will be reluctant to trust a seller even if he has no negative feedbacks.

Presentation of feedback information can affect how it is interpreted. Kuan (Kuan 2005) examines the effect of information presentation on buyers' judgment about seller trustworthiness, and finds that the same information presented in different formats is interpreted differently. Even with the same presentation format, different buyers may make different assessments of seller trustworthiness, when presented with the same information. This may be partly due to the differences in their experiences with other sellers, but there could also be some systematic differences based on age, gender, income, nationality or other demographic categories.

Differences in trust among people from different nationalities are well documented (Fukuyama 1995). Thus, there may be systematic differences in the buyer's prior beliefs about the prevalence of different types of sellers or the meaning of missing feedback. By providing information about a seller's past dealings, and by giving sellers an incentive to be trustworthy, a reputation system is designed to reduce social uncertainty. It remains to be seen, however, whether cultural differences impact the effectiveness of reputation systems. I address this question in Chapter 4, which is a laboratory experiment with student participants from one individualist and one collectivist nation.

Summarizing, the voluntary and subjective nature of feedback and its subjective interpretation are important characteristics of practical reputation systems, and pose several serious challenges to their functioning. This raises several questions of interest to researchers, designers and managers of reputation systems, and participants in environments that employ such systems. Some of these questions have already been addressed by other researchers.

In this context my contribution is twofold. First, I highlight the various effects of the voluntary and subjective nature throughout the lifecycle of feedback- provision, presentation, interpretation and use. Identifying the common theme in literature from various disciplines, I help pave the road towards a holistic understanding of practical reputation systems. I also contribute by answering three specific questions that arise due to the voluntary and subjective nature of feedback provision and its subjective interpretation.

- Does making buyers' feedback giving history available to sellers, result in higher levels of feedback provision, and thus trust and trustworthiness?
- How does a negative feedback affect a seller's future behavior, and the feedback giving behavior of the buyers who trade with him in future?
- Can a reputation system eradicate the systematic differences in trust and trustworthiness that are known to exist between different cultures?

1.5 Outline

The rest of my dissertation is organized as follows. In Chapter 2, I examine whether making buyers' feedback giving history available to sellers, results in higher levels of feedback provision, trust, and trustworthiness. This experimental study is based on a mechanism proposed by Gazzale (Gazzale 2004), which is intended to give buyers an incentive to provide feedback. Gazzale's original mechanism uses a complex information structure and sophisticated coordination schemes, which are difficult to implement in practical reputation systems. The mechanism I study is a simpler mechanism based on the same insight: if buyers are allowed to develop reputations for feedback provision, reputation effects can arise on the buyer side and provide buyers with incentives to give feedback. The focus of my study is the users' ability to understand the importance of buyers' feedback giving history. The primary requirement for reputation effects to arise on the buyer side is that sellers should discriminate among buyers based on their feedback giving reputation, and the buyers should be aware of this fact. I examine whether these requirements are met by undergraduate student subjects interacting in a laboratory setting.

Chapter 3 is an empirical study, where I study the effects of a negative feedback on sellers at eBay. When a seller gets a negative feedback, it marks a significant change in his feedback profile. There is empirical and anecdotal evidence that a negative feedback leads to a decrease in expected revenues and is a traumatic event for many eBay sellers. In this chapter I study other effects of negative feedback (besides decrease in expected revenue) by analyzing a panel of transactions for a large number of eBay sellers. In particular, I examine whether a negative feedback changes sellers' behavior in future transactions and buyers' willingness to give negative feedback in those transactions.

In Chapter 4, I conduct a laboratory experiment that deals with interpretation of feedback. I focus on the systematic differences in trust that are known to exist between different cultures, and examine whether a reputation system is able to eradicate such differences in a controlled laboratory environment. Individualism v/s collectivism is a cultural dimension that influences trusting and trustworthy behavior. Individualist people are known to be more trusting and trustworthy than collectivist people when interacting with strangers, a situation commonly faced in electronic markets. I examine whether these differences persist in the presence of a reputation system. I also examine whether there are systematic differences in how people from different cultures assess comparable reputation profiles.

Chapter 5 summarizes the findings from these three studies. It places them all in the context of the overall framework of the impacts of voluntary, subjective feedback provision and subjective interpretation of feedback profiles. It also extracts practical consequences for design and points toward fruitful directions for future research.

Chapter 2

Remain Silent and Ye Shall Suffer: Seller Exploitation of Quiet Buyers in Reputation Systems

2.1 Introduction

Consider a transaction in an electronic marketplace. The seller in such an environment faces a moral hazard. The buyer pays first and has no way to enforce that the seller fulfills the transaction as agreed. Upon receipt of payment from the buyer, the seller has no incentive to carry out his end of the contract. If the seller and the buyer expect to interact repeatedly, the folk theorem suggests that the threat of lost future custom can induce a sufficiently patient seller to act in a trustworthy manner (Fudenberg and Maskin 1986).

In many interesting e-commerce environments, two key features may render such threats empty. First, while both the buyer and seller may interact repeatedly with the e-commerce platform, a particular buyer and seller are likely to interact infrequently. As a result, the threat of lost future custom from the particular buyer is unlikely to be effective in inducing cooperation from the seller. The situation can be salvaged if other buyers condition their purchase decisions on the outcomes of the seller's past transactions. Threat of lost future business, with the current buyer or with potential future buyers, can help overcome the seller’s moral hazard.

However, transactions in electronic marketplaces are games of private monitoring- only the buyer knows whether she received the promised quality. Absent some mechanism for information sharing, there is no way for future buyers to know whether the seller acted trustworthily in the current transaction. Thus, the seller's choice of action in the current transaction has no future implications, and the buyer's threat to withhold future custom is likely to be insufficient to induce trustworthy behavior.

Many electronic marketplaces employ a reputation system to solve the moral hazard problem and induce trustworthy behavior from the seller.[3] A reputation system removes the disconnect between the present and the future, by making a record of the seller's past transaction outcomes available to potential buyers. Kandori (1992) shows that any mutually beneficial outcome (such as trusting buyers and trustworthy sellers) is attainable in games with perfect private monitoring, as long as each agent is sufficiently patient and carries a systematically updated label about his past behavior (such as a reputation score).[4] In theory, a reputation system in a marketplace where sellers interact with particular buyers infrequently but with the marketplace indefinitely, can support the same equilibrium level of trust and trustworthy behavior as a buyer and seller interacting repeatedly. Future buyers condition their behavior on sellers' current actions, thereby providing sellers with an incentive to act in a trustworthy manner in the current

---

[3] In addition to providing incentives to act in a trustworthy manner (i.e., solving the moral hazard problem), reputation system may also help in deciding who to trust (i.e., solving the adverse selection problem). In this study, we focus on the moral hazard problem.

[4] Monitoring is perfect if an agent's outcome perfectly reveals another agent's action. In our case monitoring is perfect, if a buyer is happy with a transaction, when and only when the seller acted in good faith.

transaction. Thus, reputation systems can work as a feasible and less-costly substitute for legal enforcement (Bakos and Dellarocas 2003).

In Kandori's model, information about an agent from previous interactions becomes automatically available to her current partner. This is not true for real world electronic markets, which have to rely on transaction participants voluntarily reporting transaction outcomes (i.e., providing feedback). However, in many electronic commerce environments, the incentives to share private information about a particular seller's trustworthiness with other buyers are quite weak. Information sharing is at least marginally costly, and private benefits are small as the beneficiaries are likely to be strangers, the seller's future potential customers. Feedback is thus akin to a public good, and is likely to be underprovided.

When considering the question of sufficient information provision, we would like to note two points. First, the level of information sharing necessary to induce trustworthy behavior will depend on the environment. For example, if sellers are patient and the gains from successful transactions large, even a small feedback probability may be sufficient to deter opportunistic behavior (Dellarocas 2004).

Second, it is quite plausible that non-pecuniary motives, such as altruism or emotional satisfaction, may induce some buyers to provide feedback, particularly negative feedback[5]. There is ample empirical evidence to suggest that people are willing to punish those who violate trust, even when such punishments are costly to the punisher.

[5] We borrow eBay's terminology and say that a buyer provides positive feedback if she reports that the seller shipped the item, and negative feedback if she reports that the seller did not ship the item.

For example, Fehr and Gachter (Fehr and Gachter 2000) find in their experimental study of the voluntary contribution mechanism, that community members are willing to incur a monetary cost in order to punish a free rider by reducing his payoffs.

Yet, it is not obvious that leaving feedback is as effective in eliciting emotional satisfaction as monetary rewards or punishments. Whereas monetary punishments take effect directly and impose clear damage, undermining a seller's reputation is an indirect punishment whose impact is difficult to measure. The damage, if any, takes place in the future, and may not be observable to the buyer. Furthermore, the punishment is uncertain, since it requires the participation of others, namely the seller's future potential partners. Thus satisfaction from providing feedback, and thus arguably incentives for providing this information, depend crucially on beliefs about the system's effectiveness. The incentives for leaving positive feedback are even more tenuous. For these reasons, the problem of inducing costly information provision in reputation systems invites further investigation.

A number of mechanisms have been proposed to overcome the problem of insufficient information sharing. Direct payment for feedback by the platform operator is perhaps the most straightforward. Alternatively, Li (Li 2006) proposes a system in which the seller provides rebates to those who provide feedback. Such a system is based on the insight that the trustworthy seller has greater incentives to offer such a rebate. Thus the rebate offer acts as a signal of the seller’s trustworthiness. In a market with adverse selection and moral hazard, Li demonstrates the existence of a pooling equilibrium where both good and bad type sellers offer a rebate, leading to an increase in feedback provision.

Another approach, proposed by Gazzale (Gazzale 2004), is to allow buyers to acquire a reputation for information sharing. Gazzale considers environments where sellers are provided with a measure of the buyers' feedback provision history. Equipped with this information, sellers can discriminate among buyers based on their information sharing reputation. A buyer is considered to have a good reputation if she provides sufficient feedback. Gazzale demonstrates the existence of equilibria, where:

- Buyers always trust sellers with a good reputation, such sellers are always trustworthy to buyers with a good reputation, and the buyers, in turn, reward sellers by providing positive feedback with a high probability.
- If a buyer deviates from this equilibrium strategy (by not providing sufficient feedback), she acquires a bad reputation, and future sellers punish her by cheating. Likewise, if a seller deviates from the equilibrium by cheating a buyer with a good reputation, the buyer gives him a negative feedback with a high probability, which results in him getting punished by future buyers, who refuse to buy from him for a finite punishment period.

Gazzale shows that in such equilibria, high levels of trust, trustworthiness and efficiency are attainable, even with costly information provision and less than complete feedback, as long as the players are sufficiently patient. How often a buyer must provide feedback to maintain a good reputation depends on the discount factor, profits from trade, cost of feedback provision and the length of punishment period.

In theory the mechanism with buyer feedback provision reputation shows great promise, as it can support highly efficient outcomes in equilibrium, even with incomplete feedback provision. However, the mechanism achieves these desirable outcomes using complex information structure (second order reputation information about whether the person a seller cheated had been a reliable feedback provider) and coordination schemes (punishment phase, minimum frequency of feedback, repent action etc.), which may be difficult to implement in practical reputation systems. A closer look at the buyer feedback provision reputation mechanism reveals that in spite of the complicated information structure and coordination schemes, the key idea behind the efficient equilibrium outcomes is simple: When buyer feedback provision histories are made available to sellers, reputation effects kick in on the buyer side, and provide buyers with an incentive to choose the costly action of providing feedback, in anticipation of favorable outcomes in future. Perhaps this basic effect could occur in practical settings where the full information structure supporting the equilibrium is not present.

We consider a market with a reputation system, where buyers' feedback giving histories are made available to sellers. In contrast with the mechanism in Gazzale's theoretical paper, such a mechanism can be easily implemented in most practical reputation systems. In fact, eBay already makes the feedback giving histories of all users publicly available. Even in the absence of second order information and coordination schemes like those suggested by Gazzale, sellers have the necessary information to differentiate among buyers. If sellers believe that a buyer's feedback giving history is a predictor of the likelihood of getting a feedback from her, then sellers can be selectively trustworthy with some buyers. If sellers practice such discrimination and the buyers realize this fact, then reputation effects imply that buyers should provide feedback more often to avoid being discriminated against. For this simple mechanism to achieve desired outcomes, the following conditions must be satisfied:

- Sellers discriminate among buyers using their feedback provision history. They are trustworthy towards buyers who have provided feedback regularly, but would cheat those who have not.
- Buyers are aware of the possibility of such discrimination. Thus, buyers provide feedback regularly with a high probability, leading sellers to behave in a trustworthy manner, and consequently leading buyers to trust.

To investigate whether the outcome of trust, trustworthiness and costly feedback provision can be realized in practice, we conducted a laboratory experiment. In our experimental setup, we make available only the current partner's reputation information (as a seller for not cheating and as a buyer for providing feedback). Feedback profiles used in our experimental setup are similar to those used on most practical reputation systems. Moreover, we do not enforce any explicit coordination scheme. Thus, the mechanism in our experimental setup allows for reputation effects on the buyer side, but is simple enough to be implemented in practical reputation systems.

For the buyer reputation, we only reveal whether a buyer provided a feedback in each completed transaction. We do not reveal the content of the feedback, i.e. positive or negative. We make this choice because the content of feedback, if visible, may complicate the buyers' incentive to acquire a reputation for information sharing. If the buyer always leaves positive feedback, she might be labeled a pushover by the seller and the seller might attempt to defraud her. Likewise, if she always leaves negative feedback, the seller might believe fruitless any attempt to satisfy her. Thus, the buyer might seek the ‘optimal’ ratio of positives to negatives as in Ely and Valimaki’s (Ely and Valimaki 2003) analysis, and not report truthfully.

While poor treatment of buyers who do not provide feedback has the potential to encourage feedback provision, an efficient trading game outcome depends on sophisticated strategic reasoning on the part of all participants. For reputation effects to arise on the buyer side, having a good reputation must fetch the buyers more benefit than the cost of acquiring a good reputation. In Gazzale’s theoretical model, costly feedback is a part of the equilibrium because failing to maintain information sharing reputation results in a punishment for the buyer. In Gazzale's model, a seller who punishes a reticent buyer faces no risk because future buyers can distinguish between ‘equilibrium punishments’ and ‘off-equilibrium deviations’ on part of the seller. But, this is not the case in our experimental market, or in most practical reputation systems. In the absence of second order information, if a seller punishes a reticent buyer and gets a negative feedback, he is susceptible to being punished by future buyers.

In the absence of the second order information cheating based on a buyer's feedback giving record makes sense only if the seller believes that a buyer's feedback giving history is a good predictor for the probability of getting a feedback from her. If such a belief is unwarranted, this may results into scenarios with inefficient outcomes like the following: A seller, who practices discrimination based on buyers' feedback provision history, cheats a reticent buyer and gets a negative feedback. In the next round he is matched with a buyer who has a reputation for information sharing. The seller expects that this buyer will give a feedback with a high probability, so he would not have cheated her. But the buyer chooses not to buy from him, as she is reluctant to trust the seller as a result of the negative feedback from the previous round.

On the contrary, if a buyer recognizes the importance of feedback provision history, she will seek a reputation for information sharing. When the content of feedback is not revealed, the optimal way to acquire such a reputation is to give feedback frequently regardless of its content (though not at regular intervals as that would be too predictable). However, the buyer may not realize this, and have different strategies for giving positive and negative feedback. This may be true, especially if buyers have non-pecuniary incentives such as joy of retaliation for giving negative feedback. The buyers may not realize that giving feedback only when it is negative in insufficient to build a reputation for feedback provision. Considering these various factors, and the contrasting incentives they present, we are interested in examining the effects of this mechanism on buyer and seller behavior, such as trust, trustworthiness and feedback provision.

This question deserves additional attention because of a design change made by eBay along these lines. Since 2004, eBay has made a member's feedback leaving history publicly available. Thus, not only can you access every feedback received by a member (as well as a ‘feedback score’ and the proportion of positive feedbacks), you can also access the content of every feedback left by that member (Figure 5), though there are no summary statistics about feedback giving. Our mechanism is similar to eBay’s except for the following differences:

1. If the buyer does not give a feedback on a transaction, the feedback giving history in our experiment reveals this fact. eBay, on the other hand, does not explicitly

report that a transaction did not get a feedback. Thus it is not possible to precisely calculate the ratio of feedbacks to transactions on eBay.

2. eBay displays the content of the feedback, while our mechanism only reports whether a feedback was given or not. As we discussed earlier, the display of feedback contents may provide incentives to provide dishonest feedback. Even when the buyers are restricted from giving dishonest feedback, as in our experiment, displaying the content of feedback is still problematic. When the content of feedback is available to the sellers, there is no set of beliefs about the buyers' feedback giving probabilities that can be sustained in equilibrium. In such a scenario, there would be ambiguity about the transactions that get no feedback. The sellers cannot tell whether the missing feedback would have been a positive feedback or a negative feedback. Thus the buyers cannot credibly commit to any feedback giving strategy. We thus choose not to reveal the content of feedback, and only reveal whether or not a feedback was given.

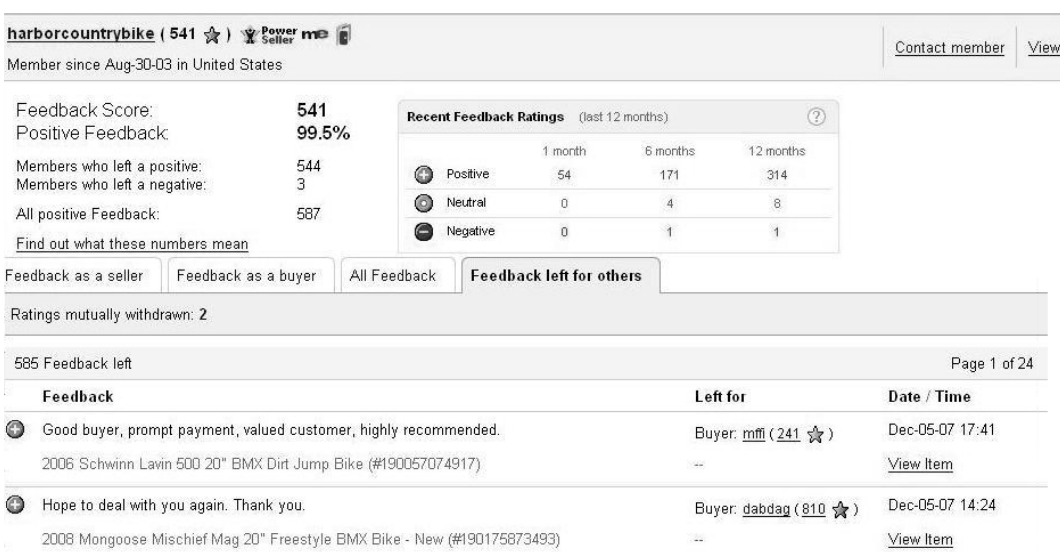

Figure 5 'Feedback left for others' on eBay

We conducted a laboratory experiment to better understand the impact, in practice, of making visible buyers' feedback provision history. There are several important questions. First, do sellers discriminate among buyers based on their feedback giving history? An important and related question is whether sellers stand to benefit from discriminating against reticent buyers. To answer this question, we examine whether reticent buyers are indeed less likely to respond with a negative feedback. Second, if some sellers discriminate against reticent buyers, will buyers realize that they are being punished for their quietness, and provide feedback more often? Or will they stop trusting the sellers and withdraw from the market altogether?

The rest of the chapter is organized as follows. The experimental design is presented in Section 2.2. In section 2.3 we present answers to the above questions based on our experimental data. We present a discussion of our results in Section 2.4, and conclude in Section 2.5.

2.2 Experimental Design

We use a lab experiment based on the design implemented by Bolton et al (Bolton, Katok et al. 2004), which we modify in order to capture the voluntary nature of feedback provision in a reputation system. We study an environment in which a seller transacts in each round with a randomly matched buyer in a modified trust game (Fig 2.2). In our modified trust game, if the buyer chooses to trust, she can (but is not forced to) reveal the transaction outcome to the seller's future partners (i.e. leave feedback). In this setting, we compare a market with information sharing reputation mechanism to a market where no such mechanism is available.

Information Condition	Cost Condition		
	No Cost (C = 0)	Low Cost (C = 2)	High Cost (C = 5)
FbHist: Buyer Feedback history available		FbHistC2 (4 sessions)	FbHistC5 (4 sessions)
NoHist: Buyer Feedback history not available		NoHistC2 (4 sessions)	NoHistC5 (4 sessions)
Automatic: Feedback automatically provided	A0 (Control) (4 sessions)		

Table 1 Summary of Treatments

In one set of treatments (denoted FbHist), sellers have access to a buyer's feedback leaving history, i.e. whether or not she revealed the actions of previous sellers. In another set of treatments (denoted NoHist), a seller in the market cannot access information about his current buyer's past behavior. In both the FbHist and NoHist treatments, we vary the cost to the buyer of leaving feedback (High Cost vs. Low Cost). In the low-cost treatments, we set C = 2, whereas in the high cost treatments we set C=5. While these costs are small relative to the value of the item, as we shall see they are rather significant relative to the gains from trade: approximately 13% in the low-cost treatments, and 33% in the high-cost treatments. In addition, we ran a baseline treatment (Automatic) in which feedback was automatically provided at no cost to the buyer. Table 1 summarizes the various treatments and the labels used to identify them.

2.2.1 Implementing the markets

At the beginning of each 12-subject session, the computer randomly assigns to half of the players the role of buyer and to the other half the role of seller. Role assignments are fixed for the entire session. A session consists of 45 rounds, which are divided into three sub-sessions of 15 rounds each. Each round consists of two subgames: the trading subgame followed by the information sharing (feedback) subgame. Figure 6 depicts the game tree.

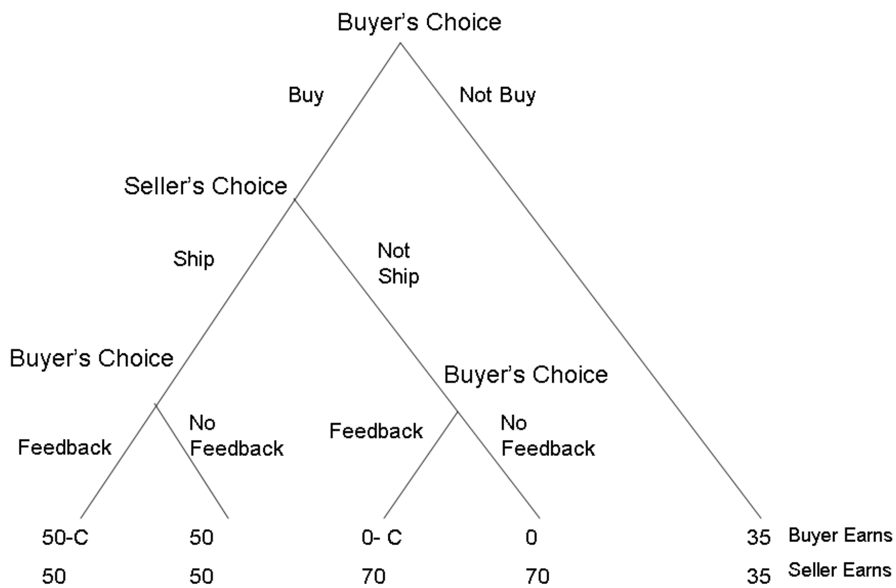

Figure 6 Game Tree

Trading subgame: In each round, both the seller and the buyer are endowed with 35 points, which is the payoff if no trade takes place. The seller offers a hypothetical item for sale at a price of 35; the buyer values the item at 50 points. Seller has no value for the item, but his cost of providing the item is 20 points. The buyer first decides whether to buy from the seller and send the payment of 35, or not buy and keep his endowment. If the buyer decides not to buy, the round is over. If the buyer decides to buy, then the seller decides whether to ship the item for a total payoff of 50, or not to ship and keep the buyer's payment along with her own endowment for a total payoff of 70. Thus, absent any information-sharing costs, successful completion of a trade will result in a surplus of 15 for the buyer and net profit of 15 for the seller. The gains from trade are about 40% of the endowment, which has been shown to be sufficient to induce high levels of transactions in at least some markets (Bolton, Katok et al. 2004). When the trading game round is over, each player can view a summary of his and his partner's actions on the computer screen.

Information sharing subgame: If the buyer decided to buy in the trading subgame, she decides whether or not to provide feedback by revealing the transaction outcome, i.e., reporting whether or not the seller shipped the item. Leaving the feedback costs the buyer C points, where C represents the opportunity cost of providing feedback [6]. The cost of feedback provision is the same for all buyers and is fixed during a session. If the buyer decides not to buy, future partners of the seller are automatically informed of this fact at no cost to the buyer.

We point out two key features of the information sharing subgame. First, in some reputation systems (such as eBay's), both the buyer and the seller rate each other. As a result, there are strategic concerns, such as the expectation of reciprocity, that affect the information sharing decision. To avoid the confounding effect due to such strategic concerns, we allow only the buyer to rate the seller. Second, we restrict the buyer's actions so that she can provide only honest feedback. The buyer can only choose whether or not to reveal the outcome of the transaction that she participated in. Thus, if the seller shipped the item, the buyer cannot incorrectly report that he did not ship.

Information: In all treatments, both the buyer and seller in a transaction are presented in every round with the seller's transaction summary and history. The transaction summary shows the following statistics about the seller's past transactions in the current subsession: the number of items sold, the number of positive and negative

[6] In an actual reputation system, there are no monetary costs of feedback provision. There is, however, an opportunity cost in the form of the time and effort required for feedback provision. In our experiment, we use the monetary cost as a proxy for the opportunity cost of feedback provision.

feedbacks received, and the number of rounds with no feedback. In addition to these summary statistics, the transaction history shows for each previous round in the current subsession whether the seller sold an item, whether feedback was left and, if feedback was indeed left, whether the seller did or did not ship the item.

In the FbHist treatments, a buyer also carries a history. Whereas a seller in the NoHist treatments cannot access information about her current buyer's past actions, we present a seller in the FbHist treatments with his current buyer's transaction summary and history. The buyer's transaction summary shows the following statistics about her past transactions in this subsession: the number of items bought, the number of times she decided to leave feedback, and the number of rounds in which she did not provide feedback. The buyer's transaction history shows, for each previous round in the current subsession, whether the buyer bought an item and, if so, whether or not he provided feedback. As discussed in the previous section, we do not reveal the content of feedback, i.e. whether it was positive or negative.

Subsessions: In order to allow players the possibility of establishing and learning a social norm, we divide each 45 round session into three 15-round subsessions. Players retain the same roles, but at the end of the first and second subsessions, all transaction summaries and histories are reset. Thus, a buyer cannot access any information about a seller's actions in a previous subsession, nor can a seller in FbHist treatment access a buyer's buying and feedback history from a previous subsession.

2.2.2 Experimental protocol

The instructions handed out to the participants are reproduced in Appendix I. The experimental protocol is described in detail in the instructions. A synopsis of the same is presented here:

We conducted 4 sessions for each of the 5 treatments. There were 12 subjects per session (48 per treatment), for a total of 240 subjects. For each treatment, two sessions were conducted at Williams College and two sessions at the University of Michigan. The experiment was programmed and conducted with the software z-Tree[7] (Fischbacher 2007). Each session consists of 45 rounds, with buyers and sellers randomly rematched between rounds. We divided each session into three fifteen-round subsessions. We made the number of rounds public knowledge in order to control for end-game effects.

We conducted four sessions of each treatment in April and May, 2006—two sessions at Williams College and two at the University of Michigan. Participants were students at each institution[8]. Sessions lasted approximately one hour on average. The payoffs in the games were in points. At the end of the session, a user's total earnings in points were converted to U.S. dollars at a pre-announced fixed rate of 100 points for 1 dollar. The average payoff was about $24, including a $3 show-up fee.

2.3 Results

[7] We thank Kan Takeuchi for programming the treatments in z-Tree and Yan Chen for the financial support for programming.

[8] Two Michigan students were graduate students; the remainder at both institutions were undergraduates.

We use three metrics to evaluate the performance of the trading game, namely trust, trustworthiness and efficiency. We define trust as the percentage of buyers who choose to buy, trustworthiness as the percentage of sellers who decide to ship (conditional on buyers choosing to buy), and trading game efficiency as percentage of transactions where the potential gain from trade is realized (Note that trading game efficiency does not include the deadweight loss of feedback provision costs). To measure the performance of the evaluation game, we use feedback rate (denoted fbRate) defined as the overall proportion of feedbacks provided (regardless of the outcome). We also use positive feedback rate (denoted posRate) defined as the proportion of feedbacks provided when the seller chose to ship, and negative feedback rate (denoted negRate) defined as the proportion of feedbacks provided when the seller chose to not ship. We begin with some descriptive statistics and an examination of the round-by-round performance of each treatment.

2.3.1 Treatment effects on performance

We begin by examining the round-by-round performance of the trading game in each treatment. Figures 7-12 depict the round-by-round evolution of trust, trustworthiness and trading game efficiency for both cost conditions. For each of these metrics, an observation in the graph represents the aggregated results for a particular round from all sessions for a treatment. For example, the statistic for trust for buyer history low cost treatment in round 10 is the proportion of buyers that choose to buy in round 10 across all four sessions for the that treatment.

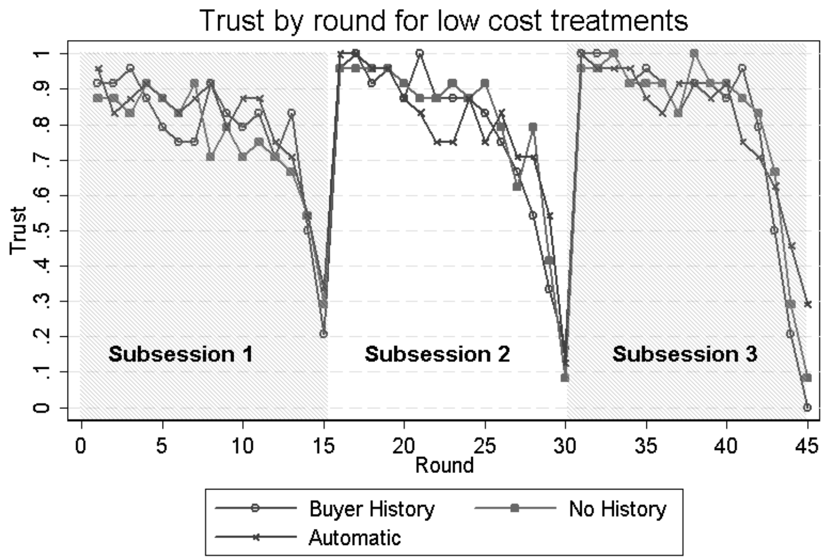

Figure 7 Trust by round for low cost treatments

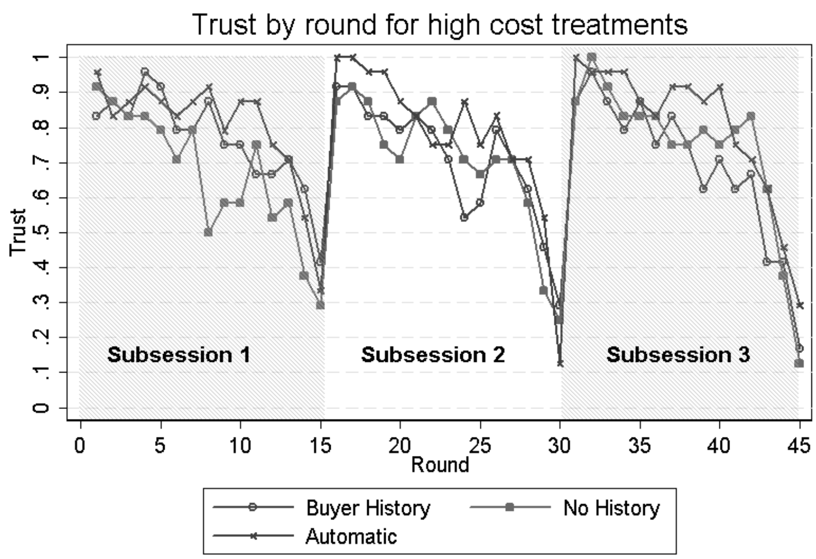

Figure 8 Trust by round for high cost treatments

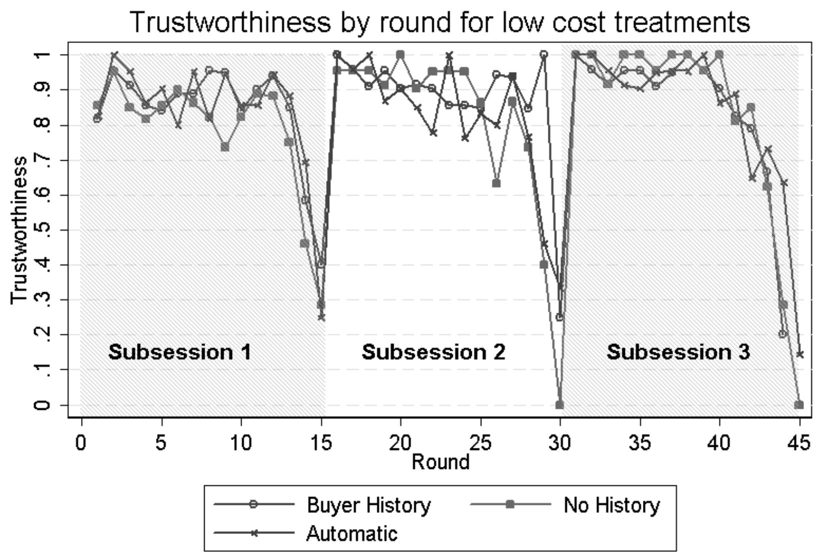

Figure 9 Trustworthiness by round for low cost treatments

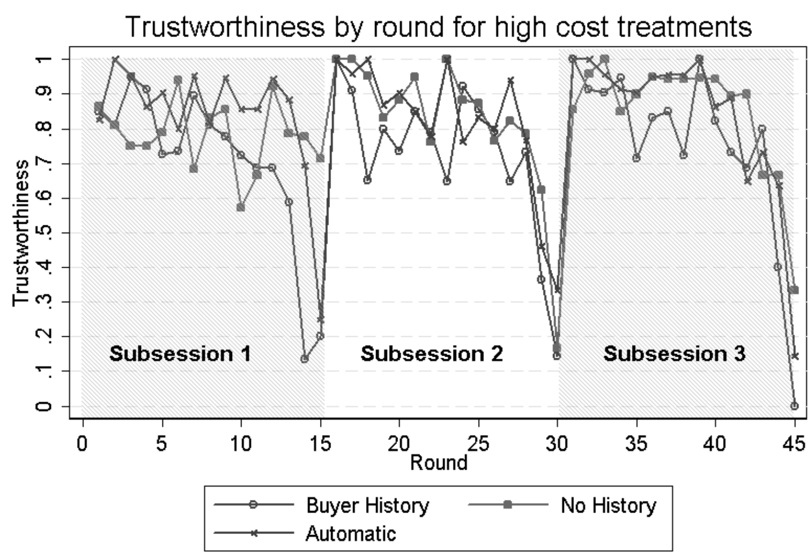

Figure 10 Trustworthiness by round for high cost treatments

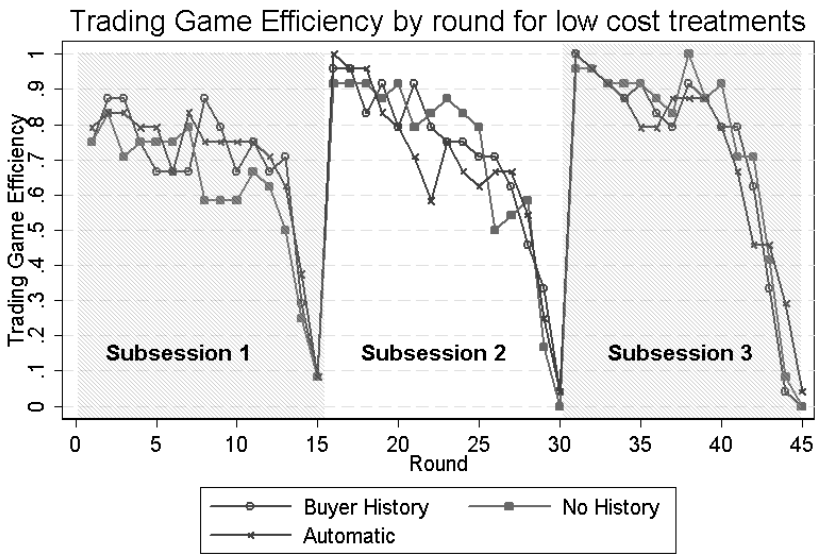

Figure 11 Trading game efficiency by round for low cost treatments

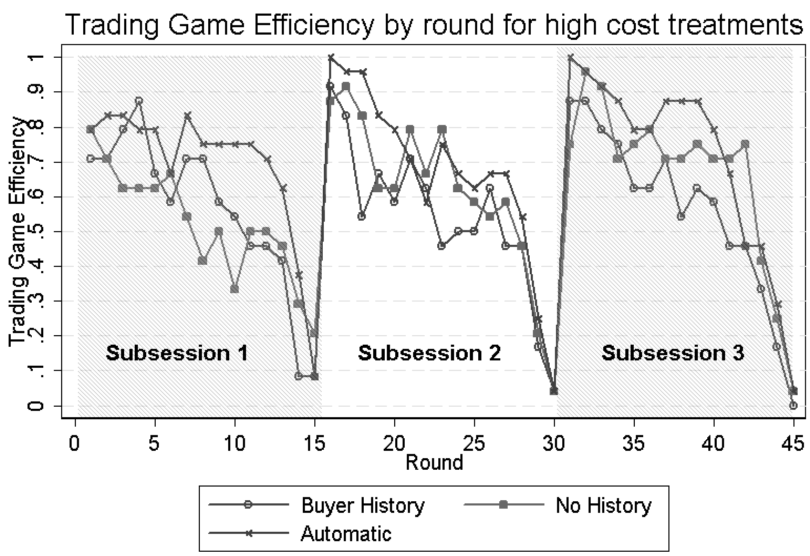

Figure 12 Trading game efficiency by round for high cost treatments

A few observations can be made immediately from the graphs. First, in all treatments, trust, trustworthiness and trading game efficiency are quite high in the first round of each subsession. They gently decrease in subsequent rounds until dropping precipitously in the final few rounds of the subsession due to the end-game effect. However, in the first rounds of subsessions 2 and 3, trust and trustworthiness are again quite high (in fact higher than the first subsession). This suggests that the precipitous drops in trust and trustworthiness in the final few rounds of a subsession are a result of situation specific strategic behavior on the part of the subjects, and not an indicator of overall market failure. Similar behavior has been observed in repeated public goods games, where high levels of cooperation are achieved when the game re-starts (Ostrom 2000). Second, the treatments with high cost of information sharing have distinctly lower trust, trustworthiness and trading game efficiency than the treatments with low cost of information sharing. Also the levels of trust, trustworthiness and trading game efficiency in the low cost treatments do not appear to be very different from those in the baseline treatment with automatic feedback. Table 2 presents a summary of the aggregate trading game performance computed using the outcomes in the first 12 rounds of each subsession in all sessions of each treatment. We drop the last three rounds in each subsession to control for end-game effects.

Treatment	Trust	Trustworthiness	Efficiency
Low Cost Buyer Feedback Provision History	0.88	0.91	0.81
Low Cost No Buyer History	0.87	0.91	0.80
High Cost Buyer Feedback Provision History	0.79	0.82	0.64
High Cost No Buyer History	0.78	0.87	0.68
Automatic Feedback	0.87	0.90	0.79

Table 2 Aggregate Trading Game Performance

Now we shall examine the round-by-round performance of the evaluation game for each treatment. Figures 13-16 depict the round-by-round record of feedback provision. Figures 13 & 14 compare the feedback rates by feedback provision history for high and low cost treatments. While figures 15 & 16 compare the feedback rates by cost for treatments with and without buyer feedback provision history.

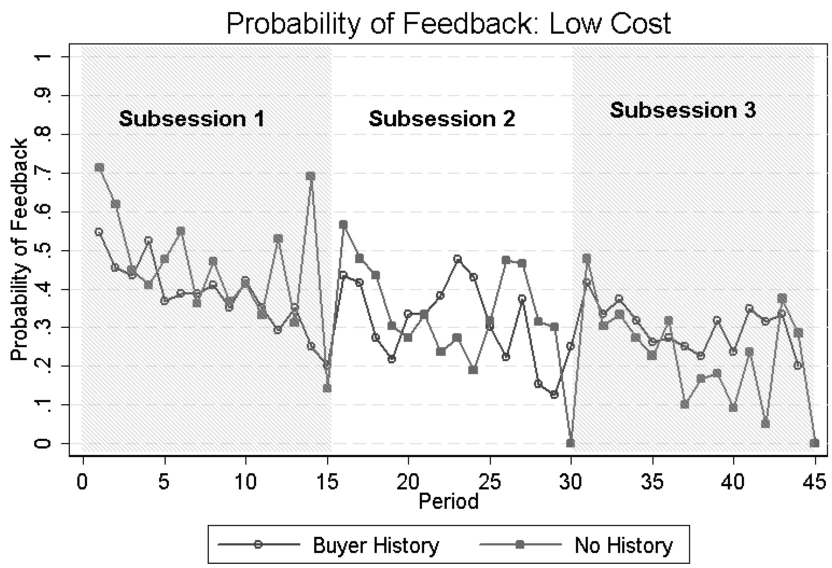

Figure 13 Comparing Feedback Rates by Feedback Provision History for Low Cost

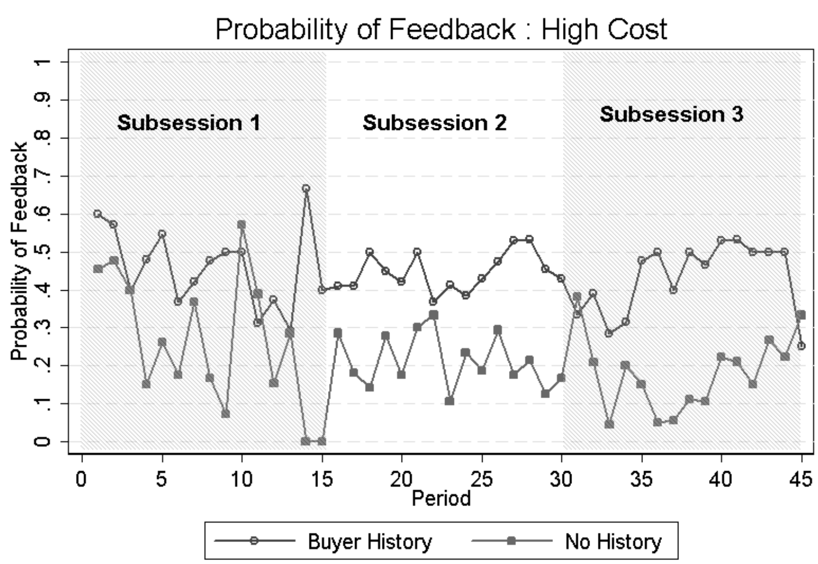

Figure 14 Comparing Feedback Rates by Feedback Provision History for High Cost

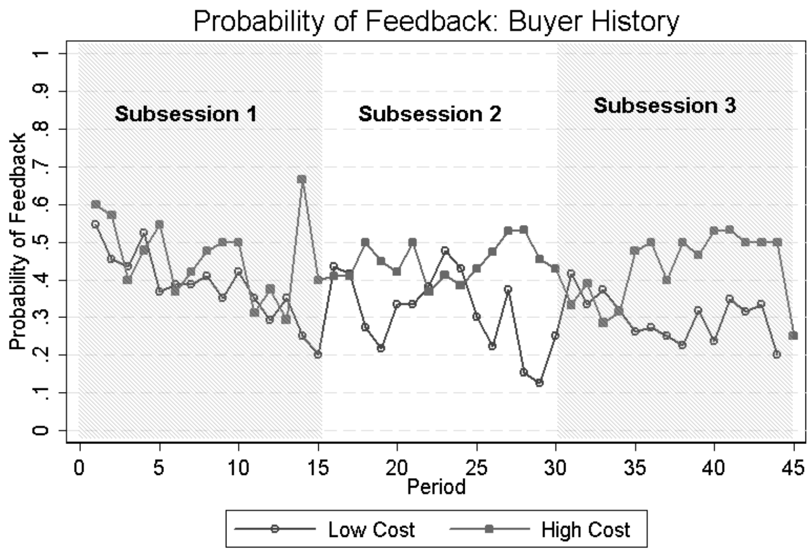

Figure 15 Comparing Feedback Rates by Cost for Buyer History Treatments

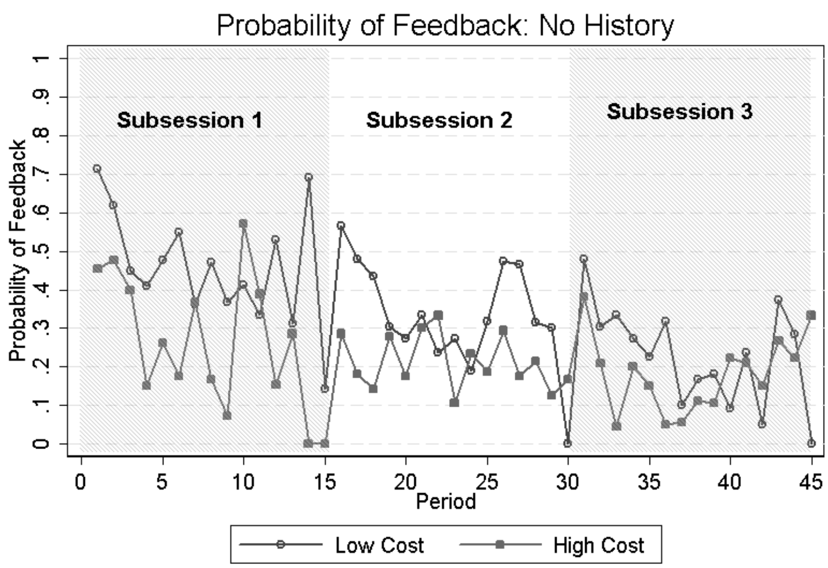

Figure 16 Comparing Feedback Rates by Cost for Buyer History Treatments

Figure 13 &14 suggest that with low cost of feedback, there is not much difference in the feedback provision rate with and without the buyer feedback provision history. However, in the high cost treatments, the availability of buyer feedback provision history results in higher feedback provision, especially in the latter subsessions. The findings from figure 15 are surprising, and somewhat counterintuitive. In the absence of buyer feedback provision history, the low cost treatment generally has higher feedback provision than the high cost treatment, as one would expect. However, when buyer feedback provision history is available, the high cost treatment has higher feedback rates, especially in the latter subsessions. If we consider positive and negative feedback separately, we find higher provision of negative feedback than positive feedback in all treatments. In the low cost condition, the presence of buyer history led to higher negative feedback but did not affect the provision of positive feedback. In the high cost condition treatment with buyer history had higher provision of positive feedback but no significant difference in the provision of negative feedback. These results are discussed in detail later in the section; the aggregate feedback rates for various treatments are presented in Table 3.

Treatment	Feedback	Positive Feedback	Negative Feedback
Low Cost Buyer Feedback Provision History	0.36	0.30	0.91
Low Cost No Buyer History	0.35	0.32	0.69
High Cost Buyer Feedback Provision History	0.45	0.39	0.70
High Cost No Buyer History	0.23	0.16	0.73

Table 3 Aggregate Feedback Rates

In the following subsections, we analyze the data and seek an answer for each of our questions.

2.3.2 Do sellers discriminate among buyers based on their feedback giving history?

The availability of buyer feedback giving history provides sellers with the opportunity to be selectively trustworthy towards some buyers, while cheating others. If the sellers think, based on a particular buyer's record, that she is unlikely to give a feedback, they have little incentive to be trustworthy towards her. Whereas if a buyer's record suggests that there's a high probability of getting a feedback from her, then sellers would not want to cheat her. In our experimental setup, sellers can view whether a buyer gave a feedback for each of her past transactions, but they cannot view the content of the feedback, i.e. whether the feedback was positive or negative. Sellers can make probability estimates for the buyer giving a positive feedback and a negative feedback, using a buyer's overall and recent feedback giving history.

We use RoundsSinceBuyerGaveFB as a measure of the buyer's recent feedback provision history. RoundsSinceBuyerGaveFB is defined as the number of transactions since the buyer last gave a feedback, before the current transaction in the current subsession. For Example, if RoundsSinceBuyerGaveFB =2 , it means that since the last time the buyer gave a feedback, there were two periods where she chose to buy, but did not give any feedback. Based on the seller's belief about the buyer's feedback giving strategy, RoundsSinceBuyerGaveFB could be used in different and contrasting ways. If the sellers believed that a buyer gave feedback at regular intervals, then a high value of RoundsSinceBuyerGaveFB would indicate that the buyer is 'due' to give a feedback, and therefore not to be cheated. On the other hand if sellers believed that lack of recent feedback indicates that a buyer has become reluctant to give feedback, they would be encouraged to cheat when the buyer has a high value of RoundsSinceBuyerGaveFB.

As a measure of the buyer's overall feedback provision history, we use BuyerPcntFbLeft. BuyerPcntFbLeft is the proportion of feedbacks left to the number of buy orders, before the current period in the current subsession. A higher proportion of feedbacks given for past transactions indicates the buyer's willingness to give feedback, and should dissuade the sellers from cheating. Thus we expect that an increase in BuyerPcntFbLeft would lead to buyers experiencing higher trustworthiness.

In addition to these factors, the seller can also use his own experience with (other) buyers in previous rounds. We use RoundsSinceSellerReceivedFB as a measure of the seller's feedback experience in the market. RoundsSinceSellerReceivedFB is defined as the number of transactions before the current transaction, since the seller last received a feedback. For example, if RoundsSinceSellerReceivedFB=3 , it means that since the last time the seller received a feedback, there were three periods where he was able to sell, but did not receive any feedback. Controlling for the buyer's feedback giving characteristics, we expect that sellers who have not received a feedback recently would be more likely to cheat.

To examine these effects, we perform a logistic regression with the probability of seller choosing to ship as the dependent variable. The regressors on both these regressions can be divided into two groups. In the first group, we have dummy variables for subsession 2 (sub2) and subsession 3 (sub3). The second group contains our variables of interest, namely RoundsSinceBuyerGaveFB, BuyerPcntFbLeft, and RoundsSinceSellerReceivedFB. We use random effects for each seller to control for the repeated observations and individual idiosyncrasies. Analogous estimates using fixed-effects models are presented in Appendix III. Barring minor differences in magnitude of effects, the results with fixed-effects conform to the results from random-effects analysis.

OUTCOME: Seller Ships	**Low Cost**	**High Cost**
SUB2 Dummy variable for subsession 2	0.151	-0.345
	(0.383)	(0.280)
SUB3 Dummy variable for subsession 3	0.485	0.092
	(0.389)	(0.295)
BHPCNTFBSLEFT Proportion of feedbacks left in the current subsession prior to the current period	-0.494	-0.124
	(0.542)	(0.414)
BUYERSINCELASTFB Number of buy orders since the last time the buyer gave a feedback (before the current period)	**-0.228**	**-0.197**
	(0.078)***	(0.065)***
SELLERSINCELASTFB Number of times the seller sold since he last received a feedback	-0.030	**-0.236**
	(0.098)	(0.075)***
CONSTANT	**3.698**	**2.513**
	(0.664)***	(0.488)***
OBSERVATIONS	687	599
NUMBER OF UNIQUEID	24	24
Log pseudo-likelihood	-170.04	-266.94
Standard errors in parentheses *** significant at 10%; ** significant at 5%; *** significant at 1%**		

Table 4 Regression Results for Seller Choosing to Ship

As seen in Table 4, the coefficient on the subsession dummy variables is not significant for either of the cost conditions. Again, for both conditions, the coefficient on BuyerPcntFbLeft is not significant, while the coefficient on RoundsSinceBuyerGaveFB is negative and highly significant. This indicates that the sellers do discriminate among buyers, but based on recent, rather than overall feedback giving history. The longer a buyer goes without giving a feedback, the less likely the sellers are to choose to ship when dealing with her. The coefficient on RoundsSinceSellerReceivedFB is not significant for the low cost condition, but negative and significant for the high cost condition. This suggests that the sellers use their own prior feedback receiving experience in the high cost condition, but not in the low cost condition.

Even though we find evidence of discrimination in both cost conditions, the effect is more pronounced in the high cost condition. For example, in the low cost condition the model predicts that sellers will be 3% more likely (97% vs. 94%) to ship when faced with a buyer who has given feedback in her last transaction, as opposed to a buyer who did not give a feedback in her last three transactions. In the high condition, the model predicts a 7% difference (89% vs 82%) in probability of sellers choosing ship for the same difference in recent feedback giving history.

2.3.3 Should sellers discriminate among buyers based on their feedback giving history?

Having established that sellers practice discrimination based on buyers' recent feedback provision history, we now move on to the next questions; i.e. 'Is this seller discrimination rational?' To answer this question, we will need to answer the following two questions:

A. What is the effect of positive and negative feedback in a seller's reputation on buyers' trust towards him?

B. Are the buyers who gave less feedback in the past less likely to give feedback in the current transaction?

Let us consider the first question, about the relationship between feedback and trust. Cheating is profitable to the seller in the short run, but if the buyer gives a negative feedback, he risks losing out on profitable trade opportunities in future. We are not trying to evaluate whether the forfeited future trade gains are sufficient to make cheating strictly suboptimal. If we wanted to answer that, we would need to consider a number of things (the number of periods in the game after the negative, length of the punishment phase, stoning, etc.). Given the data, it is not possible to assess many of these things. Hence, we conduct the simpler analysis that considers only whether the seller's reputation has some impact on the buyers' trust in him. In our experimental setup a seller's revenue directly depends on buyers trusting him[9]. Thus, a necessary condition for a feedback to be effective is that a positive feedback should lead to greater trust and a negative feedback to reduced trust from future buyers.

To examine whether our experimental market meets this requirement, we perform a random effects logistic regression for each cost condition to estimate the probability of the buyer choosing to buy. The regression results are presented in Table 5. We use SHPosFB, SHNegFB and SHNoFB are from the seller's feedback profile, which are the number of transactions where the seller received a positive, negative and no feedback resp. To control for buyer's prior experience, we use BHNumNoGets which is the number of prior transactions in the current subsession where the buyer was cheated. BuyerLastCheat is an indicator variable that represents whether the buyer was cheated in the latest period when she chose to buy. As before, we also use dummy variables for subsessions.

---

[9] Revenue also depends on the seller's own ship action. But a seller only has that choice if the buyers decide to trust him.

For both cost conditions the regression coefficient on SHPosFB is positive and significant, while the coefficient on SHNegFB is negative and significant. Thus getting positive feedback would increase and a negative feedback would reduce the seller's ability to sell in future periods. For a seller who had three positive feedbacks, the presence of a single negative caused a drop of about 3% (98% vs. 95%) in the predicted probability of buyer choosing to buy. In the high cost condition, a single negative caused a drop of about 7% (95% vs. 88%) for a seller who had three positive feedbacks. Thus our experimental market meets the condition that a positive feedback is desirable and a negative feedback unattractive to sellers.

Comparing the magnitudes of the regression coefficients it appears that about three positive feedbacks are needed to counter the effect of one negative feedback in both cost conditions. For a seller with six positive feedbacks and one negative feedback, the predicted probability of buyer choosing to buy is almost the same as the predicted probability for a seller with 3 positives and 0 negatives.

OUTCOME: Buyer Chooses to Buy	Low Cost	High Cost
SUB2 Dummy variable for subsession 2	-0.179	-0.360
	(0.371)	(0.280)
SUB3 Dummy variable for subsession 3	0.406	-0.717
	(0.409)	(0.290)**
SHPOSFB Number of positive feedback received by the seller in the current subsession prior to the current period	**0.634**	**0.576**
	(0.156)***	(0.093)***
SHNEGFB Number of negative feedback received by the seller in the current subsession prior to the current period	**-2.219**	**-1.385**
	(0.232)***	(0.166)***
SHNOFB Number of prior transactions in the current subsession where the seller received no feedback	**-0.152**	0.050
	(0.077)**	(0.066)
BUYERLASTCHEAT Dummy variable indicating whether the buyer was cheated the last time she chose to buy	**-1.377**	**-1.135**
	(0.474)***	(0.273)***
BUYERNUMNOGETS Number of prior transactions in the current subsession where the buyer was cheated	-0.453	**-0.598**
	(0.319)	(0.161)***
CONSTANT	3.897	2.647
	(0.523)***	(0.395)***
OBSERVATIONS	792	792
NUMBER OF UNIQUEID	24	24
Log pseudo-likelihood	-170.52	-295.60
Standard errors in parentheses *** significant at 10%; ** significant at 5%; *** significant at 1%**		

Table 5 Regression Results for Buyer Choosing to Buy

We now move on to the next question, i.e. whether buyers who gave less feedback in the past are less likely to give feedback in the current transaction. We perform separate regressions to estimate the probability of giving a positive and a negative feedback. As in the earlier analyses, we use BuyerPcntFbLeft and RoundsSinceBuyerGaveFB as measures of buyers' feedback provision history, and dummy variables for subsessions. In the first two columns of Table 6 we present the results of the estimation for probability of negative feedback, and the next two columns we present the results for the probability of positive feedback. For estimating the probability of negative feedback, none of our regressors have a significant effect (except the dummy for second subsession in the high cost treatment). Thus the probability of a buyer giving a negative feedback is independent of her feedback provision history. For the probability of a positive feedback, the overall proportion of feedback given in the past in a good predictor. For both types of feedback, the recent feedback history is not a good predictor of the probability of feedback.

OUTCOME: Buyer Gives a Feedback	**Negative Feedback**		**Positive Feedback**	
	Low Cost	**High Cost**	**Low Cost**	**High Cost**
SUB2 Dummy variable for subsession 2	-0.523	**0.937**	-0.546	-0.211
	(1.783)	(0.569)*	(0.399)	(0.377)
SUB3 Dummy variable for subsession 3	-1.287	0.998	**-1.146**	-0.194
	(1.828)	(0.637)	(0.428)***	(0.413)
BHPCNTFBSLEFT Proportion of feedbacks left in the current subsession prior to the current period	1.740	0.633	**3.670**	**2.591**
	(2.970)	(1.018)	(0.809)***	(0.654)***
ROUNDSINCEBUYERGAVEFB Number of buy orders since the last time the buyer gave a feedback (before the current period)	-0.122	-0.053	0.002	-0.216
	(0.255)	(0.123)	(0.130)	(0.135)
CONSTANT	**4.550**	0.391	**-2.989**	**-1.713**
	(2.220)**	(0.811)	(0.731)***	(0.709)**
OBSERVATIONS	61	121	628	486
NUMBER OF UNIQUEID	24	24	24	24
Log pseudo-likelihood	-15.68	-69.35	-166.99	-166.07

Standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 6 Regression Results for Buyer Giving a Feedback

2.3.4 Are buyers who get discriminated against due to their reticence, more likely to give a positive feedback thereafter?

The environments with buyer feedback provision history will have desirable performance characteristics if the threat of being discriminated against spurs the buyers into providing enough feedback so that cheating is unattractive for sellers. This outcome can only materialize if all the buyers understand that their feedback giving history matters, and provide enough feedback regularly to maintain a reputation for feedback provision. To acquire a feedback giving reputation, the type of feedback, i.e. positive or negative, does not matter. If the market has high average level of trustworthiness, there will be more opportunities for giving a positive feedback than for a negative feedback. To avoid being discriminated buyers must give positive feedback from time to time. Suppose a buyer does not understand this at first, and as a result gets cheated. She might retaliate to the cheating by giving a negative feedback. However, this may not be sufficient to avoid getting cheated in future. To acquire a reputation for providing feedback, she would need to give positive feedback more often. We examine whether getting cheated highlights the importance of feedback giving history, and makes the buyers more likely to give positive feedback.

We perform a random-effects logistic regression to estimate the probability of the buyer giving a positive feedback (conditional on seller choosing to ship). We use a binary indicator variable 'Cheated' to indicate whether a buyer has already been cheated in the current subsession. If the buyer gets cheated for the first time in a subsession in the n^{th} period, then cheated is set to 0 for the first 'n-1' periods, and it is set equal to 1 starting period 'n+1'[10]. In model 1, we also use BuyerPcntFbLeft, RoundsSinceBuyerGaveFB and the dummy variables for subsessions as in the earlier regressions. The regression results are presented in the first two columns of Table 7. For both cost conditions, the coefficient on BuyerPcntFbLeft is positive and significant, suggesting that buyers continue to follow their previous patterns. In this high cost condition the coefficient on Cheated is not significant. In the low cost condition the coefficient on Cheated is negative and significant, indicating that after getting cheated, buyers were even less likely to provide positive feedback after a good transaction.

[10] We are estimating the probability of getting a positive feedback, thus the n^{th} period where the buyer was cheated is not included in the regression.

OUTCOME: Buyer Gives a Positive Feedback	Model 1		Model 2	
	Low Cost	High Cost	Low Cost	High Cost
SUB2 Dummy variable for subsession 2	-0.627 (0.402)	-0.231 (0.378)	-0.636 (0.401)	-0.212 (0.376)
SUB3 Dummy variable for subsession 3	**-1.366** (0.455)***	-0.179 (0.416)	**-1.366** (0.444)***	-0.098 (0.416)
CHEATED Binary variable indicating whether the buyer was cheated previously in the current subsession	-0.758 (0.384)**	0.304 (0.345)	-0.579 (0.419)	0.386 (0.398)
BHPCNTFBSLEFT Proportion of feedbacks left in the current subsession prior to the current period	**3.890** (0.817)***	**2.715** (0.675)***	**3.865** (0.800)***	**3.330** (0.674)***
ROUNDSINCEBUYERGAVEFB Number of buy orders since the last time the buyer gave a feedback (before the current period)	0.008 (0.132)	-0.197 (0.135)		
QUIETRECENT Binary variable indicating that the buyer did not provide a feedback the last time she bought			0.190 (0.586)	0.001 (0.595)
CHEATEDXQUIETRECENT Interaction of Cheated and quietRecent			-0.966 (0.945)	0.011 (0.730)
CONSTANT	**-2.726** (0.735)***	**-1.955** (0.767)**	**-2.737** (0.722)***	**-2.564** (0.770)***
OBSERVATIONS	628	486	628	486
NUMBER OF UNIQUEID	24	24	24	24
Log pseudo-likelihood	-164.98	-165.68	-164.39	-166.86

Standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 7 Regression Results for positive feedback

We use an alternative approach in Model 2, the results of which are presented in columns 3 & 4 on Table 7. We've seen earlier that sellers discriminate among buyers based on lack of recent feedback. In Model 1, we used RoundsSinceBuyerGaveFB as an indicator of the recent feedback provision history. In Model 2, we use a binary indicator variable quietRecent to represent the same. If a buyer has not provided a feedback in the last period that she chose to buy, she is considered to be quiet for that period, and quietRecent is set to 1 her until she gives a feedback the next time[11]. We also interaction term between Cheated and quietRecent to indicate when a buyer was quiet and cheated. The regression results again find no significant effect of Cheated, quietRecent or the interaction term. Thus, it appears that getting cheated does not cause the buyers to give a positive feedback more often, at least in the same subsession.

Recall that our experimental session was divided into three subsessions, with the buyer and seller histories being cleared after each subsession. Now we examine whether getting cheated in one subsession increases the probability of buyer giving a positive feedback including in the subsequent subsessions. For this purpose we use an indicator variable CheatedEver, which indicates whether a buyer has been cheated before the present transaction, including in previous subsessions. Like in previous regressions, we also include dummy variables for subsessions, BuyerPcntFbLeft, and RoundsSinceBuyerGaveFB. The regression results are presented in Table 8. Again, the only significant coefficient is on BuyerPcntFbLeft. Getting cheated in an earlier subsession does not make a buyer more likely to give a positive feedback in subsequent subsessions. Overall, we see a trend that buyers who have given more feedback previously continue to be more likely to give feedback again in the future. This would suggest that there are inherent differences in the subjects' propensity to give positive feedback.

---

[11] quietRecent = 1 if BuyerSinceLastFb > 0.

OUTCOME: Buyer Gives a Positive Feedback	Low Cost	High Cost
SUB2 Dummy variable for subsession 2	-0.192 (0.425)	-0.591 (0.551)
SUB3 Dummy variable for subsession 3	-0.527 (0.494)	-0.665 (0.646)
CHEATEDEVER Binary indicator variable indicating whether the buyer was cheated previously including in earlier subsessions	**-1.391** (0.617)**	0.566 (0.588)
BHPCNTFBSLEFT Proportion of feedbacks left in the current subsession prior to the current period	**3.791** (0.825)***	**2.496** (0.664)***
BUYERSINCELASTFB Number of buy orders since the last time the buyer gave a feedback (before the current period)	0.022 (0.135)	-0.214 (0.133)
CONSTANT	**-2.977** (0.766)***	**-1.720** (0.714)**
OBSERVATIONS	628	486
NUMBER OF UNIQUEID	24	24
Log pseudo-likelihood	-164.19	**-165.60**
Standard errors in parentheses * significant at 10%; ** significant at 5%; *** significant at 1%		

Table 8 Regression Results including cheating from earlier subsessions

2.3.5 Do environments with buyer feedback history have higher levels of trust, trustworthiness and efficiency?

To compare the performances of the treatments with and without buyer feedback provision history, we use the Wilcoxon-Mann-Whitney rank test. To satisfy the independence requirement in the rank test, we use an individual buyer as a unit of observation when considering trust, an individual seller as a unit of observation when considering trustworthiness and a subsession as the unit of observation when considering efficiency. The results for the trading game are presented in Table 9. We find no significant differences in any of the performance characteristics for either of the cost conditions.

Performance Metric	Cost	Null Hypothesis	z-score	Prob > |z|	Number of Observations
Trust	Low	Trust(NoHist) = Trust(fbHist)	-0.052	0.9586	48 (24 in each condition)
Trust	High	Trust(NoHist) = Trust(fbHist)	-0.197	0.8441	48 (24 in each condition)
Trustworthiness	Low	Trustworthiness(NoHist) = Trustworthiness(fbHist)	-0.76	0.447	48 (24 in each condition)
Trustworthiness	High	Trustworthiness(NoHist) = Trustworthiness(fbHist)	1.469	0.1419	48 (24 in each condition)
Efficiency	Low	Efficiency(NoHist) = Efficiency(fbHist)	0.174	0.8618	24 (12 in each condition)
Efficiency	High	Efficiency(NoHist) = Efficiency(fbHist)	0.607	0.5438	24(12 in each condition)

Table 9 Trading Game Performance: Results from Wilcoxon Rank Test

In Table 10, we present results from the evaluation game. To satisfy the independence requirement, we use an individual buyer as the unit of observation. When considering all feedback (positive or negative), we find significantly higher incidence of feedback in the treatment with buyer history in the high cost condition, but no differences in the low cost condition. When considering only positive feedback, the buyer history treatment has significantly more positive feedback in the high cost condition, but there are no significant differences in the low cost condition. When considering only negative feedback, we find that the buyer history treatment has significantly more feedback in the low cost treatment, but there are no differences in the high cost treatment.

Performance Metric	Cost	Null Hypothesis	z-score	Prob > |z|	Number of Observations
Feedback	Low	ProbFb(NoHist) = ProbFb(fbHist)	0.392	0.6952	48 (24 in each condition)
Feedback	High	ProbFb(NoHist) = ProbFb(fbHist)	**-2.691**	**0.0071**	48 (24 in each condition)
Positive Feedback	Low	ProbPosFb(NoHist) = ProbPosFb(fbHist)	0.568	0.5697	48 (24 in each condition)
Positive Feedback	High	ProbPosFb(NoHist) = ProbPosFb(fbHist)	**-2.251**	**0.0244**	48 (24 in each condition)
Negative Feedback	Low	ProbNegFb(NoHist) = ProbNegFb(fbHist)	**-1.967**	**0.0492**	47 (23 and 24 resp.)[a]
Negative Feedback	High	ProbNegFb(NoHist) = ProbNegFb(fbHist)	0.306	0.7597	46 (22 and 24 resp.)[b]

a. One subject in treatment without buyer history never got a chance to give a negative feedback
b. Two subjects in treatment without buyer history never got chance to give a negative feedback

Table 10 Feedback Provision: Results from Wilcoxon Rank Test

Summarizing, we find the effects of buyer feedback provision history to be a mixed bag. We see in Table 10 that the buyer history is able to elicit higher levels of feedback at least in some conditions. Yet, higher incidence of feedback does not translate into better trading game performance characteristics. We discuss the results in the next section, suggesting possible future extensions.

2.4 Discussion

We find evidence that the sellers practice discrimination based on buyers' recent feedback provision history. They are more likely to cheat buyers who have not given a feedback recently. However, such discrimination seems unwarranted. Recent feedback history was not found to be a good predictor of the incidence of either type of feedback (positive or negative). In the high cost condition, we found that buyers were more likely to provide feedback when history was made visible. We found no evidence, however, to suggest that the buyers who got cheated understood they were being punished for their reticence. They did not increase their feedback provision on subsequent successful transactions. The net result was that the availability of feedback provision history did not affect the trading game performance on any of our three metrics, namely trust, trustworthiness and percentage of successful transactions.

There were three deviations from the rational behavior expected in the equilibrium of high feedback provision and efficient trading outcomes that are worth explaining in more depth, to understand why they might have occurred and to understand their impacts on overall outcomes. Some of the impacts of the non-rational behavior were detrimental to the intended incentives of making histories visible, while others helped to maintain the desired outcomes of trust and trustworthiness.

The first deviation is that buyers were more likely to provide negative feedback than positive even in the buyer history conditions. Averaging across the three subsessions, 87% of our buyers gave a negative feedback at each available opportunity in the low cost treatment, and 60% of the buyers gave a negative at each available opportunity in the high cost treatment. For positive feedback these numbers were considerably lower, 12% in the low cost treatment and 21% in the high cost treatment. Of course, considering that there were more opportunities to give a positive feedback, we would not expect the buyers to always give a positive feedback. A more striking statistic is how many subjects never gave a positive feedback: 40% in the high-cost treatment and 52% in the low-cost treatment. The buyers showed a higher propensity for giving a negative feedback, even though this was not the best way to acquire a feedback giving reputation. Both types of feedback were equivalent for that purpose since sellers were not shown the valence of feedbacks in buyers' history. There are two possible explanations for this result.

One explanation is that buyers gave negative feedback out of non-monetary motives, such as altruism or joy of retaliation. This explanation is consistent with other experimental literature (Fehr and Gachter 2000), which finds that people are willing to punish violators of trust, even when such punishments are costly and yield them no direct benefit. This explanation is also consistent with the outcome that negative feedback provision was also higher in the absence of feedback giving histories. As shown in Table 3, the prevalence of negative feedback is quite high in all four conditions. By chi-square tests, the difference between positive and negative feedback provision is highly significant in all four conditions.

A second possible explanation is that some buyers misinterpreted the feedback giving history and thought that future sellers would know whether the feedback given was positive or negative, or in case they gave no feedback, whether a transaction was good or bad. Thus, they might have thought it more important to establish a reputation for giving a negative feedback. If this were the case, we should find, in both cost conditions, a higher incidence of negative feedback when feedback histories were available. However, we find mixed evidence, i.e. buyer history leads to higher negative feedback (91% vs. 69%) in the low cost condition, but not in the high cost condition (70% vs. 73%).

The effect of buyers being more likely to give negative feedback than positive is that a history with few feedbacks may not be a good indicator that a buyer will not provide a negative feedback, and thus it is not safe for sellers to cheat quiet buyers.

Indeed, Table 6 showed that feedback histories were significant predictors of feedback after positive transactions, but not after cheating by the seller. If sellers realized this, the incentive effect of history for buyers to provide feedback at all would dissipate.

The second deviation from rational self interested behavior is that sellers discriminated against buyers who did not give a feedback in the recent past, even though feedback giving histories were not useful to predict whether buyers would provide a negative feedback when cheated. As Table 4 showed, the more transactions since a buyer last gave feedback, the more likely the seller was to cheat her. Thus, we conclude that sellers did not realize that histories were an unreliable indicator of willingness to provide negative feedback. Ironically, this deviation from rationality on the sellers' part compensates for the first deviation on buyers' part, and allows the system to still provide buyers with an incentive to give feedback.

The third deviation from rationality is that buyers who were discriminated against didn't start giving more positive feedback. It appears that at least in the high cost treatment, feedback histories caused many buyers to make a strategic adjustment: From Table 3, we see that positive feedback provision was 39% with feedback histories, and only 16% without. Those buyers, who did not make a strategic adjustment initially, however, seem not to have made the adjustment even after getting cheated as analyzed in Table 7. The effect of this off-equilibrium behavior is to dissipate some of the improvements in trading game efficiency. Sellers cheated reticent buyers, erasing the gains that were realized from more trustworthy behavior with other buyers.

There were two other interesting puzzles in the results. One is that presence of buyer feedback giving history was found to affect the positive feedback rate only in the high cost treatment (39% vs. 16%), but not in the low cost treatment (30% vs. 32%). One possible explanation for this finding comes from signaling theory. When feedback was inexpensive as in the low cost condition, buyers may not have felt the need to signal their willingness to give feedback. Perhaps they assumed that with the low cost of feedback, sellers would fear that buyers would provide negative feedback simply to get revenge. On the other hand, when feedback cost more, at least some buyers may have felt the need to signal their willingness to bear such costs, and thus provided positive feedback more often.

The higher incidence of negative feedback in the presence of buyer history (91% vs. 69%) only in the low cost treatment, but not in the high cost treatment (70% vs. 73%) is a more puzzling result. We considered a number of possible reasons, but why this phenomenon should occur only in the low cost condition and not the high cost condition, still remains a mystery.

2.5 Conclusion

We conclude that making buyers' feedback giving history available to sellers affects the behaviors of both buyers and sellers. Buyers' feedback giving history allows sellers to target certain buyers for untrustworthy behavior, a strategy unavailable to them otherwise. Sellers discriminate against buyers based on their recent feedback provision history. However, this discrimination does not yield the sellers the desired benefits. A buyer's feedback giving history is not a good predictor for the probability of getting a negative feedback from her. Sellers frequently end up getting a negative feedback and losing out on future benefits. We find that at least some buyers do not understand that their feedback giving history matters. Even those that understand that feedback giving history matters are not able to predict what metrics the sellers use to discriminate. On the whole, it is more difficult for the buyers to comprehend the reputation effects than the sellers. One possible solution to this problem is to not give the sellers access to buyers' feedback giving history. However, we think it is too early to give up on the potential gains in performance that can be achieved if all sellers and buyers understand and play according to the equilibrium strategy. An alternative solution would be to educate the buyers by making them aware of the importance of feedback giving history. Such an education campaign could take several forms: it could be a tutorial at the time of registration, or periodic messages reminding them that feedback history matters. Future research could examine the practical implications of such an education campaign on the behavior of buyers and sellers. If such a campaign is effective, it could reduce the amount of exploitation of reticent buyers, while simultaneously increasing the total feedback provision, leading to improved trading outcomes.

Chapter 3

Self-Selection, Slipping, Salvaging, Slacking, and Stoning: the Impacts of Negative Feedback at eBay

3.1 Introduction

The Internet brings together buyers and sellers separated by physical boundaries, opening endless avenues for trade (Bakos 1997). With the expanding reach of the Internet, electronic markets are emerging as an increasingly important factor of the economy. eBay, touted as 'The World's Online Marketplace' is among the most noteworthy examples of electronic markets.

In a trade environment like eBay, buyers have limited information about product quality and seller reliability, at the time of a transaction. Thus, electronic markets like eBay are "ripe with the possibility of large-scale fraud and deceit"(Kollock 1999). The information asymmetry between buyers and sellers can reduce market efficiency. Most directly, risk averse buyers may miss beneficial transactions. In addition, adverse selection may lead high quality sellers to abandon the market, creating a 'Lemons Market" as modeled by Akerlof (Akerlof 1970). Buyers are unwilling to pay the full price of the high quality good because they are unable to distinguish those goods from lower quality goods. Thus, the sellers of high quality goods are unable to get the full value of their items, and choose to leave the market.

The Feedback Forum at eBay, where buyers and sellers leave feedback about each other, reduces the information asymmetry, by telling buyers whether previous customers were satisfied with the seller. Both buyers and sellers can leave post-transaction feedback about each other; each feedback contains a numerical rating (+1, 0, or –1), and a text comment.

In principle, reputation systems like the Feedback Forum can improve the efficiency of marketplaces in three ways:

Signaling: A seller's feedback history can serve as a signal to buyers of how risky it is to purchase from that seller. This allows each buyer to choose which sellers to buy from, and how much to bid, based on the buyer's level of risk aversion.

Sanctioning effect: Sellers will strive to avoid negative feedback, in order to avoid adverse impact on their future sales.

Selection effect: Because buyers will be better able to distinguish high quality from low quality, high quality sellers will not leave the market. Indeed, the low quality and fraudulent sellers may be driven from the market, leaving a higher overall quality level and less risk even for those buyers who do not carefully monitor the signals about the trustworthiness of individual buyers.

Prior research has tried to document the extent of some of these effects. One study provided a model showing that past feedback is somewhat useful as a signal predicting future customer satisfaction (Resnick and Zeckhauser 2002). Many studies have examined the sanctioning effect by quantifying the impact of a seller's feedback profile on the probability that an item will sell and the price it will receive. Generally, these employ a cross-sectional methodology, comparing sales between sellers and including variables in regression models to try to control for differences among sellers and listings. One study (Resnick, Zeckhauser et al. 2006) employs a field experiment that compares sales by a single seller using different personas with different feedback profiles. See Bajari (Bajari and Hortacsu 2004), and Resnick et. al. (Resnick, Zeckhauser et al. 2006) for a survey of results. In general, the studies of the sanctioning effects suggest that a better reputation is of some value to a seller, though the studies differ somewhat in their assessments of the particular impacts positive and negative feedback in a profile.

This paper analyzes other aspects of how seller and buyer behavior change as seller feedback profiles change, other than buyer bidding behavior. In particular, we analyze self-selection by sellers, changes in seller quality levels, and changes in buyer feedback giving patterns. Our approach is based on analysis of a large panel of sellers. We track these sellers after they get a negative feedback, and examine whether they continue to list items, and when they do, whether they get positive feedback, negative feedback, or no feedback at all.

When a seller gets a negative feedback, it marks a significant change in a feedback profile. An earlier study found that neutral and negative feedback constituted only about 1% of all feedback (Resnick and Zeckhauser 2002). Anecdotal evidence from observation of eBay discussion boards and attendance at a conference of eBay users suggests that getting a negative feedback is a traumatic event for many.

Thus, we contrast what happens before receiving a negative feedback with what happens afterward. Seller behavior, both before and after receiving negative feedback, is reflected in whether and how frequently they list items. Buyer behavior is reflected in whether listed items sell. Both buyer and seller behavior are reflected in the feedback that buyers give to sellers. For example, a higher probability of negative feedback may reflect worse performance by the sellers, or it may reflect a greater willingness of buyers to report dissatisfaction, or some combination of those.

Our analysis is most similar in spirit to that in Cabral & Hortacsu (Cabral and Hortaçsu 2004). We ask many of the same questions but there are major differences in our analytical approach and we reach quite different conclusions.

3.2 The Impact of Negative Feedback

A negative feedback may result from random factors, or it may be an indicator that the quality of the seller's goods or services has declined. We will refer to such a quality decline as *slipping*.

Rather than just reflecting a change in the seller's quality, negative feedback can cause changes. The seller's own behavior may change and the community's attitude towards her may also change. Let us examine each of these in detail.

3.2.1 Seller Behavior Change

The seller may stop listing items entirely, what we will refer to as *self-selection*. This would be a sign that the system is working as intended to weed out sellers of lower quality, though it could be that some high-quality sellers who get a negative feedback as a result of a genuine misunderstanding are also being driven from the system.

A seller's quality also may change as a result of receiving a negative feedback. The seller may be more careful and conservative in describing her items, more responsive in her communication, and faster and more careful in her packaging and shipping, in an effort to *salvage* her reputation. Contrariwise, sellers may deliberately offer lower quality service after receiving a negative feedback, a practice that we will refer to as *slacking*. Like slipping, slacking will be observable in the form of less positive feedback on future transactions and more negative feedback. The difference is that, with slipping, the quality decline takes place around the time of the transaction that eventually gets negative feedback, while with slacking it happens only after the seller receives the negative feedback.

Slacking could occur for psychological reasons, if sellers get angry or discouraged after receiving a negative feedback. It could also occur as part of equilibrium behavior that creates an incentive for sellers not to get negative feedback in the first place. Cabral and Hortacsu (Cabral and Hortaçsu 2004), following Diamond (Diamond 1989), construct an equilibrium for a model where getting negative feedback provides evidence that the seller is of an opportunistic type rather than a type that always gives high effort. Once the seller has been found out as not always giving high effort, buyers expect the seller not to give high effort in the future, and it is rational for the sellers to act in accordance with that buyer belief. Whether slacking occurs for psychological reasons or as a way to maintain an equilibrium of high effort prior to receiving a negative, if we find evidence of slacking, we would consider it a negative consequence of the design of the feedback system, as it would reflect sellers offering lower quality than they were capable of after receiving a negative feedback.

3.2.2 Buyer Behavior Change

Clearly, a negative feedback reflects unfavorably on the seller. A negative feedback may result from factors beyond the control of the seller, such as a package lost in the mail. Still, natural Bayesian updating of beliefs will lead buyers to be more suspicious of a seller after she has received a negative feedback.

As mentioned earlier, our analysis will focus on how buyers' increased suspicions affect their feedback giving behavior rather than their bidding behavior. Resnick and Zeckhauser (Resnick and Zeckhauser 2002) suggest that buyers would be more willing to cast another stone at an already disreputable seller, a phenomena they call a *stoning* effect. That is, receiving a negative feedback would make the user more likely to receive another. This could occur for one of two reasons. First, buyers may be willing to forgive a single bad behavior but want to punish sellers who exhibit a pattern of bad behavior. Second, a buyer may interpret what happened in her own transaction differently depending on the suspicions raised by the seller's previous feedback. For example, if an item appears to be damaged in shipment, a previous negative feedback suggests that the damage was more likely to have been the seller's fault.

Dellarocas (Dellarocas 2001) shows that a reputation system can create the same sanctioning effect no matter how much of the seller's history buyers see. When only feedback for the most recent n transactions are displayed (or buyers only pay attention to the recent feedback), buyers can only sanction sellers for up to n rounds after they receive a negative feedback, but if the per-round punishment in terms of lost revenue is sufficiently high, the sellers will be deterred from getting a negative feedback. If the per-round punishment is limited, however, then it will be necessary to extend sanctioning over a longer period of time. Stoning can be viewed as a strategy that probabilistically extends a seller's sanctioning period after receiving negative feedback—even if many buyers are paying attention only to a seller's most recent feedbacks, because of stoning there is an increased probability of getting another negative feedback, which keeps a negative feedback in the "recent history" for longer.

Messages on eBay's message boards suggest that sellers believe buyer stoning to be a common phenomenon. On August 22, 2004, one user wrote:

> *"I have been buying and selling on eBay for almost four years and have over 1000 positive feedbacks with four negs and six neutrals. Something that I have noticed, is that the Negs/Neutrals seem to happen in spurts. I think that people - especially buyers - are reluctant to leave a Neg or Neutral for a seller who has an excellent feedback score. Likewise, they are more apt to leave the Neg or Neutral if they see that others may have recently left one as well..."*

Another veteran eBay seller replied:

"Goes without saying that period following a non-positive feedback is at high risk of getting another."

If stoning occurs, the probability of negative feedback should go up for transactions after the seller receives a negative feedback. When a transaction goes well, however, we hypothesize that the presence of a negative feedback should have no effect on whether the buyer provides a positive feedback or no feedback at all. Thus, unlike seller slipping and slacking, stoning should increase the probability that a seller receives negative feedback, but should not affect the probability that a seller receives positive feedback. Table 11 below contains a summary of the various effects of negative feedback that we'll be exploring in this paper.

Effect	Party	Description
Self-selection	Seller	Drop out after negative feedback
Slipping	Seller	Quality declines; that decline leads to negative feedback
Salvaging	Seller	Quality increases after negative feedback
Slacking	Seller	Quality declines after receiving negative feedback
Stoning	Buyer	More willing to give negative feedback

Table 11 Summary of possible effects to be explored

3.3 Data Analysis

For the analyses conducted in this paper, we selected panels of sellers and examined their transactions and feedback. The use of a panel dataset is important as it enables us to study the changes in individual users, and also account for user heterogeneity. The datasets we analyzed were derived from the following master datasets, provided by eBay in a form that stripped all personally identifiable details.

1) Items Dataset: contains transactional data for all the items listed for sale on eBay from February 1st 1999 to June 30th 1999.

2) Feedback Dataset: contains all the feedback data up to May 31st 1999.

3) Users Dataset: contains the id and registration dates for all the users who registered before June 30th 1999.

The participants' feedback profiles as of the times of transactions were not stored by eBay, but we were able to reconstruct measures of prior feedback similar to those eBay displayed to users. eBay calculates information in its feedback profiles in two ways. The first and perhaps more intuitive one, treats the feedback as the unit of analysis. A feedback profile includes a count of the total number of positive feedbacks and the total number of negative feedbacks in a user's profile. This accounting method, however, makes it relatively easy to inflate a reputation score, by having a friend leave multiple positive feedbacks.

A second accounting metric treats the partnership as the unit of analysis. At most one positive feedback from each partner counts toward the seller's count of distinct "members who left positive feedback" and similarly for neutral and negative feedback. For example, if user A sold 5 items to user B and gave a positive feedback for all the items, it would be counted only once. The composite score that eBay displays next to a user's id at various places in the site, including on auction listings, is the difference between the number of transaction partners who left positive feedback and the number who left negative feedback. For our analyses that use measures of a seller's prior feedback history at the time of a transaction, we rely on the metrics that treat the partnership as the unit of analysis.

Analysis of the content of neutrals and negatives showed that both are used primarily to indicate problematic transactions (Resnick and Zeckhauser 2002). Thus, we treat neutral feedback as negative for the purpose of our analyses.

3.3.1 Selection Effect

First let us examine whether sellers drop out when they get bad reputations, either not selling any more on eBay or switching to a new user id and starting over without any feedback. Of course, there will always be some attrition of users not continuing to sell on eBay, regardless of the state of their feedback profile. The question is whether attrition is higher among users with worse feedback profiles. To answer this question, we examined a panel of 76,956 users who joined eBay on or after February 1, 1999, received a feedback (as either buyer or seller) in the period April 11-30, 1999, and who had sold at least one item prior to receiving that feedback.

One indicator of a selection effect is whether the last feedback of a user was positive or negative. In the panel, 6.42% of users who listed no items after April 30 had a negative or neutral as their last feedback during April 11-30, while the overall percentage of negative (and neutral) feedback was only 1.37%[12]. This suggests some selection effect: the higher probability of negative feedback in last transactions indicates that users were more likely to stop selling after a negative feedback than after some other feedback.[13]

A more direct measure of the impact of feedback on whether users drop out comes from analysis of whether individual users sold again after receiving positive or negative feedback. To avoid statistical complications from repeated, overlapping measures, we randomly selected one feedback event for each user from the April 11-30 period: 2,091 were negative (or neutral), 74,865 positive. When the randomly selected feedback was negative, 53.85% of the users listed another item before June 30th. For those users whose randomly selected feedback was positive, 82.35% listed another item before June 30th. This suggests quite a large selection effect based solely on receiving a single negative feedback.

---

[12] The frequency of negative feedback was higher for this panel than in other datasets we consider in this paper because the sample is restricted to relatively new sellers, who tend to get somewhat more negative feedback.

[13] Cabral and Hortacsu [4] interpret similar data about higher percentages of negative feedback just before a seller drops out as evidence of seller profit-taking in advance of dropping out, rather than a decision after receiving negative feedback to drop out.

There may be additional effects from a user's feedback profile beyond the impact of the most recent feedback. If we look into the user's history, are sellers who received a negative feedback recently more likely to drop out? Are sellers who have received more negative feedbacks more likely to drop out, regardless of the content of the most recent feedbacks? Are sellers who have received more positive feedbacks less likely to drop out, either because they value their accumulated reputations, or because the large amount of positive feedback is an indicator of sellers who are more committed to eBay?

To test for these other indicators of selection effects, and to see whether the impact of a negative in the most recent transaction is still strong when controlling for the cumulative effect of the seller's full reputation, we conducted a logistic regression. The dataset is the random selection of one feedback for each of the 76,956 sellers. The outcome variable is whether the user listed another item. The covariates are as follows.

- *fbscore* is the score of the current feedback, 1 if positive, else 0.
- *Posr* is the number of distinct partners who gave positive feedback prior to this transaction's closing time. As is customary in other empirical analyses, we employ a log transform on the number of positive feedbacks: we expect the marginal impact of another positive feedback to decline as the user accumulates more feedback.
- *Negr* is the number of distinct partners who gave negative or neutral feedback prior to this transaction's closing time.
- *Neg5* is 1 if there is at least one negative feedback among the seller's five most recent feedbacks as of the time this feedback. It is 0 otherwise.

Table 12 contains results of the logistic regression predicting whether a seller lists an item for sale after receiving a negative feedback.

OUTCOME: Seller lists another item	
Logposr Log of seller's number of prior positive feedbacks	0.036 (0.009)***
Negr Seller's number of prior negative feedbacks	-0.120 (0.034)***
Fbscore Feedback score for the current feedback	1.294 (0.046)***
Neg5 Dummy variable indicating whether the last 5 feedbacks contained a negative	-0.555 (0.066)***
CONSTANT	0.220 (0.045)***
OBSERVATIONS	76,956
Pseudo-R2	0.013
Standard errors in parentheses *** significant at 10%; ** significant at 5%; *** significant at 1%**	

Table 12 Results of logistic regression predicting whether the user lists again

We find that the seller's earlier positive feedback has a very small but positive impact on her probability of returning as a seller. There is a decrease in the probability of return with prior negative feedback, supporting the claim that there is a selection effect from accumulated negative feedback. A recent negative feedback has a larger effect. Controlling for long-term and recent history, the feedback on the current item is still a large and significant predictor of whether the seller will return for another transaction.

To understand the effect of the covariates, consider a hypothetical user with 50 positive feedbacks. If the user has no negative feedback, the probability of return is 84%.

If the user has one negative feedback, which is not recent, the probability of return is 82%. If the user has one negative feedback which is relatively recent, but not for the last transaction, the probability of return is 73%. If the user has one negative feedback, and it is for the most recent transaction, the probability of return is 59%.

Thus, the strongest selection effect comes immediately after receiving a negative feedback, although there is a smaller ongoing, cumulative effect. We will return in the discussion section to the implications of the primacy of the most recent negative feedback and the comparatively smaller impact of accumulated positive feedback.

3.3.2 Salvaging, Slacking and Stoning Effects

In this section we look at the impact of negative feedback on the behavior of sellers who continue selling and the impact on the feedback giving practices of buyers. The analysis is again based on the transaction and feedback histories of a panel of users, in this case a panel of 9,655 users who sold at least one item meeting criteria to be described below.

One approach would be to consider only sellers who received negative feedback and to directly compare behavior before and after they received their first negative feedback. For example, Cabral and Hortacsu (Cabral and Hortaçsu 2004) show for a different dataset that sellers conducted more transactions before their first negative feedback than between receiving their first and second negative feedback.

We are concerned, however, that such an analysis would be biased, because either the dependent variable is censored or the sample is truncated. This is clearest in the analysis of number of transactions as a proxy for probability of getting negatives. If users with at least one feedback are the sample, the dependent variable for number of transactions after the negative is censored for those users who did not receive a second feedback. They would have conducted an unknown number of additional transactions before getting a negative. On the other hand, the sample of users with at least two negatives is a truncated sample--it leaves out those users who did not receive a second negative, and thus would have had, on average, more transactions after the first negative and before the second than those in the truncated sample.

Our approach is to conduct logistic regressions with the feedback on each transaction as the dependent measure. We compare outcomes for transactions that are in windows of time where the hypothesized slipping, slacking/salvaging, and stoning effects would occur. The logistic regressions also control for a seller's long-term feedback history.

The partnership is the unit of analysis. To avoid confounds from multiple transactions for the same partnership (e.g., partners may be less likely to leave second feedbacks since they know that eBay only counts one in its partner-based statistics) we analyze the outcome of only of the first transaction for each partnership. We exclude from the analysis partnerships where both parties joined eBay prior to the beginning of our transactions dataset, as we are unable to determine if the first transaction in our dataset was truly the first one for the partnership. We consider only partnerships where the user in our sample was the seller in their first transaction, since the hypothesized effects such as slipping apply more clearly to sellers.

During the time period of our data, eBay did not require users to tie feedback to particular transactions, though they had the option to do so. Even when a buyer explicitly specified a transaction, if there were multiple transactions between the partners in a short time, we think that a feedback for any one transaction was often intended to cover them all. Thus, for each partnership's first transaction, we classified the outcome based on the first feedback recorded from buyer to seller within 42 days (six weeks) of transaction closing time.

We have feedback data only through May 1999. The transactions later in the dataset are less likely to have a feedback recorded within our dataset. To avoid any truncation effect, we consider only transactions that ended on or before April 10, 1999, i.e., 42 days before the end of the dataset.

To summarize, we analyze the feedback outcomes for first transactions within each partnership, where the user from the panel was the seller in the first transaction, where at least one of the partners joined eBay after February 1, and where the first transaction closed some time in the period Feb. 1- April 30. Of the 100,761 transactions meeting these criteria, 42.34% resulted in positive feedback, 0.42% in negative (or neutral) feedback, and the remaining 57.23% received no feedback.

Slipping, stoning, and slacking or salvaging behavior all would be reflected in the outcomes of transactions near those transactions that are known to have received negative feedback. To differentiate among the effects, we define more carefully the windows around those transactions where we would expect to see the effects.

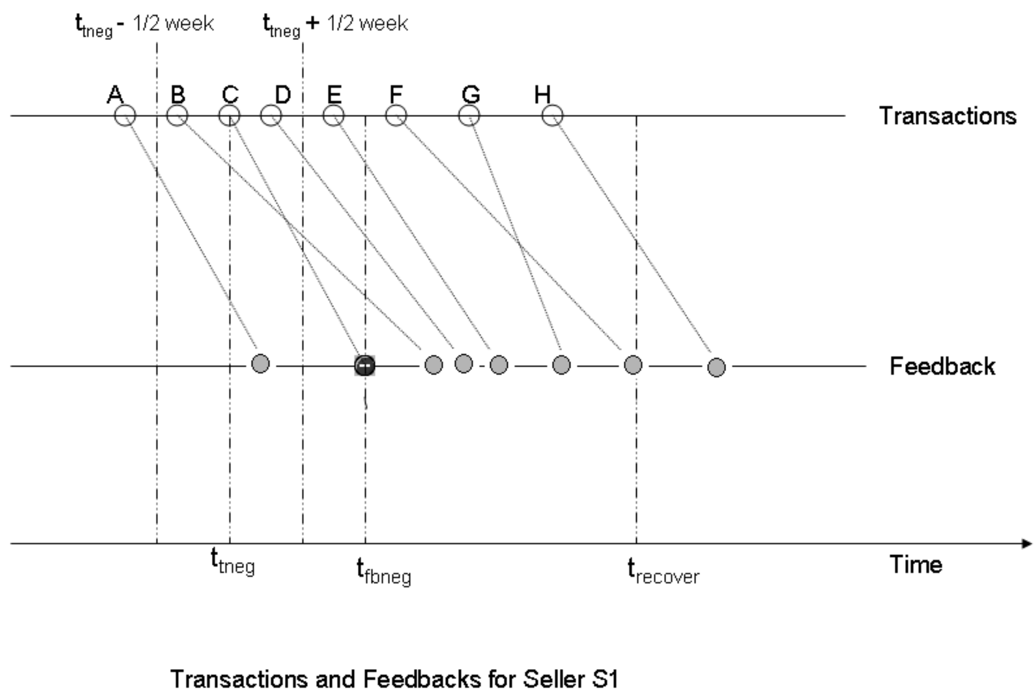

Figure 17 Windows generated from negative feedback for transaction C

Figure 17 illustrates the definition of these windows for a single transaction that received negative feedback in seller S1's history. The transactions are shown from left to right in order of their closing times. The feedbacks received by the user are displayed along the same timeline, but arranged according to the time the feedback was received. Lines connect feedbacks with the transactions they comment on. Transaction C received a negative feedback. We classify all other transactions as to whether they are in transaction C's slipping, slacking/salvaging, and/or stoning windows, in order to assess whether being in these windows affects the probability of receiving negative feedback.

The slipping window is defined for transactions that closed shortly before or after the known bad transaction (C in Figure 17). The idea is that a transaction known to get a negative may be indicative of a temporary decline in the seller's quality, perhaps due to a vacation or illness or family crisis. We (somewhat arbitrarily) assume that the window of decline lasts for a week, centered on the known bad transaction. A transaction that receives a negative is defined not to be in its own slipping window, since we will be analyzing whether being in a slipping window has an impact on feedback outcomes for transactions. A transaction is considered to be in a slipping window only if it is in the window defined by some other transaction. In the figure, transactions B, C & D close during a one week period centered on the closing time for transaction C. Transactions B & D are in C's slipping window, while transaction C cannot be in its own slipping window. If B or D received negative feedback, then transaction C would be in the slipping window for one of them, but here we describe only the windows defined by transaction C, which is known to have received a negative feedback.

The slacking/salvaging window is defined for transactions that close during a recovery period after the negative feedback for C is received. Note that this occurs later in time than the slipping window, though it may overlap if the feedback for C occurred less than half a week after transaction C closed. Intuitively, the recovery period lasts until the negative feedback becomes less salient in the user's profile. The basic format of a user profile in 1999 was similar to its current format, although percentages of positive feedback were not calculated and displayed then as they are now. Additional positive feedbacks received after the negative would push the negative comment down on the screen, eventually requiring a user to scroll or even click to another page to see it. We (somewhat arbitrarily) defined the recovery period to last until five positive feedbacks were received. Thus, any transaction that closed in the time between receipt of the first negative and the fifth subsequent positive feedback (including transactions F, G and H in

Figure 17) was classified as being in the slacking/salvaging window for that negative feedback.

The stoning window is also based on the same five feedback recovery period. A transaction is in the stoning window, however, based on the time when feedback for the transaction is received, not the time when the transaction closes. Thus, in the figure, transactions, B, D, E, F, and G, but not H, are in the stoning window defined by feedback for transaction C. Table 13 below summarizes the necessary conditions for a transaction to be a part of a slipping, slacking or stoning window (t_{tneg} is the closing time of a transaction that gets negative feedback and thus defines windows, t_{fbneg} is the time when it gets a feedback and $t_{recover}$ is the time of the fifth feedback after t_{fbneg}).

Effect	Affected Party	Condition	Transactions that satisfy the condition from Fig 17
Slipping	Seller	Transaction time within ½ week of t_{tneg}	B, D
Slacking/ Salvaging	Seller	Transaction time between t_{fbneg} and $t_{recover}$	F, G, H
Stoning	Buyer	Feedback time between t_{fbneg} and $t_{recover}$	B, D, E, F, G

Table 13 Classifying transactions as candidates for various effects

For transactions that did not receive feedback, we can only estimate whether a recovery window was active at the time when the buyer might have contemplated giving feedback. For transactions that did receive feedback, the median time before feedback was about 21 days after transaction close. Thus, if a transaction did not receive any feedback, we classify it as being in a stoning window if a recovery period was active 21 days after the transaction.

The windows are overlapping, but sufficiently distinct to enable analysis. Tables 14-16 show the limited overlap in our dataset. The vast majority of transactions are not in any of the windows defined by negative feedbacks, as we would expect since negatives are rare.

	Slacking Window		
Slipping Window	No	Yes	Total
No	89.50%	2.54%	92.04%
Yes	7.19%	0.78%	7.96%
Total	96.68%	3.32%	100,761

Table 14 Cross Tabulation of slipping and slacking windows

	Stoning Window		
Slacking Window	No	Yes	Total
No	93.48%	3.20%	96.68%
Yes	2.44%	0.87%	3.32%
Total	95.93%	4.07%	100,761

Table 15 Cross tabulation of slacking and stoning window

	Slipping Window		
Stoning Window	No	Yes	Total
No	89.54%	6.39%	95.93%
Yes	2.50%	1.57%	4.07%
Total	92.04%	4.07%	100,761

Table 16 Cross tabulation of slipping and stoning windows

As hypothesized, the three windows are all correlated with transaction outcomes, as shown in Table 17. The probability of negative feedback is higher for transactions in any of the windows than it is for transactions generally. The probability of positive feedback is lower, and the probability of no feedback is higher.

	All Transactions	Slipping Window	Slacking Window	Stoning Window
Negative	0.42%	1.02%	1.14%	1.46%
Positive	42.34%	33.59%	31.91%	28%
None	57.23%	65.39%	66.9%	70.53%
Total	100,761	8,024	3,341	4,103

Table 17 Fraction of feedback received for transactions in the windows

While the three windows identify possible effects that are temporally close to any particular negative, it could be that the observed correlations really reflect a long-term, rather than short-term effect. In particular, receiving a negative feedback may be correlated with other negative feedback because it suggests that the seller's overall quality level is lower than that of other sellers. Under this hypothesis, it is the seller's overall percentage of negative feedback that is correlated with a higher percentage of additional negative feedback.

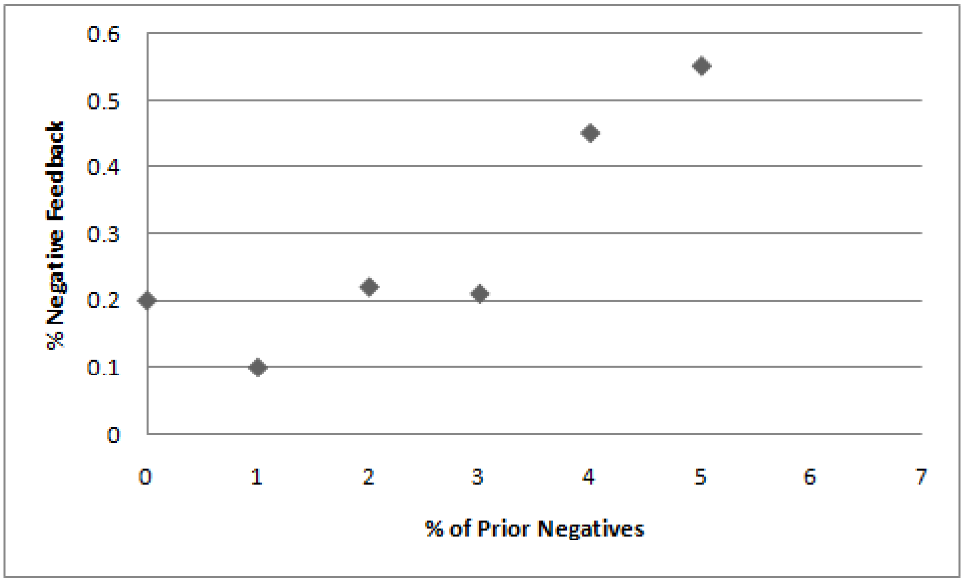

Figure 18 Percentage of prior negatives v/s percentage receiving negative feedback

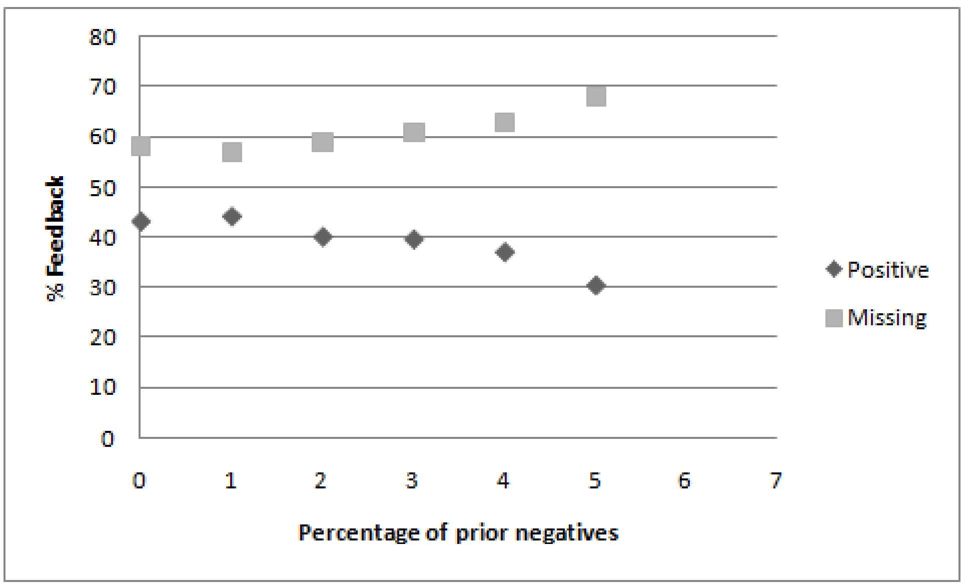

Figure 19 Percentage of prior negatives v/s percentage receiving positive/missing feedback

Figure 18 is consistent with this hypothesis. Transactions are grouped into bins based on the percentage of negative feedback the seller had received for prior transactions. Transactions with 5% or more prior negative feedback are included in the 5% bin. The y-axis indicates the percentage of the transactions in each bin that received negative feedback. Figure 19 shows, using the same bins, the impact of prior negative feedback on the probability of positive and no feedback for a transaction. Thus, we include the percentage of negative feedback as a control variable in our logistic regressions.

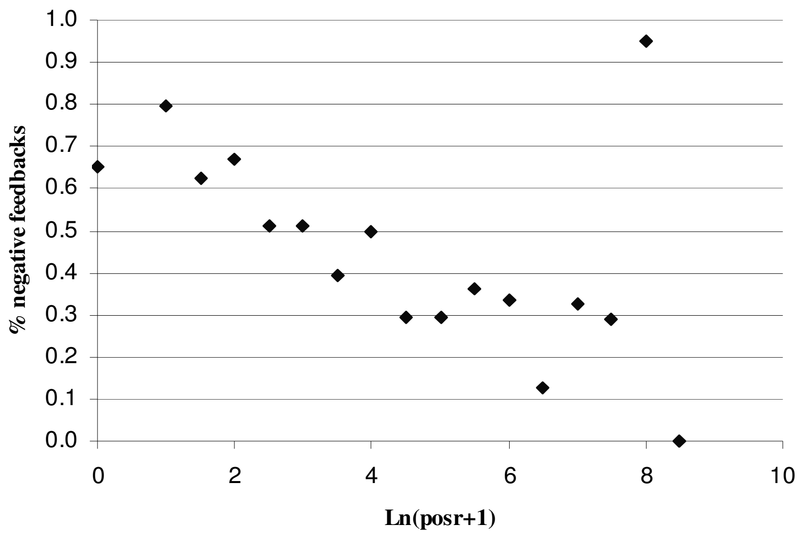

Figure 20 Log of prior positives v/s percentage receiving a negative feedback

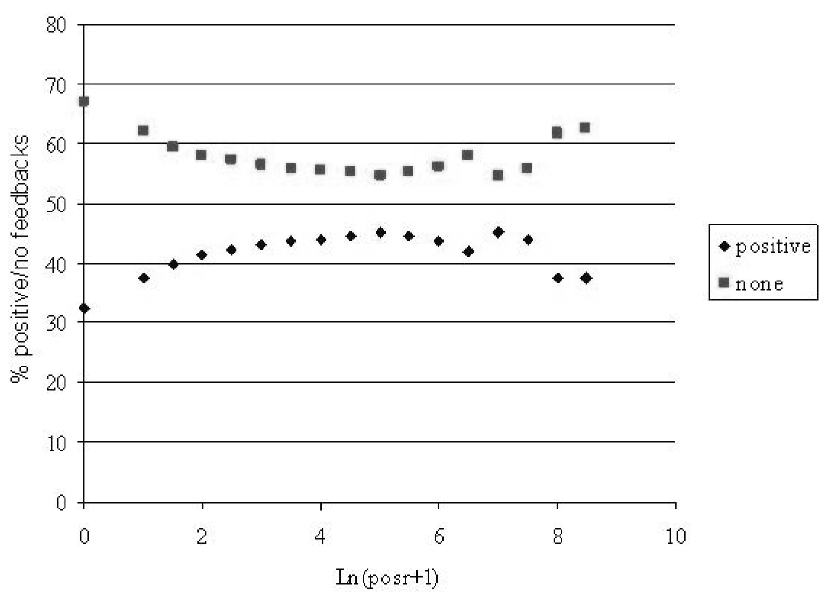

Figure 21 Log of prior positive vs. percentage receiving a positive/missing feedback

We also include a proxy for the seller's experience level as a control variable. Figures 20 and 21 illustrate the correlation between the number of prior positive feedbacks and transaction outcomes. It appears that sellers tend to improve in quality over time, but that once sellers have accumulated a very large number of feedbacks; they are less likely to get any kind of feedback, perhaps because neither they nor their customers think that additional feedback is as important. To account for this curvilinear effect, we include in the regressions a squared version of the experience variable as well as the variable itself.

To summarize, the covariates in the regressions are:

- *Slip*—1 if the transaction is in the slipping window defined by some other transaction's negative feedback, otherwise 0.
- *SlackSalvage*—1 if the transaction is in a slacking/salvaging window, else 0.
- *Stone*—1 if the transaction is in a stoning window defined by some other transaction's negative feedback, else 0.
- *pctneg*—the percentage of the seller's feedback prior the current transaction that was negative or neutral, capped at a maximum of 5%.
- *Logposr*—The log of the seller's number of prior positive feedbacks (1 is added to the count before taking logs, to ensure that the quantity is defined).
- *Logposr squared*—Logposr*logposr.

Two logistic regressions are reported in Table 18. The outcome variables are the probability of receiving a positive feedback and receiving a negative feedback. The regressions include a seller random effect, a per-user effect for the users in our panel data. Each user is assumed to have a particular innate quality level, composed of several immeasurable qualities such as honesty, proficiency in the mechanics of fulfilling the order, communication skills, and diligence. Quality is assumed to be normally distributed across the population of users[14].

OUTCOME	Positive Feedback	Negative Feedback
LOGPOSR Log of seller's number of prior positive feedbacks	0.198	-0.179
	(0.017)***	(0.081)***
LOGPOSR squared LOGPOSR * LOGPOSR	-0.016	0.001
	(0.003)***	(0.012)
PcntNfbr % of seller's prior negative feedback (capped at 5%)	-0.076	0.157
	(0.01)***	(0.04)***
SlipWindow Whether the transaction is in a slip window	-0.122	0.643
	(0.032)***	(0.151)***
SlackWindow Whether the transaction is in a slack window	-0.075	0.134
	(0.047)	(0.224)
StoneWindow Whether the transaction is in a stoning window	-0.314	0.631
	(0.04)***	(0.171)***
CONSTANT	-0.717	-5.316
	(0.025)***	(0.144)***
OBSERVATIONS	100,761	100,761
NUMBER OF UNIQUEID	9665	9665
Standard errors in parentheses *** significant at 10%; ** significant at 5%; *** significant at 1%**		

Table 18 Random effects logistic regression predicting probability of positive and negative feedback

Most of the univariate effects noted in Table 17 hold up in the multivariate, random effects regression. In particular, transactions in slipping or stoning windows were more likely to get negative feedback, and less likely to get positive feedback, even

---

[14] The regressions are conducted using stata's xtlogit command, which allows for the specification of a random effects model. A fixed effects model, while more attractive in some ways, would be inappropriate, especially for predicting negative outcomes, as it would effectively ignore any users who had no variability in their outcomes, including users who did not get any negative feedback

controlling for the other windows and the variables that capture features of the long-term feedback profile. The impact of being in a slacking window, however, is relatively small and not statistically significant; suggesting that much of the apparent effect of being in a slacking window was a spurious attribution that more properly should be attributed to the other variables.

Again to illustrate the model's predictions, consider a hypothetical user who has 100 prior feedbacks: 99 positive and 1 negative. The predicted probability of receiving a positive feedback on the next transaction is 44.6% if the transaction is in none of the windows. This probability drops to 42.8% if the transaction is in a slacking window, 41.6% if in a slacking window, or 37.1% if the transaction is in a stoning window. For this same hypothetical user, the predicted probability of getting a negative feedback on the next transaction is 0.26% if the transaction is in none of the windows. The predicted probability increases slightly to 0.29% if it falls into a slacking window but almost doubles (to 0.48%) if the transaction is in a slipping or stoning window.

3.4 Discussion

The data strongly support the existence of self-selection among sellers. Immediately after receiving a negative feedback, the chance of dropping out increases significantly and fades once the seller has received even one subsequent positive feedback, though there is still a small lingering effect. Sellers with more positive feedback are slightly less likely to drop out, but the effect is quite small.

It is not clear whether sellers who drop out are responding primarily to the psychological impact of receiving a negative or to an expected economic impact on future sales. Given the relatively small impact that a single negative seems to have on profits in the studies that have estimated this, it seems doubtful that the economic impact would be sufficient to drive sellers out of the market.

In addition, if economic impacts were critical to the decision, we would expect some of the sellers who apparently drop out to simply re-register, in order to start fresh with no feedback. This should happen only when the net value of the accumulated profile is worse than the value of a newcomer's profile. Thus, we should see a marked decline in the probability of dropping out as sellers accumulate more positive feedback. The actual effect of positive feedback on the probability of dropping out, however, was very modest, lending additional support to the idea that the psychological impact of a negative is more important than its economic impact.

While seller self-selection may be valuable to preserve the overall trustworthiness of the marketplace, from eBay's perspective it may be that the selection process is convincing too many sellers to refrain from participating. That is, good sellers may be dropping out of the system because negative feedback makes them feel unappreciated. This may explain, in part, why eBay encourages buyers to try to resolve problems with sellers directly before posting negative feedback on the system—fewer negative feedbacks may keep sellers with thin skins continuing to list items.

For sellers who remain in the marketplace, negative feedback leads to more negative feedback. Part of this change seems to be due to slipping, a temporary underlying change in the seller's quality that leads both to the observed negative and a higher probability of negatives on other nearby transactions. Another part of the change seems to be due to stoning, a change in how buyers respond to sellers.

Our initial hypothesis was that stoning would only increase buyers' willingness to provide negative feedback; they would speak up in situations where they otherwise might have submitted no feedback. The analysis, however, shows that transactions in the stoning window were also far less likely to receive positive feedback. The most plausible explanation is that seeing a negative in a seller's profile may color a buyer's interpretation of everything that happens: not only are they likely to interpret ambiguous signals as indicators of bad behavior, but they may be more prone to noticing small problems that do not merit negative feedback but dissuade them from giving positive feedback.

Once slipping and stoning and the seller's prior reputation are accounted for, transactions in a slacking window do not have a significant impact on transaction outcomes. This suggests either that sellers do not change their quality level after receiving negative feedback, or that some users slack while others try to salvage their reputations, leaving no apparent net effect.

One caveat in interpreting our results is that the windows we defined based on the hypotheses of slipping, stoning, and slacking may really be picking up some other effects that differ between the windows. The results are consistent with our hypotheses of slipping and stoning, but do not rule out other possible explanations.

Another caveat is that, even if the effects exist, we may not have correctly classified transactions with respect to the true windows for those effects. For example, the seven day length for slipping windows was chosen arbitrarily. Perhaps seller quality actually declines during a longer window, which would overlap more with the other windows, especially the stoning windows. Thus, we might incorrectly attribute to stoning what is really a slipping effect. To check this, we recalculated longer slipping windows of 14 and 20 days and re-ran the regressions. The coefficients and standard errors barely changed, even with the 20-day slip windows, which included 63% of the transactions in stoning windows.

Another measurement error may result from the stoning window classification of transactions that did not receive any feedback. As noted above, we classified such transactions based on whether a recovery period was in effect 21 days after the transaction's close. But a buyer might have considered and decided against giving feedback somewhat earlier or later than that, and been influenced by the seller's then-current profile. If, for example, too many no-feedback transactions were classified into stoning windows, it would artificially inflate the magnitude of the stoning window's effect on positive feedback. But in that case it would tend to deflate the stoning window's effect on negative feedback. Thus, it is very unlikely that both of the effects of the stoning window are spurious results of our classification of no-feedback transactions.

3.5 Conclusion

The economic theory underlying reputation systems is well-understood, at least if participants are assumed to be rational actors. A seller's feedback profile acts as a signal about the quality of their future transactions, through some combination of indicating their underlying type and the strategic actions they will take. Buyers' responses to those profiles create a sanction for sellers who do get negative feedback. This deters bad behavior and causes those sellers who cannot meet the expectations of buyers to self-select out of the system.

In practice, both buyers and sellers may be acting more on emotion and less on calculation than the economic theories account for. Some motivations not contemplated in conventional utility models, such as a desire to punish wrongdoers even if one does not gain personally from the punishment (Fehr and Gachter 2002), can lead to stoning, which helps to make a reputation system more robust. Our evidence suggests that stoning is a real phenomenon at eBay.

Theoretical models provide upper bounds on the amount of cooperation that is possible among purely self-interested actors. It may be that designs that take into account predictable actions that are not self-interested can lead to even more efficient market outcomes. For example, if buyers will punish sellers who receive a negative feedback even if the expected future behavior of those sellers is no different from that of other sellers, then in equilibrium sellers can continue to provide the same quality rather than slacking off after receiving a negative feedback.

On the other hand, psychological factors can also make a reputation have unintended and undesirable negative consequences. For example, if generally high quality sellers provide worse performance after a negative feedback because they are discouraged, or drop out entirely, efficiency declines. Such consequences may be a necessary evil in order to deter the entry of bad actors, and to encourage effort from good sellers. Our data, however, suggest that any slacking effect at eBay during the period of our study was small relative to the slipping and stoning effects. In order to understand and improve the design of reputation systems in the field, it is important to understand how they are actually functioning. Empirical analysis of the kind provided in this paper is an important first step.

Chapter 4

In the Eye of the Beholder: The Impact of Culture on Trust and Trustworthiness in the Presence of Reputation Systems

4.1 Introduction

A reputation system collects, aggregates, and distributes information about peoples' past behavior. Little is known about cross-cultural differences in how people interpret information from reputation systems and adjust their strategic behavior. This chapter presents the first experimental evidence about such cross-cultural differences. In the process, I also shed light on the question of whether apparent cross-cultural differences in trust are merely a rational response to differences in trustworthiness.

The purpose of a reputation system is to enable trust and encourage trustworthiness in situations of social uncertainty. Rousseau et al. (Rousseau, Sitkin et al. 1998) define trust as "a psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another". Gambetta (Gambetta 1988) defines trust similarly as the confident expectation that the other party will perform the transaction according to the trustor's expectations. According to Coleman (Coleman 1990), situations of trust are a subclass of situations that involve risk, where the risk one takes depends on the performance by another actor. An actor is trustworthy if he would act in a way that benefits the trustor, contingent on the trustor putting herself in a vulnerable position.

Social uncertainty exists for an actor, when her interaction partner has an incentive to act in a way that imposes a cost on her, and she does not have enough information to predict the partner's behavior (Yamagishi, Jin et al. 1998). Clearly, if an interaction partner has an incentive to not be trustworthy, a decision to trust involves greater risk, and hence will be undertaken less frequently. In those cases where the interaction partner is trustworthy, a lack of trust prevents a mutually beneficial transaction, to the detriment of both parties. Thus, trust is sometimes considered a form of social capital, because it enables wealth-generating interactions to take place that might otherwise be missed (Fukuyama 1995; Putnam 2001).

According to Kollock (Kollock 1994), high social uncertainty promotes committed partnerships, which result in repeated interactions between particular trading partners. Repeated interaction between trading partners enables trust in two ways. First, each person can use evidence from past interactions to assess the partner's character. Second, the expectation of future interactions creates a strategic incentive for good behavior in the present, what Axelrod poetically described as the shadow of the future (Axelrod 1984). Anticipating the strategic effect on the trustee of the shadow of the future, it may become rational to trust an interaction partner in the present.

Not all of the potentially beneficial transactions, however, are between partners who will transact repeatedly. For example, one study of the electronic marketplace eBay found that 89% of buyer-seller pairs conducted just one transaction during a five-month period, and just 18% of the transactions were between partners who had previously transacted (Resnick and Zeckhauser 2002). To facilitate transactions among strangers, a different source of trust is needed, where trust is not based on long-term committed partnerships.

One source of trust without committed partnerships is broad cultural norms. Hall & Hall (Hall and Hall 1990) define culture as a system for creating, sending, storing and processing information. Hofstede (Hofstede 2003) considers culture as "the collective programming of the mind that distinguishes the members of one group from another". To study the similarity and differences between national cultures, Hofstede conducted a survey of matched samples of IBM employees from more than 50 countries, along with a series of follow up studies on other samples. Based on this empirical research, Hofstede proposed a framework of five independent dimensions of national culture. He claims that each of these dimensions deals with a basic problem with which all societies have to cope, and to which their approaches vary.

Hofstede's cultural dimensions are: Power Distance; Uncertainty Avoidance; Individualism vs. Collectivism; Masculinity vs. Femininity; Long-term versus short term orientation. A detailed description of these cultural dimensions is beyond the scope of this research. We will focus mainly on the cultural dimension of individualism vs. collectivism. We are interested in this dimension because it is considered to have an impact on economic behavior and trust in particular.

The cultural dimension individualism vs. collectivism is related to the integration of individuals into groups. Individualism stands for a culture in which people are connected by loose ties, and individuals are expected to look after themselves and their immediate families. Collectivism stands for a culture in which individuals are connected by strong ties and integrated into strong, cohesive groups, which pervade most of their activities. In general, countries in Europe and North America are high on individualism, while countries in Asia, South America, and Africa are high on collectivism (Hofstede 2003).

Yamagishi (Yamagishi 1988; Yamagishi, Cook et al. 1998) argues that in a collectivist culture people respond to social uncertainty by forming committed, long standing groups that are mostly closed to outsiders. Trustworthy behavior is expected within groups, but is not extended to out-group individuals. In essence, the shadow of the future created by repeated interactions with a single partner covers all members of the in-group, even those with whom a particular person may interact rarely. In individualist cultures, on the other hand, people exhibit a more generalized trust; trustworthy behavior is expected most of the time, even from strangers. In this sense, individualistic cultures might be better described as universalist (Yamagishi, Cook et al. 1998).

Differences in generalized trust across cultures are well documented. Social surveys like the World Values Survey have reported significant variability in generalized trust across nations based on responses of carefully chosen samples of respondents. The World Values Survey asks the respondents a series of questions concerning trust in various political and social institutions, as well general questions such as 'Generally speaking, would you say that most people can be trusted or that you can't be too careful in dealing with people?'

While the surveys report differences in trust for a large number of cultures using attitudinal measures, Glaeser et al (Glaeser, Laibson et al. 2000) demonstrate that such attitudinal measures can diverge from behavioral measures taken from actual trust decisions. There are several smaller studies that compare trusting and trustworthy behavior. These studies typically compare the behavior in experimental settings using games such as the Prisoner's dilemma game, or the Trust game.

Hayashi et al (Hayashi, Ostrom et al. 1999) report findings from Watabe et al (Watabe, Terai et al. 1996), who studied differences in trust among Japanese and American subjects using a one-shot Prisoner's dilemma game. Both these studies found Japanese subjects to be significantly less trusting compared to the American subjects. Buchan and Croson (Buchan, Croson et al. 2002) compared the behavior of subjects from the United States, China, Japan and Korea in a modified version of the trust game that elicits trust and trustworthiness in a group setting. They reported significant differences in the levels of trust and trustworthiness across cultures. They found that the Americans and Chinese were more trusting, compared to the Japanese and Koreans, while the Chinese and Koreans revealed themselves to be more trustworthy, compared to the other two groups. Wilinger et al (Willinger, Keser et al. 2003) compared the behavior of French and German subjects in a one shot trust game, and found the German subjects to be more trusting but as trustworthy as the French subjects. Barr (Barr 2003) had subjects from Zimbabwean villages participate in a one-shot trust game and found significantly less trust among subjects from villages where long standing ties between families were disturbed due to resettlement.

The convergence of computer and telecommunication technologies resulted in the explosive growth of the Internet during the 1990s and brought about a revolution in the business world. Electronic marketplaces like eBay brought together 'hundreds of thousands of businesses and millions of individual customers by providing a platform for interaction and trading'(Bakos Yannis 1997). Ba and Pavlou (Ba and Pavlou 2002) argue that the revolutionary feature of electronic markets is that they enable the exchange of goods and services without ever meeting the trading partner. Of course, the strangers who act as trading partners may not be trustworthy.

By matching strangers for potentially beneficial trades, electronic markets enhance the need for generalized trust. Many electronic marketplaces use a reputation system to enable trust. A reputation system can gather information about participants' past behavior, both objective data such as the number of transactions and their value, and subjective feedback from interaction partners. For example, following each transaction, eBay allows both the buyer and seller to submit a multiple-choice rating (positive, neutral, negative) and a one-line review of the trading partner. Both the individual reviews and summary counts of the number of ratings of each type are visible to potential future trading partners.

Just as repeated transactions in a committed long-term partnership yield information about who has been trustworthy and has an incentive to be trustworthy, a reputation system provides the information and incentives even for one-time interactions. A buyer can examine the past feedback profile of a seller he has not purchased from before, and decide not to buy, or to offer a lower price to compensate for the possibility that the seller will not be trustworthy. For the seller, the expectation that her feedback profile will be visible to future buyers creates an incentive to be trustworthy. In presence of a reputation system, the future casts a shadow even if the seller does not expect any further interactions with that particular buyer.

Observational and experimental evidence indicates that many buyers use seller reputation profiles as a signal of trustworthiness and adjust their bidding behavior accordingly. Resnick et al (Resnick, Zeckhauser et al. 2006) describe a controlled experiment in which an established seller profile led to 8.1% more revenue for sales of similar items ("matched pairs"), and also summarize results from many cross-sectional observational studies of eBay sales. Since reputations affect buyer behavior, and hence future seller profits, sellers do indeed have some incentive to behave in a trustworthy manner in order to get positive feedback ratings.

In a controlled laboratory setting as well, reputation systems have an effect. Bolton et al (Bolton, Katok et al. 2004) used a multiple-round trust game to study the effects of reputation on trust. They compared three experimental markets: partners, strangers and reputation. In the partners market, the same pair interacted with each other repeatedly, while in the strangers and the reputation markets, the subjects were matched up randomly in each round. In the reputation condition, trustor had access to information about the trustee's past behavior, while no information was available in the strangers' market. The reputation market had significantly higher trust and trustworthiness than the strangers' market, but fared poorly on both counts compared to the partners market. Bolton et al also found that in the reputation market the trustor's own experience affects their trust in addition to information about the trustee's reputation. In particular Bolton et al found that previous experience of being cheated (by another trustee) diminishes trust significantly.

The literature leaves two questions unanswered that we attempt to answer here. First, do reputation systems eliminate the behavioral differences in trust and trustworthiness that are known to exist between individualistic and collectivist cultures? Second, if such differences exist, are the differences in trust merely a response to differences in trustworthiness of interaction partners in a particular setting, or are there additional differences in how people from different cultures assess comparable reputation profiles?

4.2 Experimental Design

We investigated these questions in a laboratory setting. Student subjects from an individualist culture (the United States) and from a more collectivist culture (India) played a sequence of trust games against randomly matched subjects from their own culture. Subsequently, subjects assessed the trustworthiness of a set of eBay seller profiles.

4.2.1 Trust Game

Twelve subjects arrived for each experimental session. Six were randomly assigned to the role of seller, the other six to the role of buyer. Once assigned, they maintained their roles throughout the trust game.

In each of 15 rounds, a seller and a buyer were randomly matched to play a simultaneous move trust game. It is conventional to use a sequential-move trust game, where the buyer first decides whether to buy and the seller only decides whether to ship when the buyer chooses to buy. For reasons explained below, our experiment employed a simultaneous-move game instead. The seller committed to a shipping strategy, a decision about whether she would ship if the buyer chose to buy. In the event that the buyer did not buy, the seller's decision had no impact on the outcome. The strategies and the payoffs (in experimental currency) in the game are depicted in Table 19 below. The number of rounds was announced in advance to observe the endgame effects and to maintain experimental control. (Bolton, Katok et al. 2004)

		Seller	
		Ship	Not Ship
Buyer	Buy	50, 50	0, 70
	Not Buy	35,35	

Table 19 Simultaneous Move Trust Game

Before choosing a buy or ship action, each buyer was asked to predict the action of his or her partner. This provides a subjective assessment of risk, to complement the behavioral measure taken from their actions.

To aid in trust assessments, each buyer saw a reputation profile of the seller she was matched with. The seller's profile included the total number of positive and negative feedbacks and a round-by-round history. For those rounds where the seller's partner chose to buy, the history showed that the seller received positive feedback if she shipped the item, negative feedback if she didn't. For those rounds where the seller's partner chose not to buy, the seller's action was not revealed in his subsequent reputation profile; the entry for those rounds merely stated, "not sold". Note that reputation profiles automatically and accurately recorded a seller's history—there were no considerations of strategic feedback reporting as might occur in a system where buyers choose whether to leave feedback. Sellers received no information about the buyer's history, so they could not predict the buyer's action based on the buyer's history. Sellers were, however, reminded that the buyer was looking at the seller's history.

After each round, buyers and sellers were told how many points they earned. Those points came from two sources. First, each subject received a payoff for the transaction outcome, according to Table 19. Second, each received a payoff based on the accuracy of their prediction about their partner's action. Subjects were rewarded according to a proper scoring rule, detailed in the appendix, so that a buyer maximized her expected payoff by reporting his true beliefs about the probability that the seller would ship. Since sellers chose shipping strategies, contingent on buyer actions, it was possible to evaluate the accuracy of a buyer's prediction of the seller's action even when the buyer did not buy. This was the reason for implementing a simultaneous move game rather than a sequential game.

4.2.2 eBay Profile Assessment

After 15 rounds of the game, the experiment was over. The subjects answered a short demographic questionnaire. They then completed a second task. Each subject was asked to assess twelve eBay feedback profiles. Subjects were reminded that on eBay, not every transaction results in feedback, so the percentage of positive feedback received may be higher or lower than the actual percentage of happy customers. For each profile, the subject was asked to state a probability that the seller with the stated profile would complete a transaction as agreed in a timely manner. The subject was also asked to give a brief textual explanation for why they chose the particular number. The subjects were given 15 minutes to complete the questionnaire.

4.2.3 Experimental Protocol

There were seven sessions in all, with twelve subjects in each session. For logistical reasons, we conducted all our experiments in the United States. All subjects were undergraduate or graduate students at the University of Michigan. In four of the sessions, all the subjects were born in the United States and had lived there for most of their lives. For three of the sessions, subjects were born in India and had lived most of their lives in India. When recruiting Indian subjects, we wanted subjects who had little exposure to cultures outside India. We conducted six sessions of our experiment in first month of the Fall 2006 semester; 63% of our Indian subjects were incoming students who had only been outside India for about a month. To get enough subjects we also included some students who had been in the United States for up to 18 months. U.S. subjects were similarly restricted to those who had not spent more than 18 months outside the U.S.

Although, we tried to make our subjects as similar as possible, there were other demographic differences between the groups[15]. 27 out of the initial 36 US subjects were female, while 27 out of the 36 Indian subjects were male. Majority (25 out of 36) of the Indian subjects were graduate students, while almost all (31 out of 36) of the initial US subjects were undergraduates. A reviewer of an early draft of this work raised a concern that the effects we attribute to culture were in fact due to the differences between graduate and undergraduate students. To check for this possibility, in May 2007, we conducted an additional session with only US graduate students. 3 out of the 12 subjects in the additional session were male and 9 were female. Data from all seven sessions are pooled in the analysis.[16]

---

15 We tested for the effect of academic status and gender as covariates, and found that the effects were not significant.

16 Results using only the first six sessions were similar to those with the full seven sessions; the addition of the seventh session more definitively rules out graduate student status as a confound.

The total time for each session was between 45 and 60 minutes. Subjects, on average, earned $20.51 for their participation, which included a $3.00 show up fee and $10.00 for completing the eBay profile assessments. The experiments were conducted at the Institute for Social Research at the University of Michigan in September 2006 and one session in May 2007. The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). Appendix III provides more details of the experimental protocol, including screenshots from the software used to implement the trust game.

4.3 Results

The first research question was whether the presence of a reputation system is sufficient to erase differences between cultures in the level of trust and trustworthiness exhibited in a trust game. The answer is a clear "no."

In the context of the trust game, trust refers to a buyer choosing to buy from a seller, while trustworthiness means a seller choosing to ship. For this analysis we aggregate the decisions made across all sessions in all the rounds.

As seen in Table 20, the aggregate levels of trust and trustworthiness are higher in the US group compared to the Indian group. There is a high correlation between levels of trust and trustworthiness as would be expected. The aggregate differences between the two groups are statistically significant using a two sample Wilcoxon rank-sum test. (p-value < 0.01).

	USA	India
Percentage of choosing buy (Aggregate Trust)	77%	56%
Percentage of choosing ship (Aggregate trustworthiness)	74%	49%
Total Number of Decisions	360	270

Table 20 Aggregate Trust and Trustworthiness

The round-by-round trading patterns in the two cultures can be seen in the plots of trust in Figure 22, perceived trustworthiness in Figure 23, and actual trustworthiness in Figure 24. As before, trust is defined as the percentage of buyers that chose to buy, while trustworthiness is the percentage of sellers that chose to ship. Perceived trustworthiness is the buyers' assessment of the probability that the seller will ship in that round. In all three figures, data for a trading period have been aggregated, within a culture, across all subjects and hence all experimental sessions.

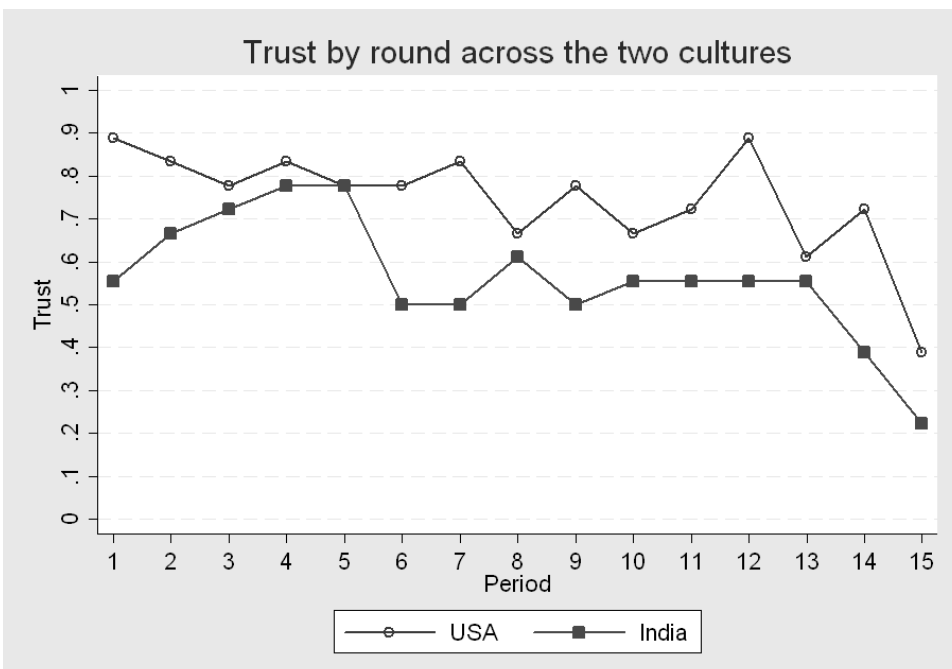

Figure 22 Trust measured as the probability of the buyer choosing to buy

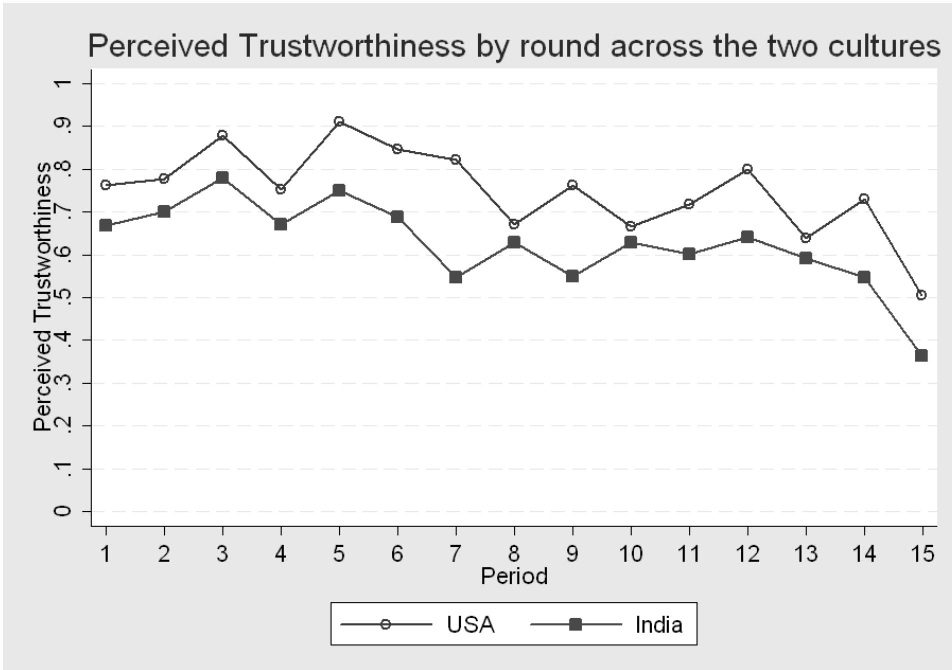

Figure 23 Perceived trustworthiness as reported by the buyers

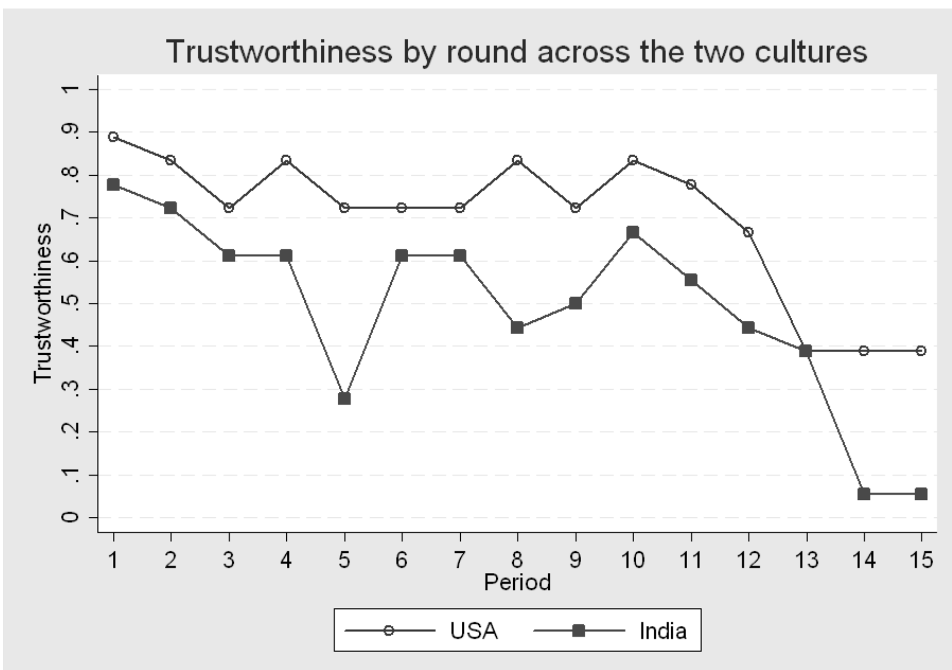

Figure 24 Trustworthiness as the proportion of shipping conditional on buying

There are similarities and differences in the trading patterns of the two groups. The US sessions start with fairly high levels of trust and trustworthiness and remain fairly high for the first 12 rounds. Average trust in the US sessions for the first 12 rounds is

0.80, while average trustworthiness is 0.81. On the other hand, the Indian sessions begin with fairly low levels of trust, which increases till the 5th round, and then dips sharply following a drop in trustworthiness in the fifth period. The average trust in the first 12 rounds in the Indian sessions is 0.6, while the average trustworthiness in the first 12 rounds is 0.56. For both groups, there is a drop-off in trust and trustworthiness in the last few rounds, which can be attributed to the end-game effect—subjects realize that a reputation will have little importance in the future as the game draws to a close.

The logistic regression reported in Model 1 in Table 21 confirms statistically, at an individual level of analysis, the differences in trustworthiness between the two groups of subjects. The outcome variable is the sellers' decision to ship or not ship. A dummy variable codes for the last three rounds, to capture the endgame effect, and an ordinal variable codes the round number for the first twelve rounds. Sellers are significantly less likely to ship in the last three rounds, and trustworthiness declines slightly during the first twelve rounds.

OUTCOME: Seller Ships	Model 1
CULTURE 0 for US, 1 for India	-2.306 (3.08)*
PERIOD Period for the first 12 periods	-0.096 (2.59)*
LAST3 Dummy variable indicating the last 3 rounds	-3.224 (7.93)*
GENDER 0=Male, 1=Female	0.195 0.36
GRAD Dummy variable indicating whether the subject is a graduate student	0.514 1.25
CONSTANT	2.656 (4.15)*
OBSERVATIONS	630
NUMBER OF UNIQUEID	42
Standard errors in parentheses *** significant at 10%; ** significant at 5%; *** significant at 1%**	

Table 21 Logistic Regression for Seller's Ship Decision

Similarly, the logistic regression in Model 2 of Table 21 statistically confirms the difference in trust between the two groups. Indians trust (buy) less compared to their U.S. counterparts, and in the experiments, they are also less trustworthy than the U.S. subjects.

The models include random effects for the individual subjects, to account for the correlated errors introduced by repeated measures from the same subjects. Versions of the models with fixed effects for the sessions showed no significant effects for the sessions within each culture, so we have ignored session effects in the remainder of the analysis.

Buyer Decision: Buy or not buy	Model 2	Model 3	Model 4
CULTURE 0 = USA , 1 = India	-1.162 (3.79) *	-1.111 (2.47) **	-1.186 (2.09) **
PERIOD Period for the first 12 periods	0.03 (1.33)		
GRAD Dummy variable indicating graduate student	0.734 (2.67) *	0.305 (1)	0.058 (0.14)
GENDER 0= Male, 1= Female	0.336 (1.13)	0.327 (0.98)	0.333 (0.75)
PartnerPlus Seller's past positive feedback		0.147 (2.87)*	0.305 (3.20)*
PartnerMinus Seller's past negative feedback		-0.731 (6.47)*	-0.469 (2.76)*
CultureXPartnerPlus Interaction of culture and PartnerPlus		0.079 (1.2)	-0.021 (0.20)
CultureXPartnerMinus Interaction of culture and PartnerMinus		0.215 (1.38)	0.109 (0.46)
LAST3 Dummy variable indicating last 3 periods		-0.960 (2.77)*	-0.653 (1.72)+
ExperiencePlus Prvious good experience			-0.136 (1.2)
ExperienceMinus Previous bad experience			-0.725 (2.75)*
CultureXExperiencePlus Interaction of culture and ExperiencePlus			0.153 (0.97)
CultureXExperienceMinus Interaction of culture and ExperienceMinus			0.255 (0.83)
Constant	0.710 (2.42) **	1.447 (3.74)*	1.877 (3.66)*
Observations	630	630	630
Number of UniqueID	42	42	42

Absolute value of z-statistics in parentheses

+ significant at 10%; ** significant at 5%; * significant at 1%

Table 22 Regression for Buyer's Buy Decision

Our second research question was whether the differences in trust between the two groups reflect a difference in broad outlook, or whether they can be explained by differences in the local environment of the experiment. We answer this question in two stages.

First, we consider whether the trust differences can be explained as a purely rational response to differences in trustworthiness of the interaction partners. If this were the case, then Indian and U.S. subjects, when paired with partners having the same reputation histories, ought to make the same trust decisions. In the experiment, the Indian and U.S. subjects did not have partners with the same reputation histories. The sellers' profiles can be statistically controlled, however.

Model 3 adds two independent variables, one for the number of positive feedbacks in the seller's profile and one for the number of negative feedbacks in the seller's profile. There is a high degree of multi-collinearity between these variables and the period number[17], so we drop the period number from the model. The model also includes interaction terms between culture and the new independent variables, to account for the fact that Indians may be affected differently by the sellers' feedback profiles. As expected, the more rounds in which a seller was trustworthy, the more the buyer trusts; the more rounds in which a seller failed to ship, the less the buyer trusts. Even controlling for the seller's history, however, there is a significant difference between Indian and U.S. subjects in their trust decisions. For example, when faced with a seller who has four positive and two negative feedbacks, prior to the last three rounds, the model predicts that a U.S. subject will buy (trust) 72.25% of the time, while an Indian subject will buy 58.78% of the time. The interaction effects between culture and the partner's past feedback profile are not statistically significant.

---

17 There is not perfect collinearity, because a seller does not get any feedback for a round in which the buyer chose not to buy.

The work of Bolton et al. suggests that a buyer's own history in an experiment may influence her trust decisions. The adage, "once burned, twice shy" may influence a buyer even when she is paired with a different seller than the one who cheated her in the past. Even if both U.S. and Indian subjects respond to getting cheated by sellers in the same way, the fact that more Indian subjects were cheated during the experiment might explain the differences in their trust decisions, without invoking some larger cultural difference in how they make trust decisions.

Model 4 in Table 22 adds two additional covariates to the logistic regression. ExperiencePLUS counts the number of previous rounds where the buyer bought and the seller shipped. ExperienceMINUS counts the number of previous rounds where the buyer bought but the seller did not ship. Finally, to assess whether Indian and U.S. subjects responded differently to being cheated, we included the interaction terms, culture*ExperiencePLUS and culture*ExperienceMINUS.

The results indicate that, buyers are indeed influenced in their trust decisions by their own history of good treatment, even though that treatment came mostly from sellers other than the one they are currently matched with. There was not a significant effect from the number of previous transactions that went well. The interaction terms of culture and own experience did not have a significant effect on the buyer's buy decision. Controlling for all these factors, however, a difference between the two groups remains, and is significant only at the 0.05 level. Thus, there is some difference between U.S. and Indian buyers in making trust decisions that is not explained by their different experiences with sellers in the experiment. Apparently, different levels of trust between the two cultures are not simply a response to the trustworthiness of their immediate environments.

A similar pattern of results holds in regression models predicting the buyers' direct assessments of seller trustworthiness, rather than the buyers' decisions to buy. The outcome variable here is the probability buyers assign to the event that a particular seller will ship. Models 5-7 in Table 23 have the same independent variables as models 2-4 above. Since buyers' assessments of those probabilities are naturally linked to their buying decisions, it is not surprising that the signs on all the coefficients are the same. In these models, however, some of the coefficients are not statistically significant.

Buyer's belief of seller's trustworthiness	Model 5	Model 6	Model 7
CULTURE 0 = USA , 1 = India	-0.151 (3.20) *	-0.135 (2.58) *	-0.136 (2.75) *
PERIOD Period for the first 12 periods	0.005 (1.73)+		
GRAD Dummy variable indicating graduate student	0.1 (2.41)**	0.024 (0.68)	0.004 (0.13)
GENDER 0= Male, 1= Female	0.064 (1.41)	0.05 (1.28)	0.039 (1.17)
PartnerPlus Seller's past positive feedback		0.018 (4.01)*	0.02 (2.36)**
PartnerMinus Seller's past negative feedback		-0.102 (9.28)*	-0.099 (6.14)*
CultureXPartnerPlus Interaction of culture and PartnerPlus		0.01 (1.55)	0.012 (1.08)
CultureXPartnerMinus Interaction of culture and PartnerMinus		0.045 (2.87)*	0.059 (2.51)**
LAST3 Dummy variable indicating last 3 periods		-0.122 (3.51)*	-0.109 (2.99)*
ExperiencePlus Prvious good experience			0 (0.02)
ExperienceMinus Previous bad experience			-0.019 (0.9)
CultureXExperiencePlus Interaction of culture and ExperiencePlus			0.021 (1.36)
CultureXExperienceMinus Interaction of culture and ExperienceMinus			-0.047 (1.64)
Constant	0.667 (14.80)*	0.761 (17.54)*	0.78 (19.88)*
Observations	630	630	630
Number of UniqueID	42	42	42

Table 23 Regression for Buyer's assessment of seller's probability of shipping

Results from the subjects' assessments of eBay profiles also confirm a cultural difference in risk assessments. In these tasks, subjects had no information about whether their trust decisions were correct, and thus there was no impact of personal experience. Table 24 shows the results of a regression model predicting the probability that the subject will assign to a profile, controlling for percentage of positive feedbacks in the profile. To be sure that the experiences in the previous trust experiment were not directly carrying over to their assessments of the eBay profiles, Model 9 includes a control for the number of times the subject was cheated by a seller during the experiment. Consistent with the findings in the experiment, the Indian subjects assigned lower probabilities of good seller behavior than the U.S. subjects did, about 4% lower on average. For a seller with no prior feedback the American subjects assessed the likelihood of trustworthy behavior to be 51.9%, while Indian subjects assessed a new seller to be trustworthy with a probability of 41.8%.

Stated probability of shipping from survey	**Model 8**	**Model 9**
PERC Percentage of positive feedback	0.005 (29.14)*	0.005 (18.80)*
CULTURE 0=USA, 1= India	-0.041 (2.42)**	-0.071 (2.01)**
GRAD Dummy variable indicating graduate student	0.007 (0.44)	0.033 (1.25)
GENDER 0=Male, 1=Female	-0.033 (0.22)	-0.023 (0.94)
ExperiencePlus Prvious good experience		-0.003 (0.59)
ExperienceMinus Previous bad experience		0.007 (0.73)
Constant	0.45 (21.95)*	0.468 (8.27)*
Observations	1008	504
Number of UniqueID	84	42
Absolute value of z-statistics in parentheses		
+ significant at 10%, ** significant at 5%, significant at 1%		

Table 24 Regression for Probability Assessment

4.4 Discussion

The presence of a reputation system appeared to increase trust and trustworthiness for both Indian and U.S. subjects. Buyers assigned higher probabilities of good behavior to sellers with better feedback profiles and were more willing to risk transacting with them. Moreover, in the last few rounds, when the potential future benefits of a good reputation no longer cast a long shadow, trust and trustworthiness declined markedly. This suggests that the higher levels of trust and trustworthiness before the last few rounds were due in part to the presence of the reputation system. For both cultures in our experiment, the reputation system was able to sustain fairly high levels of trust and trustworthiness prior to the last few rounds. It could not, however, eliminate the differences in trust and trustworthiness between the two cultures. Trust and trustworthiness both were consistently lower among Indian subjects than U.S. subjects throughout the experiment.

It is not clear why a difference in trustworthiness between the cultures remained in spite of the incentives created by the reputation system. It could be that, due to cultural norms or less strategic analysis, some or all U.S. subjects were more trustworthy than they would have been if they were merely optimizing their payoffs in response to the difference in expected future revenues from having a good reputation. In the last round, for example, when the future value of a reputation was zero, 40% of U.S. sellers were trustworthy anyway, while less than 10% of Indian sellers were. It is possible, however, that the lower trustworthiness in the Indian sessions was in part a rational response by the sellers to the lower trust. Because of the lower trust, a good reputation is less valuable for an Indian seller than for an American seller. Based on the data available, these effects cannot be studied further; further experimentation is needed.

Buyers assessing the risk of interacting with a particular seller can incorporate three types of information: their personal prior beliefs about the trustworthiness of people in general; information about the characteristics or history of actions of people in the current interaction environment; and information about the characteristics and history of actions of the current interaction partner. Do people from different cultures differ in their personal prior beliefs, reflecting differences among cultures in generalized trust? Do they differ in how they incorporate information about their immediate environments?

In the experiment, U.S. and Indian buyers actually operated in different environments. Sessions were culturally homogeneous, so U.S. buyers were paired with U.S. sellers and Indian buyers with Indian sellers. Subjects could easily assess the cultural makeup of the subject pool because they saw each other across the room. Thus, it is not possible to isolate with certainty the sources of differences in buyer trust between the two subject pools.

The results suggest, however, that there was a difference in generalized trust between the two populations. First, in the first round, before subjects had received any information about the trustworthiness of their partners, U.S. subjects were more trusting than their Indian counterparts. Second, in the assessment of eBay profiles for new sellers with no prior feedback, U.S. subjects were more trusting. Third, regression models provided statistical controls for information about the pool of participants (the subject's own history of treatment by sellers in previous rounds) and about the current interaction partner (the seller's past history). The regression models indicate that buyers with similar information were less trusting if they came from the Indian subject pool.

Future experiments could eliminate differences in the local environment, and thus isolate the effects of differences between subjects in generalized trust. For example, both

U.S. and Indian buyers could be paired with the same pool of sellers (U.S., Indian, or mixed).

Our results were less clear about cultural differences in how subjects interpret information about the pool of interaction partners and about particular interaction partners. The only interaction term that was statistically significant showed that Indian subjects were less affected than U.S. subjects by negative feedback in their partners' profiles. This effect was significant only in the analysis of the buyer beliefs about seller trustworthiness, not in the trust decisions. On the other hand, in assessing eBay feedback profiles, Indian subjects imposed a greater penalty on new sellers who had no prior feedback. Further research is needed to better understand cultural differences in how people incorporate information about a local environment and about individual participants in it.

One caveat in interpreting the results is that the U.S. sessions had a larger percentage of female subjects. We have statistically controlled at the individual level for gender, and found no significant differences in trust or trustworthiness. It is possible, however, that the percentage of female subjects has an effect at the session level, affecting both male and female participants equally, so that it does not show up at the individual level of analysis. We did not conduct enough sessions with variability in gender balance among them to conduct a multilevel analysis. Thus, we cannot rule out the possibility that some of the effects we attribute to culture may be due to differences in gender balance between different experimental sessions.

Another possibility is that the differences in trust and trustworthiness arise out of an income effect. At the time of our experiment, a majority (63%) of our Indian subjects had been in the United States for less than a month. It is plausible that they considered the earnings from the experiment to be higher stakes than the American subjects did. Prior experimental research has found significant reduction in trust as the stakes increase (Johansson-Stenman, Mahmud et al. 2005). With our current data, we cannot rule out the possibility that the income effect was a driving factor in our results. To eliminate this confound, we would recommend that if logistically feasible, future researchers should conduct experiments in the subjects' native countries and design the experiment taking into account the purchasing power of the local currencies (Henrich, Boyd et al. 2001).

One more caveat in interpreting the results is that the trust game in our experiment was a simultaneous move game, instead of the classic sequential game. Prior work by McCabe et al (McCabe, Smith et al. 2000) shows that this distinction affects behavior. They found higher levels of cooperation in a sequential move game than in a simultaneous move game. This raises the possibility that the sequentiality of the game affects the Indian subjects differently than the Americans. We are not aware of prior research that examines such cross-cultural differences. Whether Indians would be more trusting and trustworthy in a sequential environment is a question for future research.

One other potential limitation is that we provided incentives both for correct reports of beliefs about a seller's actions, and for the buyer's choice of action to trust or not. Given the payoff structures, a subject maximized expected payoffs by reporting beliefs honestly and buying whenever he believed the seller would ship with probability greater than 0.7. It is possible, however, that some risk averse subjects chose to hedge their bets. For example, a subject could report low confidence in the seller's trustworthiness, yet still buy, reasoning that either the belief or the action was bound to be correct and be compensated. In fact, when subjects reported beliefs below 0.7, in 30% of the cases they bought anyway. When they reported beliefs above 0.7, in 10% of the cases they did not buy. Overall, reported beliefs were not consistent with actions for Indian subjects in 19% of cases and for U.S. subjects in 15% of cases. Some of the inconsistencies may reflect errors in calculating that the optimal cutoff was 0.7. Others may reflect rational risk spreading. If subjects did hedge their bets, it would result in incorrect measurement of either the attitudinal or behavioral measure of trust, or both. Most likely, the additional noise would just make it harder to detect culture differences in trust, and we found similar patterns for both measures of trust. Still, a replication of this study with an experimental protocol that does not provide rewards for belief reports could rule out this potential confound.

Further research is also needed to replicate this study and test it across other cultures. We have attributed the observed differences between Indian and U.S. subjects to a cultural difference along the individualist/collectivist dimension. Tests with other individualist and collectivist cultures are needed in order to determine whether this attribution is correct, or whether some other dimension of difference between the U.S. and India is the root cause.

4.5 Conclusion

A better understanding of the interaction between culture and the impact of reputation could have important implications for the design of reputation systems. If culture primarily affects the prior beliefs that people have about the trustworthiness of partners without established feedback profiles, then it may be valuable to provide information about the general level of trustworthiness of people with such profiles. For example, eBay could indicate the percentage of all sellers that received a positive feedback on their first transaction. If there are also cultural differences in the interpretation of long term feedback profiles, then a different approach is needed to teach people how to interpret these profiles. Similarly, if sellers in some cultures are not sufficiently accounting for the effect of bad feedback on their future interactions, it may be helpful to make the effect salient in some manner.

Through experimental investigation, we found that reputation systems do not diminish cross-culture differences in trust and trustworthiness in a situation of social uncertainty. Many important questions remain, however. Do cultural factors merely affect the initial interactions in an environment and then those initial interactions affect later interactions? Or do larger cultural norms continue to exert an influence on behavior beyond the effect of the norms and expectations created within an online marketplace or community? Research opportunities are vast, with important implications for the design of systems.

Chapter 5

Conclusion

Many electronic marketplaces employ reputation systems to foster trust and trustworthiness among their participants. Most reputation systems rely on voluntarily provided feedback by a user's transaction partner as the sole means for other traders to learn about the transaction outcome. When feedback is voluntarily provided, it is almost certain to be incomplete, i.e. some transactions will not get a feedback. Further, there are no objective guidelines about what constitutes a satisfactory outcome. As a result, the content of feedback given, and whether a feedback is given at all are both subjective decisions. A consequence of the incomplete and subjective nature of feedback provision is that the interpretation of feedback is also subjective: how one interprets a particular feedback profile will depend on one’s prior beliefs about the distribution of trustworthiness among sellers, and beliefs about feedback giving strategies among buyers. The voluntary and subjective provision of feedback and its subjective interpretation raise a number of questions, of interest to researchers and practitioners of reputation systems. In this dissertation I addressed three such questions.

In Chapter 2, I conducted an experimental investigation of a mechanism designed to give buyers incentives for feedback provision. The mechanism considered here is based on a mechanism proposed by (Gazzale 2004). It starts from the intuition that if buyers' feedback giving histories are made available to sellers, reputation effects will arise on the buyer side and act as incentives for buyers to provide more feedback. In this chapter I examined whether such a mechanism had the desired effects on buyer and seller behavior in practice. I found that the mechanism had mixed success.

Making buyers' feedback giving history available to sellers affected the behavior of both buyers and sellers. Sellers discriminated among buyers based on their feedback giving history. They were found to be more trustworthy towards buyers who had provided feedback in the recent past. However, this discrimination did not yield the sellers the desired benefits. A buyer's feedback giving history was not found to be a good predictor for the probability of getting a negative feedback from her. A majority of the subjects in our experiment gave a negative feedback with a high probability, even when it was costly, and even when they had not recently given a feedback. As a result, sellers who cheated based on a buyer's feedback giving history frequently ended up getting a negative feedback, and lost out on future benefits. At least some of the buyers did not understand that their feedback giving history mattered. As a result they did not give enough positive feedback, and were targeted by sellers for untrustworthy behavior. Even the buyers who understood that the importance of feedback giving history could not predict what metric the sellers would use to discriminate.

As a consequence of all the above, the buyer feedback giving history mechanism did not result in higher levels of trust, trustworthiness or gains from trade. While effects of reputation on sellers' behavior have been long established, we conclude that it is more difficult for the buyers to comprehend the reputation effects of their own feedback giving patterns.

There are two practical implications of this result, one pertaining to this particular type of mechanism and the other more general. For mechanisms that rely on buyer reputations to solicit feedback, the findings from our study sound a warning bell. If some buyers do not understand the importance of feedback giving reputation, making buyer history publicly available could make such buyers more susceptible to being cheated and possibly lead to a drop in overall trustworthiness of the market. The system designers can respond to such a situation either by not making this information available, or by undertaking a campaign to make the buyers aware of the importance of feedback giving histories. Such a campaign could take several forms, it could be a tutorial at the time of registration, or periodic messages reminding them that feedback history matters. Alternatively, as suggested by Dellarocas (Dellarocas, Ming et al. 2003), users who give feedback frequently could be highlighted using distinctions similar to those currently used for sellers with a good reputation. Studying the practical implications of such an education campaign is a fruitful area for future research. Laboratory and field studies could be used to determine the exact form that such an education campaign would need to take in order to be effective.

In our study we examined the implications of one particular implementation of buyer feedback giving reputation. Studying the effects of other implementations (e.g. buyer history where the valence of feedback given is visible) could important implications for designers and also yield insights into the feedback giving process.

Another fruitful direction for future research would be to examine seller and buyer behavior in a setting with three or more outcomes. In a binary outcome trust game, the buyer prefers a good outcome to no transaction, and no transaction to a bad outcome.

In such an environment, we found that almost all buyers reported bad outcomes with a high probability. Consider an alternative environment where a seller can cheat a buyer so that she gets a neutral or mildly dissatisfactory outcome. Empirical estimation of a model of feedback giving by Dellarocas (2007) found that buyers who are neutral or mildly dissatisfied rarely give feedback. Laboratory studies can be conducted to examine the validity of this finding, and to examine the effects of a buyer reputation mechanism on buyers' feedback giving behavior, in particular for a neutral outcome.

The other implication of this result is not restricted to this particular mechanism. Many mechanisms that intend to elicit cooperation, feedback provision etc., achieve desirable outcomes assuming sophisticated strategic reasoning on part of the players. In reality the players may lack the ability to comprehend the equilibrium strategies, leading to inefficient outcomes contrary to the theoretical predictions. The results from this experiment highlight the need to supplement theoretical results with laboratory experiments using human subjects. This naturally suggests certain directions for future research. Other mechanisms that work well in theory could be tested in the laboratory, particularly to examine the validity of assumptions about information structure and strategic sophistication.

In Chapter 3, I conducted an empirical analysis of transaction and feedback data from eBay, to understand the impact a negative feedback in the seller’s profile on his own future behavior, and behavior of the buyers he gets matched with. I found evidence that after sellers got a negative feedback; they were less likely to list another item, at least for a while. For sellers who continue to sell, the presence of a negative feedback in the recent past made them more likely to get another negative feedback. Among the various possible causes that would lead to such an outcome, we found evidence for stoning, i.e. buyers were harsher when dealing with sellers who already had a negative feedback in the recent past. Our results highlight the subjective nature of feedback provision, whereby buyers use a seller's reputation not only to predict the outcomes of transaction in future, but also to interpret the transaction outcomes they experience, and it affects their willingness to be publically critical.

This result is important for researchers estimating the impact of a negative feedback. Stoning makes a negative feedback even more undesirable for sellers. Not only does a negative feedback result in reduced revenues, but it also attracts other negative feedback, thus possibly extending the severity and the length of the punishment period for the seller. Considering the preponderance of positive feedback on eBay, stoning can have a desirable side effect of preventing sellers with a good reputation from engaging in sporadic cheating. But there's an undesirable consequence as well. If a seller gets a negative feedback out of bad luck (e.g. postal damage) or as retaliation, he might have to endure an extended punishment period as a result of stoning.

Stoning can also have other undesirable consequences for the marketplace. It can lead to information herding, whereby sellers who have a recent negative feedback are punished more severely, and sellers who do not have a recent negative feedback are absolved of some transgressions. In other words, the presence of stoning also suggests an unwillingness to cast the first stone at a seller who has a clean reputation. The larger problem here is the subjective nature of feedback provision, of which stoning is just one example. It would be better for the marketplace if the buyers reported their experience without getting affected by what others say.

Reputation system designers could use this insight when designing the interface for feedback provision. If there are separate tabs for positive and negative feedback, a buyer can easily view all negative feedbacks received by a seller, and she may be harsher with the seller than if she had not viewed the negative feedback left by others[18]. If this is the case, the marketplace will be better off if they do not list positive and negative feedback separately. Such a design, however, would make it harder for buyers to be influenced by seller feedback profile during the buying phase as well as, the feedback provision phase. Another alternative would be to not provide an easy way to navigate from the feedback provision form to a user's feedback profile page. Further research needs to be conducted to understand how different designs and presentation schemes affect buyers' willingness to provide negative feedback. One such study could be a laboratory experiment where buyers are given a moderately negative outcome, and asked to provide a feedback when presented with seller reputations that vary in content and presentation format.

In Chapter 4, I addressed the subjective interpretation of feedback. Buyers use the information in a seller's feedback profile to make an assessment of the seller's trustworthiness. Besides the information in the feedback profile, buyers' perceptions are also affected by factors such as their risk aversion, prior experience with other sellers, and their beliefs about the distribution of various seller types in the population. On some

---

[18] It is possible that this is the reason eBay does not have separate tabs for positive and negative feedbacks. When eBay acquired Bazee.com, an online auction market in India, they changed the design of the feedback profile page, to remove separate tabs for positive and negative feedback. Of course, eBay may simply have preferred to make the negative feedback less prominent in order to increase the overall buyer faith in the marketplace.

of these factors, there may be systematic differences based on race, gender, nationality etc. Differences in trust and trustworthiness among people from various cultures are well documented. I conducted a laboratory experiment to examine whether a reputation system can mitigate such differences. I used the individualism vs. collectivism, a dimension of culture that is considered to have an effect on trust. People from individualistic culture are generally more trusting and trustworthy than people from collectivist cultures, when dealing with strangers. The results from our experiment show that a reputation system could not eradicate these differences in trust and trustworthiness. We also found systematic differences between people from the two cultures in their interpretation of comparable reputation profiles.

This result assumes significance considering the continuing spread of electronic markets across the globe. eBay has portals in 29 countries, and there are a number of other country specific electronic markets, such as taobao in China, that rely on a reputation system. A better understanding of the effect of culture on behavior in reputation systems could have important implications for the designers and managers of these reputation systems.

Cultural biases could affect beliefs and behavior in a number of ways. People from different cultures might have different beliefs about the distribution of different types in the marketplace, or about the proportion of good and bad transactions among the missing feedbacks. They might also have different expectations about the norms to be followed, and possibly also about what constitutes a satisfactory transaction. Designers would be well advised to design their systems taking into account culture specific attributes of the target user population. For example, if cultural biases primarily affect the prior beliefs that people have about the trustworthiness of partners without established feedback profiles, a site could indicate the percentage of all sellers that received a positive feedback on their first transaction. Presentation format could also be designed taking into account the cultural differences in the interpretation of long term feedback profiles.

My contribution in this context is to highlight the fact that due to the subjective interpretation of feedback, culture can affect behavior in reputation systems in a number of important ways, and to motivate further research. Further research, using a broader sample of cultures, needs to be conducted to better understand the effects of culture on behavior in reputation systems, and to see what particular presentation formats and education campaigns can enhance the effectiveness of reputation systems in increasing trust and trustworthiness among strangers.

Summarizing, the voluntary and subjective nature of feedback and its subjective interpretation have important implications for the designers, managers and researchers of reputation systems. There are vast opportunities for research, leading to better understanding of practical reputation systems and ultimately to better design.

Appendix I

Experimental Instructions and Screenshots for Chapter 2

A. Experimental Instructions

Introduction

You are about to participate in a decision process in which one of numerous alternatives is selected in each of 45 rounds. This is part of a study intended to provide insight into certain features of decision processes.

During the session you will play a game that gives you an opportunity to earn cash. At the end of the session, you will be paid your earnings plus a $3 show-up fee. If you make good decisions you may earn a considerable amount of money. Decisions and payments are confidential: No one will be told your actions or the amount of money you make.

Procedures and Payoffs

You and the other participants in the room are the players in the game. Half of the players are randomly assigned to play the role of "buyers" whereas the other half play "sellers." The experiment is divided into three sessions, with each session lasting 15 rounds.

You will maintain the same role throughout the experiment.

In each round, each player is randomly matched with another player to trade a (fictional) commodity. First the buyer chooses to either buy or not buy. If the buyer chooses not buy, then the game ends and both players receive 35 points. If the buyer chooses to buy, then the game continues and the seller makes a decision to ship or not ship. Ship pays each player 50 points, while not ship pays the buyer 0 points and the seller 70 points.

After all sellers have made a decision, the buyers who have bought an item in the current round decide to either leave feedback or not leave feedback. Leaving feedback costs the buyer 5 points. If he chooses leave feedback, then all future buyers matched with this seller in this session will know whether or not this seller shipped the item in this round. If the buyer chooses not leave feedback, then future buyers matched with this seller in this session will not know whether or not the seller shipped the item in this round. If the seller did not get to move because the buyer did not buy from her, then the computer reveals the fact automatically.

Note that your earnings in each round depend only on the actions of you and the player with whom you are matched. Also, in each round, you are randomly matched with one of the players of the other role.

Figure A on the last page of the instruction illustrates the game procedure in a game tree.

The exchange rate in the game is 1 dollar per 100 point.

Information

Throughout the game, you will be able to view the information on your transaction history on the left side of the screen. The information about your current match will be displayed on the right side of the screen.

Each player can view his or her total earnings in the current session. The following set of information about a seller will be available to herself and her current match:

• The number of items sold (the number of times the buyer chose to buy from her) in the current session

• The number of positive and negative feedbacks received in the current session

• The number of rounds with no feedback in the current session

• For each round in the current session, a seller can see whether she sold the item, her shipping decision and whether feedback was left. A buyer can see a round-by-round history of the seller with whom he is currently matched: for each round in the current session, whether that seller sold an item, and, if the buyer in that round left feedback, whether the seller shipped the item.

The following set of information about a buyer will be available to himself and his current match:

• The number of times buying in the current session

• The number of times providing feedback in the current session

• For each round in the current session, a buyer can see whether he bought an item, whether the item was shipped to him, and whether he left feedback. A seller can see a round-by-round history of the buyer with whom she is currently matched: whether he bought an item and whether he left feedback for each round in the current session.

The transaction information is reset in each session. That is, after 15 rounds, a new session will start and all of the previous transaction information is no longer available. This means that your match cannot view any information about you from previous sessions.

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

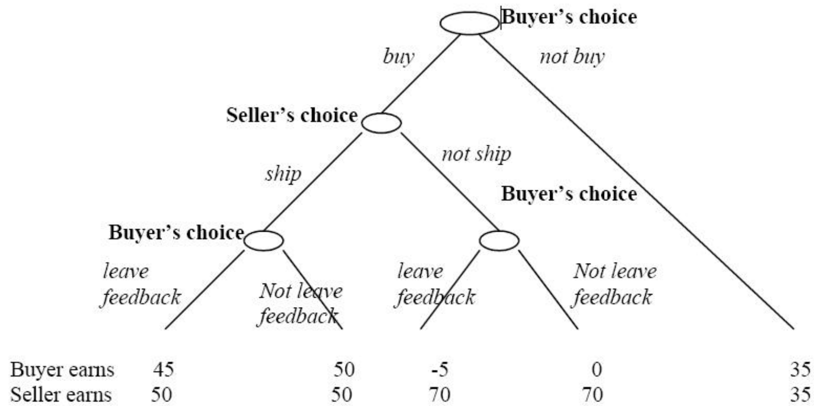
Buyer's choice
buy
not buy
Seller's choice
ship
not ship
Buyer's choice
Buyer's choice
leave feedback
Not leave feedback
leave feedback
Not leave feedback
Buyer earns 45 50 -5 0 35
Seller earns 50 50 70 70 35

B. Typical Screen

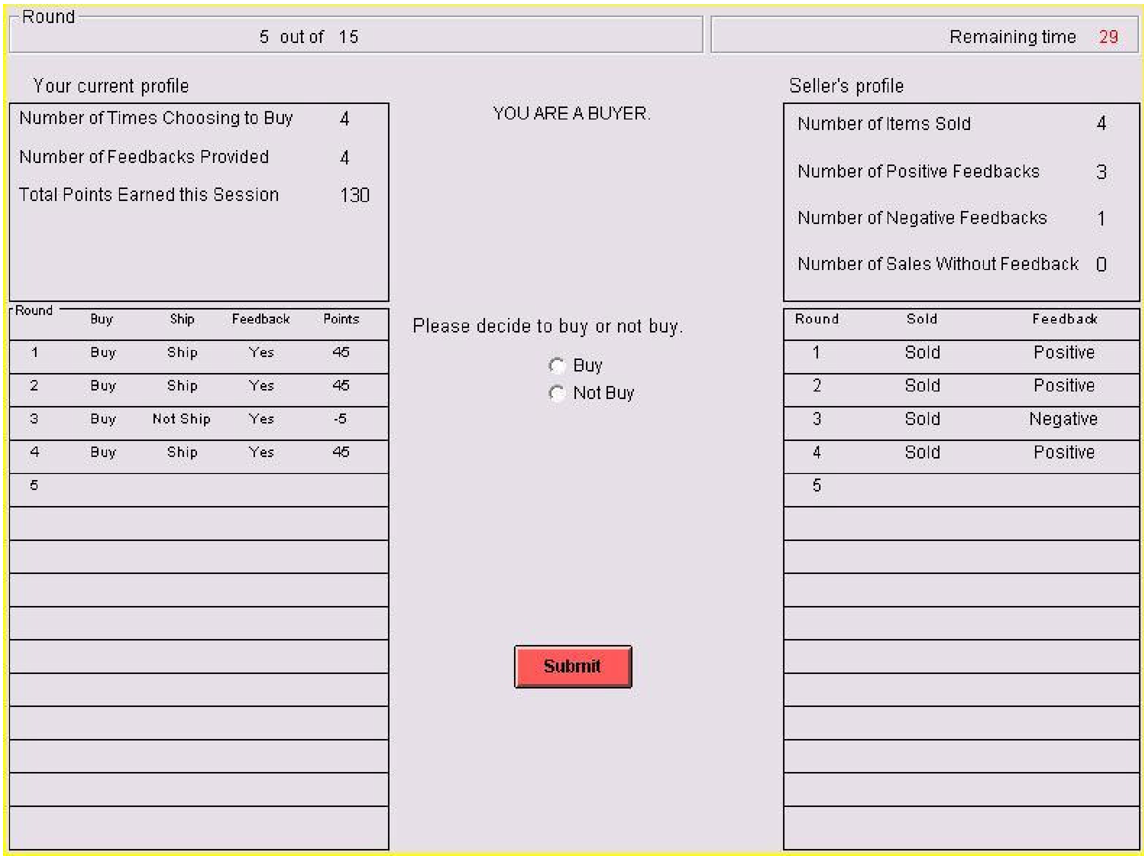

Appendix II

Experimental Instructions and Screenshots for Chapter 4

A. Experimental Instructions

The printed instructions handed to the subjects are reproduced here.

Introduction

This is a study intended to provide insight into certain features of decision processes.

There are two parts of this study. In Part I, you will participate in a game with other participants in the room. Part II of the study is a questionnaire.

In Part I you will play a game that gives you an opportunity to earn cash. If you make good decisions you may earn a considerable amount of money. Decisions and payments are confidential: No one will be told your actions or the amount of money you make. At the end of the session, you will be paid your earnings plus a $3 show-up fee. In Part II you will be given 20 minutes to answer questions from a questionnaire. Completing the questionnaire earns you an additional $10.

Part I

Procedures and Payoffs:

You and the other participants in the room are the players in the game. Half of the players are randomly assigned to play the role of "buyers", whereas the other half play "sellers." You will maintain the same role throughout the experiment.

In each round, each player is randomly matched with another player to trade a (fictional) commodity. The buyer chooses to either buy or not buy. At the same time the seller chooses whether she will ship or not ship, if the buyer decides to buy from her. If the buyer chooses not buy, both players receive 35 points. If the buyer chooses to buy and the seller chooses to ship each player receives 50 points. If the buyer chooses to buy and the seller chooses to not ship the buyer gets 0 points and the seller gets 70 points. The following table illustrates the payoffs for different outcomes.

The exchange rate in the game is 1 dollar per 100 points.

Seller / Buyer	Ship	Not Ship
Buy	50,50	0,70
Not Buy	35,35	

In each cell, the first number is the buyer's payoff, and the second number is the seller's payoff

After both players have made their decision, the computer reveals your match's decision and your payoff for the current round. The computer also records the seller's ship decision automatically, and reveals it in the form of a feedback in the rounds where the buyer chooses to buy.

If the buyer chooses to buy, and the seller chooses to ship, the seller will get a positive feedback for this round. If the buyer chooses to buy and the seller chooses to not ship, the seller will get a negative feedback. If the buyer chooses to not buy in this round, the seller will not get any feedback for this round.

Information

Throughout the game, you will be able to view the information about your transaction history on the left side of the screen. The information about your current match will be displayed on the right side of the screen.

Each player can view his or her total earnings in the session. The following set of information about a seller will be available to herself and her current match:

• The number of items sold (the number of times the buyer chose to buy from her)

• The number of positive and negative feedbacks

• For each round, a seller can see whether she sold the item, her shipping decision and the points earned in the round. A buyer can see a round-by-round history of the seller with whom he is currently matched: for each round, whether that seller sold an item, (that is whether the buyer chose to buy from her), and if so, the feedback for that round.

The following set of information about a buyer will be available only to himself:

• The number of times buying

• For each round in the current session, a buyer can see whether he bought an item, whether the item was shipped to him, and the points earned in the round.

Note that sellers do not have any information about a buyer's previous actions.

Predicting your current match's choice

At the beginning of each round, before you make your choice of buy/not buy or ship/not ship, you will be given an opportunity to earn additional money by predicting the choice of your match.

If you are a seller, you will be asked to predict a probability (from 0 to 1) that the buyer will choose to buy from you. On the other hand if you are a buyer you will be asked to predict the probability that the seller will choose to ship. At the end of each round, we will compare your prediction with the choice actually made by your match. You will earn additional points between 0 to 10 for your prediction. You can maximize your payoff by stating your true belief about what you think your match will do. The payoffs for prediction are calculated as follows:

Suppose you are a buyer, and you predict that your match will choose ship with a probability 0.9

Now, suppose your match actually chooses Ship. In that case your payoff will be:

Prediction Payoff when Seller chooses ship = 10-10*(1- Predicted Probability of Ship)2

= 10 – 10*(1-0.9)2 = 10-10*(0.1)2= 9.9

If your match actually chooses Not Ship. In that case your payoff will be:

Prediction Payoff when Seller chooses not ship= 10-10*(Predicted Probability of Ship)2

= 10- 10*(0.9)2 = 10-8.1 = 1.9

Suppose you are a seller, and you predict that your match will choose buy with a probability 0.75

Now, suppose your match actually chooses Buy. In that case your payoff will be:

Prediction Payoff when Buyer chooses buy = 10-10*(1- Predicted Probability of buy)2

= 10 – 10*(1-0.75)2 = 10-10*(0.25)2= 9.375

If your match actually chooses not buy. In that case your payoff will be:

Prediction Payoff when buyer chooses not buy= 10-10*(Predicted Probability of buy)2 = 10- 10*(0.75)2 = 10-5.625 = 4.375

Note that your prediction payoff does not affect, or depend on your payoff in the actual game. The prediction payoff depends only on how accurately you are able to predict your match’s choice. At the end of each round you will be able to see your prediction payoff for that round. This is true even when the buyer chooses not buy. In that case, the seller’s ship decision is not revealed to the buyer, but the buyer and the sellers still earn points for the accuracy of their predictions.

Since your prediction is made before you know what your match will choose, the best thing you can do to maximize the expected size of your prediction payoff is to simply state your true beliefs about what you think your match will do. Any other prediction will decrease the amount you can expect to earn as a prediction payoff.

The consent form explains your rights as a subject as well as the rules of confidentiality that will be adhered to regarding your participation.

B. Screenshots

Typical Buyer Screen

A. Prediction Task

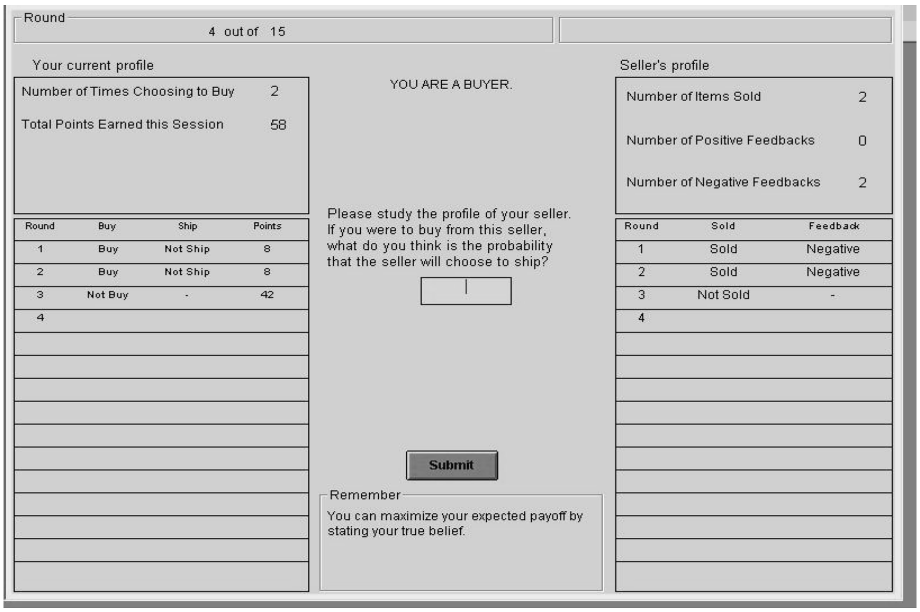

B. Trading Task

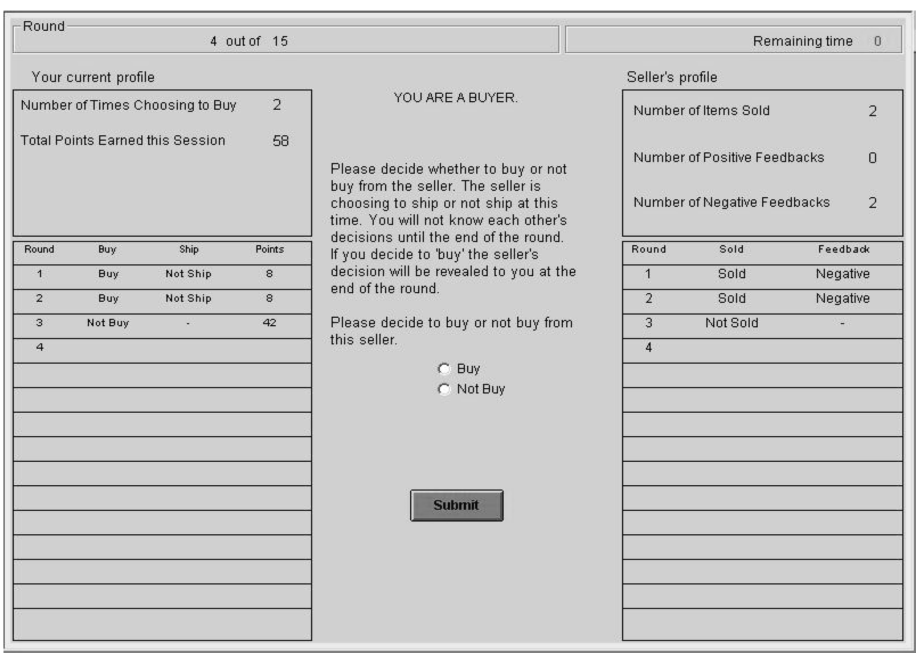

Typical Seller Screen

A. Prediction Task

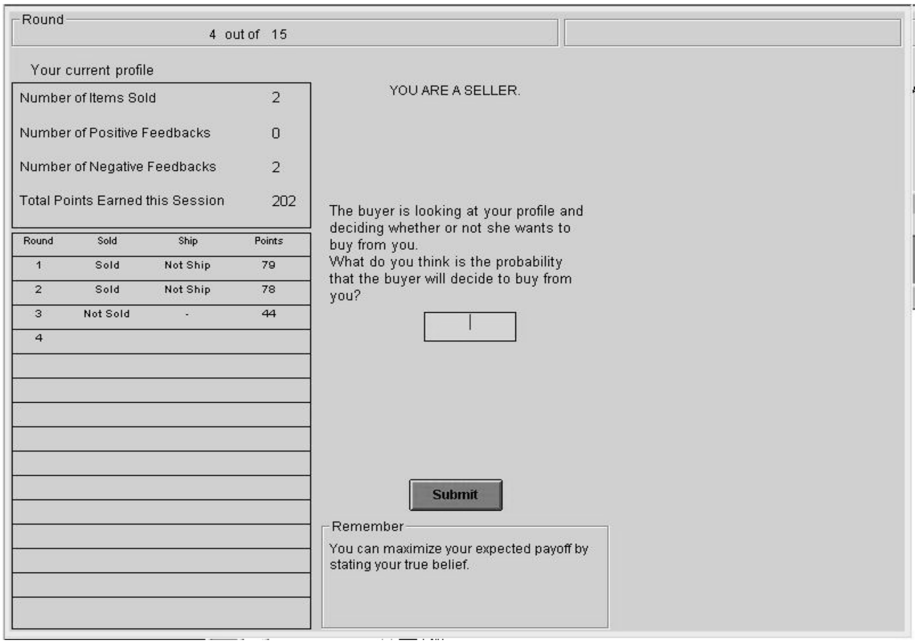

B. Trading Task

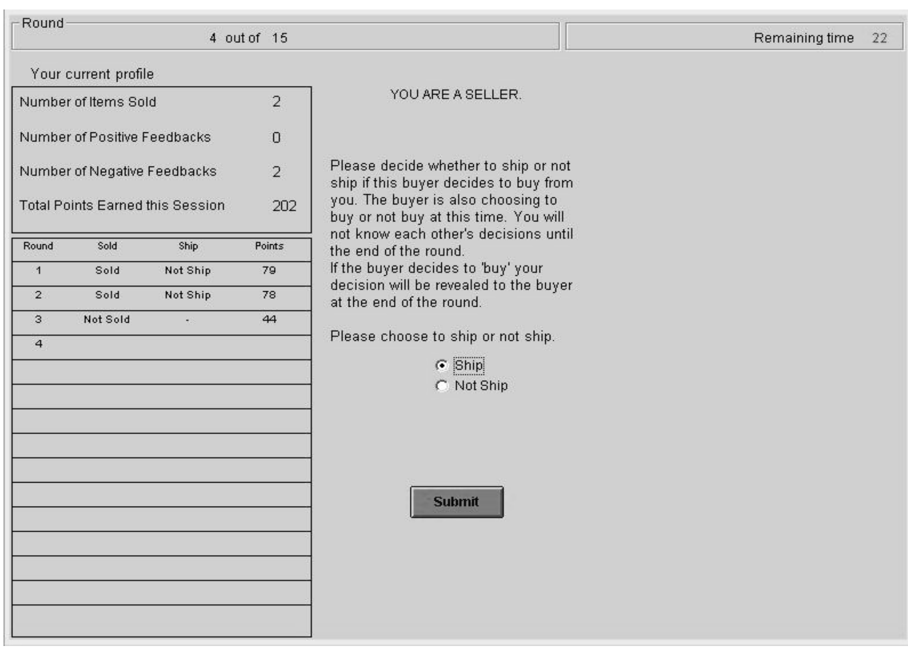

Appendix III

Regression Results with Fixed Effects

OUTCOME: Seller Ships (Using Fixed Effects)	**Low Cost**	**High Cost**
SUB2 Dummy variable for subsession 2	0.204	-0.381
	(0.390)	(0.280)
SUB3 Dummy variable for subsession 3	0.516	0.035
	(0.389)	(0.294)
BHPCNTFBSLEFT Proportion of feedbacks left in the current subsession prior to the current period	-0.467	-0.122
	(0.543)	(0.415)
BUYERSINCELASTFB Number of buy orders since the last time the buyer gave a feedback (before the current period)	-0.238	-0.201
	(0.079)***	(0.066)***
SELLERSINCELASTFB Number of times the seller sold since he last received a feedback	-0.121	-0.268
	(0.103)	(0.077)***
OBSERVATIONS	433	458
NUMBER OF UNIQUEID	1619	1619
Log pseudo-likelihood	-115.63-199.03	-115.63-199.03
Standard errors in parentheses *** significant at 10%; ** significant at 5%; *** significant at 1%**		

Appendix IV

eBay Profile Assessment Questionnaire

Consider yourself to be a buyer on eBay. eBay users can only be identified by their pseudonyms. After each transaction, the buyer and seller are allowed to, but not required to rate each other by leaving a feedback. Each feedback consists of a rating (positive, negative, or neutral), and a short comment. A positive rating carries a feedback score of +1, a neutral rating carries a feedback score of 0, and a negative rating carries a feedback score of -1. These ratings are summarized in individual member profiles. In this study we will consider a short version of eBay member profiles with only positive and negative feedbacks. The member profile for member "user111" is shown below:

ID: user111
Feedback Score: 48 Positive Feedback: 96%
Number of positive feedbacks: 50 Number of negative feedbacks: 2

Feedback Score: This is the sum of all the feedback scores received by the user. In the example shown above, the feedback score is: (50)*(1) + (2)* (-1) + 1(0) = 48

Positive Feedback: This is the fraction of the number of positive ratings to the total number of positive and negative ratings received by the user expressed in a percentage format. In the above example, positive feedback = 48/ (48+2) = 96%

Number of positive feedbacks: Total number of positive ratings received by the user.
Number of negative feedbacks: Total number of negative ratings received by the user.

Now suppose that you are interested in purchasing a digital camera, which is available for $150. You will be presented with a series of profiles (with dummy id user111), of sellers who are selling this item. For each of the profiles, you will be asked to evaluate the probability that the seller will ship the product as described in a timely manner.

We have empirical evidence that about half the transactions on eBay get feedback. What we don't know is whether the buyers who did not leave a feedback were happy with their transactions or not. Thus a seller's feedback profile may not perfectly reflect the proportion of buyers that were satisfied with her performance. We would like to know how you interpret different seller profiles. There are no correct or incorrect answers in this part of the study; we are interested in your interpretation of different feedback profiles. We would also like you to give a brief description of your decision process.

1. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 90 Positive Feedback: 95%
Number of positive feedbacks: 95 Number of negative feedbacks: 5

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

2. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 498 Positive Feedback: 99.2%
Number of positive feedbacks: 499 Number of negative feedbacks: 1

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

3. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 8 Positive Feedback: 90%
Number of positive feedbacks: 9 Number of negative feedbacks: 1

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

4. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 6 Positive Feedback: 80%
Number of positive feedbacks: 8 Number of negative feedbacks: 2

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

5. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 98 Positive Feedback: 99%
Number of positive feedbacks: 99 Number of negative feedbacks: 1

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

6. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 0 Positive Feedback: -
Number of positive feedbacks: 0 Number of negative feedbacks: 0

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

7. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 10 Positive Feedback: 100%
Number of positive feedbacks: 10 Number of negative feedbacks: 0

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

8. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 80 Positive Feedback: 90%
Number of positive feedbacks: 90 Number of negative feedbacks: 10

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

9. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 450 Positive Feedback: 95%
Number of positive feedbacks: 475 Number of negative feedbacks: 25

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

10. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 480 Positive Feedback: 98%
Number of positive feedbacks: 490 Number of negative feedbacks: 10

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

11. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 40 Positive Feedback: 90%
Number of positive feedbacks: 45 Number of negative feedbacks: 5

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

12. A seller is selling this camera; the seller's profile is the following:

ID: user111
Feedback Score: 48 Positive Feedback: 98%
Number of positive feedbacks: 49 Number of negative feedbacks: 1

If you purchased the camera from this seller, what do you think is the probability that the seller will ship the product as described in a timely manner? (On a scale of 0 to 1)

Please describe in 1-2 sentences why you chose the above number

Bibliography

Axelrod, R. M. (1984). The Evolution of Cooperation. New York, Basic Books.

Ba, S. and P. Pavlou (2002). "Evidence OF the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior." MIS Quarterly **26**(3): 243-268.

Bakos Yannis, J. (1997). "Reducing Buyer Search Costs: Implications for Electronic Marketplace." Management Science 43(12).

Bakos, Y. and C. Dellarocas (2003). Cooperation Without Enforcement? A Comparative Analysis of Litigation and Online Reputation as Quality Assurance Mechanisms, SSRN.

Barr, A. (2003). "Trust and expected trustworthiness: experimental evidence from zimbabwean villages." Economic Journal 113(489): 614-630.

Berg, J., J. Dickhaut, et al. (1995). "Trust, Reciprocity, and Social History." Games and Economic Behavior **10**(1): 122-142.

Bolton, G. E., B. Greiner, et al. (2008). Engineering Trust: Stretegic Behavior and the Production of Reputation Information.

Bolton, G. E., E. Katok, et al. (2004). "How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation." Management Science **50**(11): 1587-1602.

Buchan, N., R. Croson, et al. (2002). "Swift neighbors and persistent strangers: A cross-cultural investigation of trust and reciprocity in social exchange." The American Journal of Sociology 108(1): 168.

Coleman, J. (1990). Foundations of Social Theory, Harvard University Press.

Cripps, M., G. J. Mailath, et al. (2004). Disappearing Private Reputations in Long-Run Relationships. PIER Working paper, SSRN.

Dellarocas, C. (2004). Sanctioning Reputation Mechanisms in Online Trading Environments with Moral Hazard, SSRN.

Dellarocas, C. (2005). Reputation Mechanisms. Economics And Information Systems. T. Hendershott, Elsevier **1**.

Dellarocas, C. and C. A. Wood (2007). "The Sound of Silence in Online Feedback: Estimating Trading Risks in the Presence of Reporting Bias." MANAGEMENT SCIENCE: mnsc.1070.0747.

Diamond, D. W. (1989). "Reputation Acquisition in Debt Markets." Journal of Political Economy **97**(4): 828-62.

Ely, J. C. and J. Valimaki (2003). "Bad Reputation." The Quarterly Journal of Economics **118**(3): 785-814.

Falk, A. and U. Fischbacher (2006). "A theory of reciprocity." Games and Economic Behavior **54**(2): 293-315.

Fehr, E. and S. Gachter (2000). "Cooperation and Punishment in Public Goods Experiments." American Economic Review **90**(4): 980-994.

Fischbacher, U. (2007). "z-Tree: Zurich toolbox for ready-made economic experiments." Experimental Economics **10**(2): 171-178.

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

Fudenberg, D. and E. Maskin (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." Econometrica **54**(3): 533-54.

Fukuyama, F. (1995). Trust : social virtues and the creation of: Social virtues and the creation of prosperity. New York :, Free Press.

Gambetta, D. (1988). Trust : Making and Breaking Cooperative Relations, Oxford: Basil Blackwell.

Gazzale, R. (2004). Giving Gossips Their Due: Information Provision in Games with Private Monitoring, Ch. 1, Doctoral Dissertation,University of Michigan.

Glaeser, E. L., D. I. Laibson, et al. (2000). "Measuring Trust." The Quarterly Journal of Economics 115(3): 811-846.

Hall, E. T. and M. R. Hall (1990). Understanding Cultural Differences, Intercultural Press.

Hayashi, N., E. Ostrom, et al. (1999). "Reciprocity, Trust, And The Sense Of Control: A Cross-Societal Study." Rationality and Society 11(1): 27-46.

Henrich, J., R. Boyd, et al. (2001). "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies." The American Economic Review 91(2): 73-78.

Hofstede, G. (2003). Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations, Sage Publications.

Johansson-Stenman, O., M. Mahmud, et al. (2005). "Does stake size matter in trust games?" Economics Letters 88(3): 365-369.

Klein, B. and K. B. Leffler (1981). "The Role of Market Forces in Assuring Contractual Performance." Journal of Political Economy **89**(4): 615-41.

Kollock, P. (1994). "The Emergence of Exchange Structures: An Experimental Study of Uncertainty, Commitment, and Trust." America n Journal of Sociology 100(2): 313-345.

Kuan, K. (2005). Reputation in the Eye of the Beholder: The Role of Information Presentation in Reputation Systems. Stephen M. Ross School of Business, University of Michigan. Ph.D.

Li, L. I. (2006). Reputation, Trust, & Rebates: How Online Markets Can Improve Their Feedback Mechanisms, Institute for Mathematical Behavioral Sciences.

McCabe, K. A., V. L. Smith, et al. (2000). "Intentionality detection and "mindreading": Why does game form matter?" Proceedings of the National Academy of Sciences 97(8): 4404-4409.

Ostrom, E. (2000). "Collective Action and the Evolution of Social Norms." Journal of Economic Perspectives **14**(3): 137-158.

Putnam, R. D. (2001). Bowling alone : the collapse and revival of American community. New York :, Touchstone.

Resnick, P. and R. Zeckhauser (2002). "Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System." The Economics of the Internet and E-Commerce, Advances in Applied Microeconomics **11**.

Resnick, P., R. Zeckhauser, et al. (2000). "Reputation Systems: Facilitating Trust in Internet Interactions." Communications of the ACM 43(12).

Resnick, P., R. Zeckhauser, et al. (2006). "The value of reputation on eBay: A controlled experiment." Experimental Economics 9(2): 79-101.

Rousseau, D., S. Sitkin, et al. (1998). "Not so different after all: A cross-discipline view of trust." Academy of Management. The Academy of Management Review 23(3): 393.

Watabe, M., S. Terai, et al. (1996). "Cooperation in the One-shot Prisoner's Dilemma Based on Expectations of Reciprocity." Japanese Journal of Experimental Social Psychology 36: 183-196.

Willinger, M., C. Keser, et al. (2003). "A comparison of trust and reciprocity between France and Germany: Experimental investigation based on the investment game." Journal of Economic Psychology 24(4): 447-466.

Yamagishi, T., N. Jin, et al. (1998). "In-group Bias and Culture of Collectivism." Asian Journal of Social Psychology 1: 315-328.

Yamagishi, T. (1988). "The Provision of a Sanctioning System in the United States and Japan." Social Psychology Quarterly 51(3): 265-271.

Yamagishi, T., K. S. Cook, et al. (1998). "Uncertainty, Trust, and Commitment Formation in the United States and Japan." The American Journal of Sociology 104(1): 165-194.