Self-Interest, Reciprocity, and Participation in Online Reputation Systems[1]

Chrysanthos Dellarocas · Ming Fan · Charles A. Wood

Sloan School of Management, Massachusetts Institute of Technology,
Cambridge, MA 02139
Department of Management Science, University of Washington Business School,
Seattle, WA, 98195
Department of Management, Mendoza College of Business, University of Notre Dame,
Notre Dame, IN, 46556
dell@mit.edu · mfan@u.washington.edu · cwood1@nd.edu
February 2004

Abstract

Reputation systems are emerging as an increasingly important component of online communities, helping elicit good behavior and cooperation among loosely connected and geographically dispersed economic agents. A deeper understanding of the factors that drive voluntary online feedback contribution is crucial to the long-term viability of such systems and of the online communities that rely on them. This paper contributes in this direction by offering what we believe to be the first in-depth study of the motivations of trader participation in eBay's reputation system. To examine these questions, we analyze data from 51,452 eBay rare coin auctions. We find evidence suggesting that the high levels (50-70%) of voluntary online feedback contribution on eBay are *not* strongly driven by pure altruism. Rather, we analytically and empirically demonstrate that the expectation of reciprocal behavior from partners increases reputation system participation from self-interested eBay buyers and sellers. We develop a random-effects probit model that sheds light on the drivers of feedback submission in individual transactions, and find that participation levels rise, then decline as users accumulate experience within the eBay community.

Keywords: Online Community, Reputation Systems, Altruism, Reciprocity, Self-interest

---

[1] Dellarocas acknowledges the financial support of the National Science Foundation under CAREER Grant No. 9984147. The authors are grateful to Tunay Tunca, Heim Mendelson, Sanjeev Dewan and the participants of the 2003 Workshop on Information Systems and Economics (WISE) for helpful comments.

1. INTRODUCTION

Reputation systems are emerging as an increasingly important component of online communities, helping elicit good behavior and cooperation among loosely connected and geographically dispersed economic agents (Resnick et al. 2000; Dellarocas 2003a). For example, eBay's feedback mechanism is the primary means through which eBay elicits honest behavior and, thus, facilitates transactions between strangers over the Internet (Resnick and Zeckhauser 2002)[2]. Several other online communities also rely on reputation mechanisms to promote trust and cooperation. Examples include eLance (online community of freelance professionals), Slashdot (online discussion forum where reputation scores help prioritize and filter postings), and Epinions (online consumer report forum where user feedback helps evaluate the quality of product reviews).

The success of online reputation systems depends on the sustained voluntary contribution of feedback by community members. Dellarocas (2003b) studies the effects of incomplete feedback submission on eBay-like reputation systems. His theoretical model predicts that, given a seller profit margin, there is a minimum degree of participation (fraction of buyers who submit feedback) that is needed in order for the mechanism to be effective in deterring sellers from cheating. Conversely, for each degree of participation, there is a minimum profit margin that is needed in order for reputation to sustain cooperation. Higher participation, thus, both increases the equilibrium levels of cooperation that are induced by reputation mechanisms, as well as makes such mechanisms effective trust building devices in a wider range of markets.

Existing research on online reputation systems treats user participation as exogenous. Noticeably missing are analyses of the motivation for traders to leave comments. Feedback

---

[2] See Appendix A for a brief overview of eBay's feedback mechanism.

submission is costly to the providers but benefits the whole trading community. Standard economic theory predicts that people are not inclined to contribute voluntarily to the provision of such *public goods* but, rather, they tend to free ride on the contributions of others (Bowles and Gintis 2002). Nevertheless, empirical results from eBay show that buyers submit ratings to more than 50% of transactions (Resnick and Zeckhauser 2002; Wood, Fan and Tan 2003).

A deeper understanding of the factors that drive voluntary participation in online feedback mechanisms is crucial to the long-term viability of such systems and the success of online communities that rely on them. This paper contributes in this direction by offering what we believe to be the first in-depth study of the motivations of trader participation in eBay's reputation system.

To examine these questions, we analyze data from 51,452 eBay rare coin auctions (collected in 2002) and develop theory-driven empirical models that estimate the drivers of trader participation in eBay's reputation mechanism. The patterns of behavior in our data set indicate that self interest is an important motivating force behind the high levels (50-70%) of feedback submission on eBay. eBay encourages both partners of a transaction (buyers and sellers) to rate each other. Our data shows that some eBay users exhibit reciprocity towards partners who have rated them first. But this also creates a selfish motivation to rate one's partner in order to increase the probability of eliciting a reciprocal response. The combined effect strengthens the propensity to participate in eBay's feedback mechanism.

We further develop a discrete choice model that sheds light on the drivers of feedback submission in individual transactions. We find that experienced users tend to rate more frequently. We attribute this to learning effects that lower the cost of rating, as well as to an increased sense of belonging to the eBay community. We *do not* find evidence for crowding-out effects typical of public good experiments. This indicates motivation for leaving comments *is not* strongly motivated by pure altruism targeted towards the specific transaction partner. Rather, a large component for this motivation is a "warm glow" feeling of adhering to the community norms (or perhaps a desire to contribute to the health of the community as a whole).

Overall, we find that the motivation to participate in eBay's review system is multifaceted, ranging from self-interest and reciprocity to "warm glow" feeling of contribution. Therefore, eBay and similar online communities can usefully consider mechanism enhancements that provide higher incentives for participation.

The rest of this paper is organized as follows. Section 2 discusses background work in the theory of public goods. Section 3 examines the drivers of feedback participation at the population level and establishes the presence of a strong component of self-interest. Section 4 introduces an empirical model for predicting feedback participation at the transaction level. Section 5 discusses our results. Finally, Section 6 concludes and discusses directions for future work.

2. BACKGROUND

This study considers online feedback as a public good, in that its submission incurs cost to the provider but benefits the entire trading community. Some of the fundamental questions about the organization of society and markets center on issues raised by the presence of public goods. Economic theory predicts that when many people share the use of public goods, there is an incentive to overuse ("tragedy of the commons"), whereas when people share the obligation to provide them, they tend to undersupply. Interestingly, however, the general consensus of experimental results is that people tend to contribute to public goods at higher levels than theory predicts (Ledyard 1995).

One explanation for this apparent paradox is altruism. Altruistic behavior is explained by the assumption that an agent's utility is positively correlated to the utility of the receiver of the agent's actions. The hypothesis that people are altruistic has a long tradition in economics and has been used to explain charitable donations and the voluntary provision of public goods (Becker 1974). The altruism hypothesis predicts that charitable contributions are subject to the strong neutrality results of public-goods models: government funding of public goods is expected to crowd-out private contributions. Although evidence from sociology, economics, political science, and social psychology shows that altruism is part of human nature, recent research reveals that the pure altruism model lacks predictive power in many situations. Several authors have proposed combining a "joy-of-giving" (sometimes also referred to as "warm glow") motive with altruism to create a model of impure altruism (Cornes and Sandler 1984, 1994; Andreoni 1989, 1990).

Different from altruism, reciprocity represents a pattern of behavior where people respond to friendly or hostile actions with similar actions even if no material gains are expected (Fehr and Schmidt 2000). Rabin (1993) provides a theoretical basis for reciprocal behavior by adopting the concept of "psychological game theory" (Geanakopolos, Pearce and Stacchetti 1989). In psychological game theory utilities depend not only on terminal-node payoffs but also on players' beliefs about other players' intentions. Thus, the payoff of a given terminal node to player A will be higher if player A believes that B's intentions towards him have been kind and lower if he believes that her intentions have been unkind. Levine (1998) offers a similar solution to explain why the same players behave kindly in some games and unkindly in other[3].

Finally, the contribution to public goods can sometimes be explained through purely selfish motives. For example, Glazer and Konrad (1996) propose a signaling theory of charity, where people contribute to charity to signal their social status. On eBay, there are several plausible selfish motivations for feedback submission (elicitation of repeat business, strategic use of praise to elicit like response from partner, etc.)

---

[3] Rabin's theory has been defined only for two-person, normal-form games. Dufwenberg and Kirchsteiger (2003) generalized Rabin's theory to *N*-person extensive form games.

Our objective is to test to what extent each of the above theories can explain feedback submission in online communities. This is accomplished by formulating hypotheses that correspond to the predictions of each theory, and testing these against a rich data set of actual eBay transactions.

3. Theoretical Models of Online Feedback Contribution

Drawing upon the above-mentioned theories of public goods, this section introduces theoretical models to explain voluntary online feedback contribution, both at the population level as well as at the level of individual transactions

3.1 Population-wide Drivers of Participation

To study participation in online reputation systems, we examined 51,452 rare coin auctions that took place from April 24, 2002 to September 11, 2002 on eBay. These auctions include items from 6,242 sellers and 16,405 buyers. We only considered auctions that resulted in a transaction (i.e. auctions that received at least one bid and where the secret reserve price, if it exists, was met). Our dataset includes auction information (e.g., ending time, selling price, comments left for an item, etc.), seller information (e.g., seller rating, seller reputation score, etc.), and bidder information (auctions bid upon, comments left by bidders, bidder reputation score, etc.). Table 1 summarizes some key descriptive statistics of our data set. We observe that participation is substantial: almost 90% of transactions receive comments from at least one trader. Sellers are 10% more likely to leave a comment than buyers. Furthermore, sellers are almost twice more likely than buyers to comment first.

Table 1. Descriptive Statistics of the Dataset

	Number of Auctions	% of Total Auctions
Number of Auctions	51,452	
Auctions where seller left comment	39,942	77.63
Auctions where buyer left comment	34,932	67.89
Auctions where both left comment	29,448	57.23
Auctions where none left comment	6,026	11.71
Auctions where seller commented first	30,524	59.33
Auctions where buyer commented first	14,902	28.96

We argue that the empirically observed pattern of ratings on eBay indicates the presence of a strong component of self-interested behavior rooted on the reciprocal nature of eBay ratings. A lot of eBay users seem to submit feedback primarily motivated by the desire to increase the probability that their partners will reciprocate. We define a *user* as the trader being studied (either the buyer or the seller). The *partner*, then, is the trading partner of the user. (i.e., if the user is the buyer in an auction, then the partner is the seller, and if the user is the seller in an auction, then the partner is the buyer.)

To construct our argument, assume that there are three distinct types of eBay traders:

- *Self-interested traders* rate only if doing so incurs some concrete economic benefits whose expected value exceeds the cost of rating.
- *Altruists* derive satisfaction from rewarding their partners with a rating. We assume that altruists rate with a fixed probability p.
- *Strong reciprocators* never rate first and rate with probability q if they have received a rating from their partner.

In this paper we ignore negative comments and we focus on understanding the motivations behind leaving a positive comment (as opposed to no comment at all)[4]. We also assume that both buyers and sellers receive benefits from a higher reputation score, as has been reported in much online auction research (e.g., Ba and Pavlou 2002; Bajari and Hortacsu 2003; Houser and Wooders 2000; Livingston 2002; Melnik and Alm 2002; McDonald and Slawson 2002). We further assume that the timing of the decision to rate follows a random process. That is, nature sets an alarm clock for each agent that goes off at some random time after the end of a transaction. The agent then wakes up, observes the state of the world (most notably, whether her partner has already submitted a rating) and decides whether to submit one herself.

Let us now examine the qualitative patterns of rating behavior that theory would associate with the presence of each possible combination of the above types in the population of users. There are seven cases:[5]

I. *Self-interested types only*. In this case, no trader would have incentives to rate and thus theory expects that there would be no voluntary participation on eBay's feedback mechanism. This prediction is clearly inconsistent with our data set.

II. *Reciprocators only*. In this case, all users would wait to receive a rating from their partners before submitting a comment. As before, the resulting equilibrium would be one where nobody rates. This prediction is also inconsistent with our data set

---

[4] Neutral and negative comments on eBay are very rare. In the data set of Resnick and Zeckhauser (2002) only 1.2% of buyers and 0.5% of sellers left neutral or negative comments. In our data set the corresponding numbers are similarly low (0.7% of buyers and 0.5% of sellers).

[5] It is possible, though not likely, that an eighth case exists where an entire population is devoid of self-interested types, reciprocators, and altruists. This eighth case would result in no participation in the reputation system, and is not considered important to this research.

III. *Altruists only*. In this case, we would expect a fraction p of users to submit ratings. Furthermore, in the absence of strategic or reciprocity considerations, a user's rating behavior should be independent of the presence or absence of a rating from one's partner. In other words, it should be:

$$\Pr[UserRates] = \Pr[UserRates \mid PartnerRatedFirst]$$

IV. *Altruists and self-interested types only*. Altruists would behave as before, self-interested types would free-ride (i.e. not rate). The resulting behavior would be qualitatively similar to Case III.

V. *Altruists and reciprocators only*. In this case altruists would behave as before and reciprocators would rate with probability q if and only if they receive a rating from their partner. The following proposition then holds[6]:

***Proposition 1*:** If all eBay traders are either altruists or reciprocators and $q > p$ then:

$$\Pr[UserRates] < \Pr[UserRates \mid PartnerRatedFirst]$$

VI. *Self-interested and reciprocators only*. The presence of reciprocators in the population of users provides an incentive to self-interested users to submit feedback, hoping to elicit a like response from their partners and thus increase their own eBay reputation score. This incentive disappears if the partner has already submitted a rating. We therefore expect that only self-interested traders will rate first and only reciprocators will rate second. The following proposition holds true:

---

[6] All proofs are in Appendix B.

Table 2. Relationships between Probability of Rating Conditional on Partner's Behavior and Associated Conclusions

$\Pr[UserRates] = \Pr[UserRates \mid PartnerRatedFirst]$	Provides evidence for no reciprocation or low levels of reciprocation (Cases III or IV)
$\Pr[UserRates] < \Pr[UserRates \mid PartnerRatedFirst]$	Provides evidence for no self-interest or low levels of self-interest. (Cases V or VI)
$\Pr[UserRates] > \Pr[UserRates \mid PartnerRatedFirst]$	Provides evidence for a strong component of self-interest (Cases VI and VII)

Proposition 2: If eBay traders are either self-interested or reciprocators and the fraction of self-interested eBay traders in the population is sufficiently high then:

$$\Pr[UserRates] > \Pr[UserRates \mid PartnerRatedFirst]$$

VII. *All three types*. The addition of altruists in the above mix would increase the fractions of both first and second movers. As before, the relationship between $\Pr[UserRates]$ and $\Pr[UserRates \mid PartnerRatedFirst]$ depends on the proportion of self-interested traders in the population, with $\Pr[UserRates] > \Pr[UserRates \mid PartnerRatedFirst]$ indicating a high proportion of self-interested traders.

The above analysis shows that the relationship between the fraction of users who rate and the fraction of users who rate *conditional on their partner having rated them first* can be used to derive broad conclusions about the mix of motives behind voluntary feedback submission on eBay. The three possible cases are summarized in Table 2. The probabilities revealed from the analysis of our data set are presented in Table 3.

Immediate observation shows that $\Pr[UserRates] > \Pr[UserRates \mid PartnerRatedFirst]$ for both buyers and sellers. This is indicative of high levels of self-interest in the mix of motivations to submit feedback on eBay[7]. Application of a non-parametric t-test confirms that this relationship is statistically significant for both classes of users. The difference is particularly pronounced for sellers, indicating perhaps the presence of a higher fraction of traders that are motivated by self-interest in submitting ratings for buyers.

Table 3. Probabilities of Participation for the Dataset

Pr[*SellerRates*]	0.78
Pr[*BuyerRates*]	0.68
Pr[*SellerRates* | *BuyerRatedFirst*]	0.63
Pr[*BuyerRates* | *SellerRatedFirst*]	0.65

3.2 TRANSACTION-LEVEL DRIVERS OF PARTICIPATION

Having established that the mix of incentives to leave feedback on eBay includes a substantial component of self-interest, in this section we construct a more detailed model for understanding the drivers of voluntary feedback submission at the transaction level. We formulate a user's decision to leave feedback for a partner as an expected utility maximization problem. Specifically, we model the utility of contributing feedback as:

$$U = -c + U_{\textit{self-interest}} + U_{\textit{altruistic}}$$

where:

- c is the cost of contributing feedback, capturing the effort required to log on to the user's account, navigate to the feedback submission screens, and type a comment for the partner. We hypothesize that the cost c declines with experience because of a *learning curve effect.*
- $U_{\textit{self-interest}}$ is the expected utility a user receives from improving her situation, either in the form of expected future reciprocation from the partner or from receiving better service on future transactions.

[7] Note that our test does not rule out the simultaneous presence of altruism in the mix of motives.

- $U_{altruistic}$ is the expected utility a user receives from the "joy-of-giving" and the happiness of the partner.

The utility of not leaving feedback is zero. Therefore, the user will only leave a comment when U is positive, or, equivalently, when:

$$\Delta U_{self\text{-}interest} + \Delta U_{altruistic} - c > 0$$

where $\Delta U_{self\text{-}interest}$, $\Delta U_{altruistic}$ represent the incremental "selfish" and "altruistic" utility a user receives from contributing feedback. In the rest of the section we decompose $\Delta U_{self\text{-}interest}$ and $\Delta U_{altruistic}$ into components that can be related to observable characteristics of users, partners, and transactions[8].

<u>Selfish Motives</u>

The reciprocal nature of eBay ratings provides a selfish motive for submitting feedback. Let x denote a user's current score. If the partner submits a positive comment, her score will increase to x+1[9]. Let p denote the probability that a partner will submit a comment. Let Δp be the incremental probability that he will do so if the user moves first (e.g. because he feels obligated to reciprocate). Finally let δ be the user's discount factor of future earnings, reflecting the frequency of transactions (or, equivalently, the probability that the user will exit eBay following this transaction) and $S(\delta, x)$ the expected present value of the user's future transactions, conditional on her score at the beginning of the next transaction being equal to x.

---

[8] As before, in the following analysis we ignore negative comments and we focus on understanding the motivations behind leaving a positive comment (as opposed to no comment at all).

[9] An important exception is when the two partners have already transacted (and rated one another) in the past. To discourage fraudulent schemes where two colluding partners artificially inflate their scores by repeatedly "buying" from each other in staged auctions for the purpose of (positively) rating one another, eBay only counts one rating per user towards a partner's score. In Section 4 we show that this rule has a negative impact on comment submission frequency between repeat partners.

The expected future surplus if the user does *not* leave feedback can then be written as $S_{no} = pS(\delta, x+1) + (1-p)S(\delta, x)$ whereas the expected surplus if the buyer *does* leave feedback is equal to $S_{yes} = (p + \Delta p)S(\delta, x+1) + (1-p-\Delta p)S(\delta, x)$. The incremental expected surplus due to leaving feedback is thus equal to:

$$\Delta U_{self-interest}(\Delta p, \delta, x) = S_{yes} - S_{no} = \Delta p[S(\delta, x+1) - S(\delta, x)] \qquad (1)$$

Clearly it is $\partial S(\delta, x)/\partial \delta > 0$ (the less one discounts the future, the higher the present value of future gains). Theoretical arguments and previous empirical studies (Livingston 2002; Wood et al. 2003) have shown that the marginal impact of an eBay member's score on revenues is positive but declining (traders whose scores are already high have little to gain from one additional positive rating). We therefore assume that $S(\delta, x+1) - S(\delta, x)$ is positive but declining in *x*.

Altruistic Motives

A complementary motivation for voluntary feedback submission on eBay can be based on arguments of altruism. Previous literature has identified two different "flavors" of altruism (Ribar and Wilhelm 2002): Pure altruism assumes that the user's utility is positively correlated to the increase in the partner's expected utility that results from the buyer's action. Impure altruism assumes that a user simply receives utility ("joy-of-giving", "warm glow") from doing a good deed; the amount of utility is uncorrelated with the effect of the action on the partner's utility.

The analysis of Section 3.1 is inconclusive about the presence and type of altruism that motivates eBay users. Our model helps shed light on this question. Let *y* denote the partner's score. If the user leaves a comment, the partner's score will increase to *y*+1. Let $U(y)$ be the user's assessment of the partner's present value of future transactions, conditional on his score at the beginning of the next transaction being equal to *y*. In a pure altruism model, let κ be the factor by which the partner's expected increase in utility increases the user's utility. Drawing upon the theory of psychological games (Geanakoplos, Pearce and Stachetti 1989) and Levine's (1998) model of altruism and spitefulness (see Section 2) we decompose this factor into two components: $\kappa = \kappa_{pure} + \kappa_{reciprocal}$ where κ_{pure} represents unconditional altruism and $\kappa_{reciprocal}$ represents the additional utility a user gets from helping partners who have already submitted a positive rating, and who have, thus, exhibited kind behavior towards her. The buyer's increase in utility due to leaving feedback is then equal to:

$$\Delta U_{altruistic}(\kappa, y) = \kappa[U(y+1) - U(y)] \qquad (2)$$

As before, we assume that $U(y+1) - U(y)$ is positive but declining in *y*, capturing the fact that sellers whose scores are already high gain little from one additional positive rating. Theory, therefore, predicts that a purely altruistic motivation for submitting feedback will be inversely proportional to the seller's current score[11]. Impure altruism would not exhibit that effect; under the "joy-of-giving" hypothesis the propensity to leave feedback would be independent of the partner's score.

<u>Cumulative Predictions</u>

Table 4 summarizes the predictions of our theory regarding the impact of various attributes of the user, partner and transaction on the probability that the user will submit a rating. The rest of the section discusses these relationships in more detail.

First, a user's *transaction frequency* has a positive correlation to that user's discount factor and, thus, to the present value of whatever gains the user associates with a higher score on future transactions. According to this reasoning, users with high transaction frequencies are expected to

---

[11] This phenomenon is analogous to the "crowding-out" effect that has been associated with charitable contributions.

assign higher "selfish" incremental utility ($\Delta U_{self-interest}$) to rating their partners (expecting that their rating will elicit a reciprocal response that will increase their score) and, thus, to exhibit higher probabilities of feedback submission.

Table 4: Summary of Theoretical Predictions

Attribute	Theory Predictions	Impact on Utility	Expected overall Impact on Pr[*user rates*]
User's transaction frequency	• Increases user's discount factor	• Increases $\Delta U_{self-interest}$	Positive if selfish motives are present.
User experience (User score)	• Decreases user's cost of rating • Decreases user's incremental returns from one additional rating	• Increases U • Decreases $\Delta U_{self-interest}$	Positive if learning effects are stronger than diminishing returns to user score; negative otherwise
Partner score	• Decreases partner's incremental returns from one additional rating	• Decreases $\Delta U_{altruistic}$	Negative if altruism is directed towards partner.
Relative timing of ratings (partner commented first)	• Decreases user's selfish motives for rating • Increases user's psychological utility of reciprocation	• Decreases $\Delta U_{self-interest}$ • Increases $\Delta U_{altruistic}$	Positive if sufficiently strong reciprocation behavior is present; negative otherwise.
Repeat activity (with same partner)	• Decreases user's incremental returns from one additional rating • Decreases partner's incremental returns from one additional rating	• Decreases $\Delta U_{self-interest}$ • Decreases $\Delta U_{altruistic}$	Negative.

A *user's feedback score* (as a proxy for experience) is expected to negatively correlate with the effort required to submit a rating. Experienced users develop techniques (such as cut-and-paste of older comments) that reduce the burden of feedback submission. At the same time, experienced users have lower expected benefits from earning one additional rating. This reduces their selfish utility of rating their partner.

We see, therefore, that a user's score is expected to affect the probability of submitting a rating in two opposite directions. The cumulative effect depends on the relative strength of these two components and cannot be determined a priori from theory. A negative cumulative effect would provide evidence for the presence of selfish motives, whereas a positive cumulative effect would indicate the presence of strong learning effects (but would say nothing definite about selfish motives).

The higher a *partner's score*, the lower the incremental utility that the partner derives from receiving one additional rating from the user and the lower the pure altruistic utility that the user expects to feel from rating the partner. A negative relationship between a partner's score and a user's probability of rating would thus provide support for pure altruism directed towards the partner.

The *relative timing* of ratings affects both selfish and altruistic motives. If a partner has already submitted a rating before the user, this is expected to decrease the selfish motives of the user (because there is no reciprocal rating to be expected as a reward for the user's action). On the other hand, and according to the tenets of psychological game theory, the partner's "act of kindness" of submitting a rating first is expected to increase the user's "reciprocation factor" $\kappa_{reciprocal}$ and thus the user's altruistic utility component. The cumulative effect will be positive if and only if sufficiently strong altruistic motives are present. Therefore, if the probability that a user submits a rating when her partner has already submitted one first is substantially above zero, this provides evidence for the existence of altruistic motives. The opposite relationship provides evidence for the presence of self-interested motives.

A final factor that affects the utility of rating is whether this is a *repeat transaction* between a pair of traders who have previously met (and have rated one another) in the past. eBay's current rules do not count additional ratings exchanged between the same pair of traders towards each trader's score. These rules diminish both the selfish as well as the altruistic incentives to submit a rating. Theory, therefore, predicts a negative relationship between the probability of rating and repeat transactions between the same traders.

Drawing upon the preceding discussion, Table 5 summarizes all possible relationships of five key attributes of users, partners and transaction with participation and lists the corresponding conclusions that one can derive from the observation of each relationship.

Table 5. Summary of Possible Relationships and Associated Conclusions.

Attribute	Observed effect on participation	Conclusion
Activity	Positive	Evidence of self-interest.
	Negative	*No conclusion.*
User Score	Positive	Existence of learning effects.
	Negative	Diminishing returns of eBay score.
Partner's Score	Negative	Evidence of pure altruism.
	Positive	*No conclusion.*
Partner Commented First	Positive	Evidence of reciprocity.
	Negative	Evidence of self-interest.
Repeat Activity	Negative	eBay's "unique rating rules" negatively affect participation.
	Positive	*No conclusion.*

4. EMPIRICAL MODEL AND RESULTS

The preceding discussion suggests the following generalized empirical model:

User Participation = f (PartnerCommentFirst, UserActivity, UserScore, PartnerScore, UserRepeatBusiness)

Many researchers discuss that the effect of certain variables, such as reputation score, diminish as the score increases (e.g., Ba and Pavlou 2002, Kauffman and Wood 2004). As such, transformations are necessary to increase the effect of reputation score variables and activity variables at low levels, and decrease the effects of these variables at high levels. Thus, using the nominal value for these variables could result in a specification error, where the coefficient of the effect decreases as the nominal value of the variable increases. We use natural logs of reputation score and participants' activity in our model.

Control Variables. We also include variables that must be controlled for as we test our hypotheses. The following points summarize our motivations:

- It is reasonable to assume that the final auction price can have an effect on the comments given by both the buyer and the seller. An auction item's picture may influence a buyer's expectations of the product, and it is reasonable to assume that transactions that meet or exceed expectations are more likely to receive positive comments. We include both *auction price* and a dummy variable of picture in our model.
- Many eBay users are "hybrid" traders, that is, they participate in some eBay auctions as buyers and in other as sellers. eBay reputation scores currently do not distinguish between comments received by a user when acting as buyer and comments received when acting as seller. It is plausible that users who engage in selling at a higher rate than buying will economically benefit from a high reputation more than users that engage more in buying and that this might systematically alter the behavior of such users. To control for this, we add a variable, *Selling Activity*, to our empirical model. This variable captures the percentage of a user's selling activity in our data set.
- Past research in marketing indicates that the propensity to engage in word-of-mouth communication depends on the quality of the experience: people are more likely to engage in communication if their experience was extreme (very good or very bad) than if it was average (e.g., Anderson 1998). Several researchers have therefore hypothesized that the decision to leave positive feedback on eBay is often a signal of above-average satisfaction with a transaction (Dellarocas 2001; Resnick and Zeckhauser 2002). Since direct observation of transaction quality is unobservable in this data set, we make an assumption that quality is a somewhat persistent quality of a trader and thus attempt to control for transaction quality using a proxy variable,

Partner Quality, which is measured by the percentage of positive comments received by the partner during his entire history on eBay.

Random-Effects Probit Model. In our proposed model, the dependent variable is binary (i.e., a user leaves a comment or does not). Hausman and McFadden (1984) recommend a Probit model with tests using binary dependent variables because of the few assumptions required of probit and the reliability of the coefficient estimates. Since each market participant could conduct multiple transactions during the study period, there could be several observations for each participant. Therefore, our dataset follows a panel structure. In addition, since the number of transaction a participant conducted could be different, we have an unbalanced panel dataset. An unbalanced panel data set can result in heterogeneity if an individual participants act in a systematically different manner (e.g. a participant's propensity to leave a comment differs from person to person).

Two approaches frequently used to address problems of unobserved heterogeneity are fixed-effects and random-effects models. Fixed-effects models treat the unobserved effects as a constant over time while random-effects models treat the heterogeneity as randomly drawn from some underlying probability distribution. It has been shown that estimates computed using fixed-effects models can be biased for panels over short periods (Heckman 1981, Hsiao 1986, Gulati 1999). This is not a problem with random-effects models. Thus, we use maximum likelihood estimation to estimate a random-effects probit model developed by Butler and Moffitt (1982).

Table 6. Definition for Variables Used in Equations (3) and (4)

Variable	**Definition**
Buyer Participation	A discrete decision variable used in Equation (3) to indicate whether an auction received a comment from the buyer.
Seller Participation	A discrete decision variable used in Equation (4) to indicate whether an auction received a comment from the seller.
Buyer First	A dummy independent variable used in Equation (4) to indicate if a buyer comment for an auction existed at the time when the seller left a comment.
Seller First	A dummy independent variable used in Equation (3) to indicate if a seller comment for an auction existed at the time when the buyer left a comment.
Buyer Activity	Logarithm of the number of auctions that the buyer participated in.
Seller Activity	Logarithm of the number of auctions that the seller participated in.
Buyer Reputation Score	Logarithm of the buyer reputation score reported by eBay at the ending time of the auction.
Seller Reputation Score	Logarithm of the seller reputation score reported by eBay at the ending time of the auction.
Buyer Repeat Business	Average number of auctions per unique seller, measuring a buyer's tendency to have repeat business with a seller.
Seller Repeat Business	Average number of auctions per unique buyer, measuring a seller's tendency to have repeat business with a buyer.
Buyer Quality	The ratio of a buyer's percentage of positive comments and the average percentage of positive comments of all buyers.
Seller Quality	The ratio of a seller's percentage of positive comments and the average percentage of positive comments of all sellers.
Buyer Selling Activity	Percentage of the number of selling auctions of total number of auctions participated.
Seller Selling Activity	Percentage of the number of selling auctions of total number of auctions participated.
Auction Price	Auction final selling price.
Picture	A dummy control variable to indicate whether an auction has a picture of the coin.

In a random-effects Probit model, the residual term can be specified as:

$$\varepsilon_{it} = u_i + v_{it}$$

where u_i is unobserved effect for participant i and v_{it} is the idiosyncratic error, with both components normally distributed with zero means and independently of one another.

Final Models. The random-effects probit model for buyer participation is:

$$\begin{aligned}\Pr(\text{BuyerParticipation}_{it} = 1) = \Phi(\,&\alpha + \beta_1\,\text{SellerFirst}_{it} + \beta_2\,\text{BuyerActivity}_{it} \\ &+ \beta_3\,\text{BuyerReputationScore}_{it} + \beta_4\,\text{SellerReputationScore}_{it} + \beta_5\,\text{BuyerRepeatBusiness}_{it} \\ &+ \beta_6\,\text{SellerQuality}_{it} + \beta_7\,\text{BuyerSellingActivity} + \beta_8\,\text{Price} + \beta_9\,\text{Picture}_{t} + u_i\,),\end{aligned} \tag{3}$$

where Φ is the standard normal cumulative distribution function (CDF) and u_i is unobserved individual participant's error effect discussed earlier. Similarly, the model for seller participation is:

$$\begin{aligned}\Pr(\text{SellerParticipation}_{it} = 1) = \Phi(\,&\alpha + \gamma_1\,\text{BuyerFirst}_{it} + \gamma_2\,\text{SellerActivity}_{it} \\ &+ \gamma_3\,\text{SellerReputationScore}_{it} + \gamma_4\,\text{BuyerReputationScore}_{it} + \gamma_5\,\text{SellerRepeatBusiness}_{it} \\ &+ \gamma_6\,\text{BuyerQuality}_{it} + \gamma_7\,\text{SellerSellingActivity} + \gamma_8\,\text{Price} + \gamma_9\,\text{Picture}_{t} + u_i\,),\end{aligned} \tag{4}$$

The variables used in Equations (3) and (4) are described in Table 6.

Multicollinearity. One of the main concerns with the analysis is the possibility of multicollinearity. Most analyses make the assumption that independent variables are independent of each other, and that one independent variable cannot predict another independent variable. Tables 7 and 8 provide correlation analysis of those variables. The largest pairwise correlation between independent variables for Equation (3) is 25.6% and for Equation (4) is 41.0%. According to Kennedy (1998, p. 187), collinearity should be a concern if the pairwise correlation is above 80%. As such, we are confident that no single variable is related enough to other variables in our model to cause coefficient instability. While we show that pairwise correlation will not be a source of multicollinearity, it is possible that a *single* independent variable can be predicted by a *combination* of other independent variables.

Table 7. Correlation Matrix and Description Statistics for Variables in Equation (3)

Variable	Mean	St. Dev.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Seller First (1)	.593	.491								
Buyer Activity (2)	2.898	1.355	-.031							
Seller Reputation Score (3)	4.538	1.139	.080	.045						
Buyer Reputation Score (4)	3.612	1.124	-.012	.256	.045					
Buyer Repeat Business (5)	1.758	2.288	-.040	.171	.023	-.056				
Seller Quality (6)	.992	0.021	.151	.016	.146	.018	-.028			
Buyer Selling Activity (7)	.099	.190	-.011	.008	.008	.309	-.006	.005		
Price (8)	52.96	217.11	-.011	-.038	.022	.021	-.021	.001	.020	
Picture (9)	.810	.393	.016	.022	.049	.001	.014	.059	-.008	.018

To test for this, we employ two methods: a condition index test and a variance inflation factor test. The condition number test returns a numeric value that is indicative of the level of multi-collinearity. Greene (1999) suggests any condition index greater than 20 may be indicative of multi-collinearity, while Kennedy (1998) suggests a more relaxed criterion of a condition index greater than 30 to be indicative of multi-collinearity. Our tests return a condition index of 15.8 for Equation (3) and 17.5 for Equation (4), thus indicating a lack of multi-collinearity and ensuring that our coefficient estimates are stable.

Table 8. Correlation Matrix and Description Statistics for Variables in Equation (4)

Variable	Mean	St. Dev.	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Buyer First (1)	.290	.454								
Seller Activity (2)	3.535	1.347	-.019							
Buyer Reputation Score (3)	3.612	1.124	.076	.023						
Seller Reputation Score (4)	4.538	1.139	-.045	.378	.045					
Seller Repeat Business (5)	1.391	.491	.020	.409	.013	.011				
Buyer Quality (6)	.992	0.030	.070	.005	.148	.015	-.001			
Seller Selling Activity (7)	.661	.258	.062	.316	-.007	-.080	.132	-.005		
Price (8)	52.96	217.11	-.002	-.053	.021	.022	-.061	.0004	-.024	
Picture (9)	.810	.393	-.007	.055	.001	.049	-.013	-.003	.091	.018

We also examine the severity of multi-collinearity using the variance inflation factor (VIF) suggested by Neter, et al. (1996). Hocking (1996) suggests that a rule of thumb for problem multi-collinearity is that if any VIF is greater than 10, there is evidence of multi-collinearity. In our analysis, all VIFs are less than 1.2 for Equation (3) and less than 1.5 for Equation (4), thus indicating Both the condition number test and the VIF test indicate that there are no undue influences of multi-collinearity on our empirical model.

Results. Tables 9 and 10 show the results of the random-effects probit model that estimate buyer participation and seller participation, respectively. In Tables 9 and 10, ρ indicates the total variance contributed by the participant-specific variance u_i. If ρ is zero, then the participant level variance is not important and we can simply use a pooled probit model. The significance of the ρ coefficient in both models suggests that eBay participants possess different propensities in leaving reviews for others, and that a random-effects model is appropriate. Both probit models are significant as indicated by the Chi-square test using their log-likelihood values.

Table 9. Probit Estimation on the Probability of Buyer Participation

Independent Variable	Coefficient	St. Error	t-stat
Constant	-6.566	0.495	-13.27***
Seller First	-0.174	0.023	-7.56***
Buyer Activity	0.092	0.015	6.23***
Buyer Reputation Score	0.287	0.017	17.10***
Seller Reputation Score	0.040	0.010	3.99***
Buyer Repeat Business	-0.103	0.005	-18.93***
Seller Quality	6.341	0.502	12.63***
Buyer Selling Activity	0.882	0.097	9.09***
Price	0.00001	0.00005	0.15
Picture	0.003	0.029	0.10
ρ	0.808***	0.004	
Log Likelihood	-20941.441		
Chi-square	1237.44***		

Note: *** significant < 0.1%; ** significant < 1.0%; * significant < 5.0%.

Table 10. Probit Estimation on the Probability of Seller Participation

Independent Variable	Coefficient	St. Error	t-stat
Constant	-4.022	0.273	-14.72***
Buyer First	-0.195	0.020	-9.80***
Seller Activity	0.066	0.131	5.03***
Seller Reputation Score	0.180	0.014	13.21***
Buyer Reputation Score	0.085	0.008	10.33***
Seller Repeat Business	-0.543	0.031	-17.68***
Buyer Quality	4.684	0.268	17.50***
Seller Selling Activity	-0.099	0.077	-1.28
Price	0.000002	0.00004	0.04
Picture	0.185	0.033	5.59***
ρ	0.633***	0.006	
Log Likelihood	-17461.593		
Chi-square	1106.20***		

Note: *** significant < 0.1%; ** significant < 1.0%; * significant < 5.0%.

The results indicate that *partner first* has a negative effect on user participation, suggesting that self-interest is more significant than reciprocity on feedback contribution decisions. A participant's *activity level* has a small positive effect on both buyer and seller participation. This is consistent with our theoretical prediction that users with high transaction frequencies derive higher expected value from contributing feedback. In comparison, *reputation score* has a strong positive effect on user participation for both buyers and sellers. Also, *partner score* has a mild positive effect on user participation for both classes of users; this provides negative evidence for the presence of pure altruistic motives. *Repeat business* has a significant negative effect on user participation. It is interesting to note that the negative coefficient of repeat business is much larger for sellers than for buyers, which may suggest that sellers are more rational and self-interested.

Although we make no specific theoretical claims about the control variables, we do notice that the level of *selling activity* (percentage of transactions where the user acted as a seller) is not significant for the seller, but is significant for the buyer. This is consistent with an assertion that buyers who also conduct significant selling activities are motivated to leave comments in hopes of a reciprocal reaction from their partner in order to improve their selling activity at a later date. We also notice that *partner quality*, proxied by the percentage of comments that are positive and indicative of the quality of transactions from this partner, is positively correlated to participation. Since we cannot directly observe the transaction quality information, we encourage future research that delves into the effects of transaction quality on participation in a reputation system.

5. DISCUSSION

Our empirical results confirm many of the theoretical predications and indicate that, together with transaction frequency, eBay reputation score, which, to a large degree, measures user experience, has a positive effect on buyer and seller participation. On the other hand, receipt of a rating from one's partner and repeat business negatively affect participation rate. In this section, we discuss the detailed impacts of the some of the above variables.

<u>Reputation Score</u>

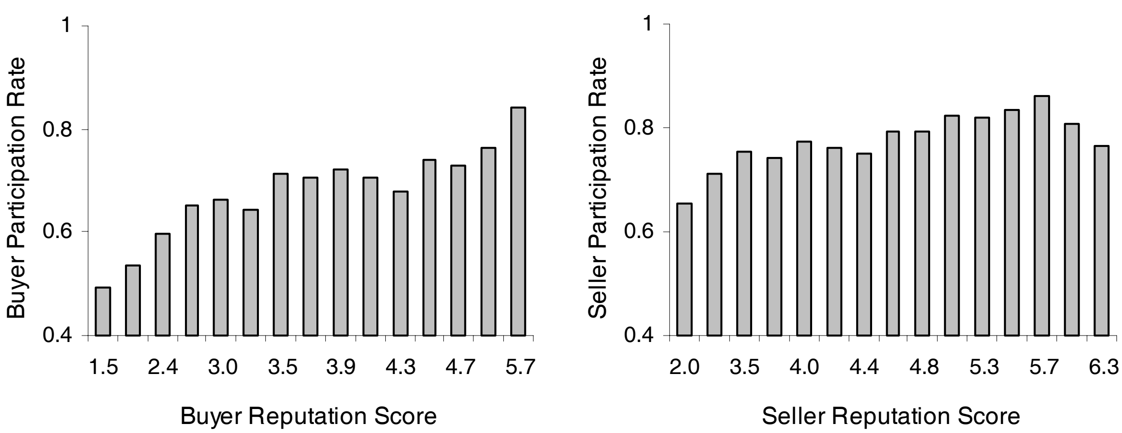

Figure 1. The Effect of Reputation Score on Participation

We calculated the moving average of users' participation rate for different levels of reputation score. The results are depicted in Figure 1. The size for the moving window is constant with each bar representing a sub-sample size of 3,430 auctions.

As reputation scores increase, both buyers and sellers tend to increase their rate of feedback submission. One explanation is that as users accumulate more experience, commenting costs decrease, causing buyers and sellers to be more willing to increase their commenting. However, seller participation decreases at the extreme upper end of the seller reputation score, indicating either a diminishing return of learning effects or simply a decay in participation over time. An alternative explanation for the above phenomena can be based on arguments of loyalty to the eBay community: as users spend more time on eBay, their sense of belonging to the eBay community goes up and, therefore, their participation tends to increase. An inverse-u-shaped relationship between length of membership and community involvement has been observed by several researchers investigating aspects of social capital and civic engagement in voluntary associations (e.g. Putnam 2000). Further research is needed to determine whether the positive correlation between user score and participation is primarily due to learning effects or to an increased sense of loyalty and belonging in the eBay community as well as to resolve the negative impact of reputation score on participation of high-end sellers.

Partner Reputation Score

A pure altruism model would have predicted a negative relationship between partner score and user participation. The absence of such a relationship in our data indicates that, to the extent that altruistic motives exist on eBay, they are of the "impure" nature - users are primarily motivated by the "warm glow" of adhering to the norms of the community (Figure 2). Alternatively, altruism on eBay is not targeted towards the partner but rather towards the community as a whole.

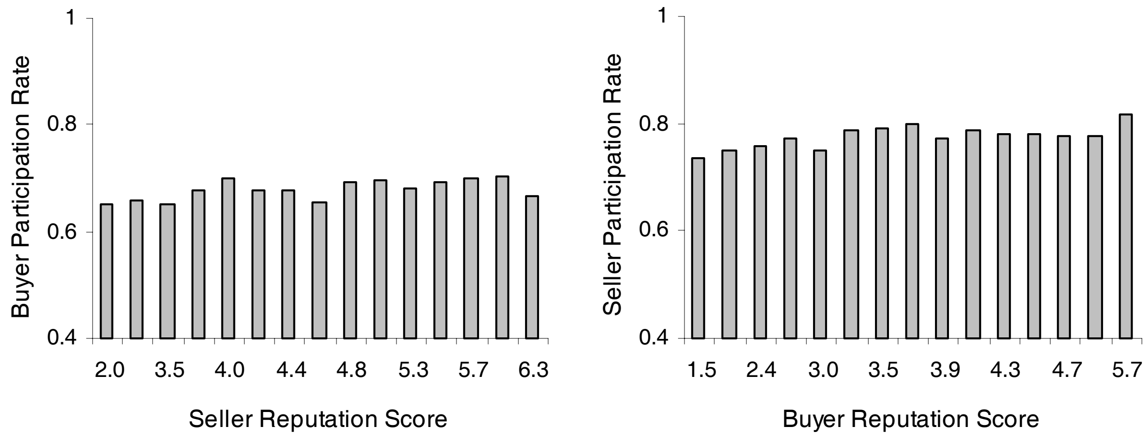

Figure 2. The Effect of Partner Reputation Score on Participation

Activity Level

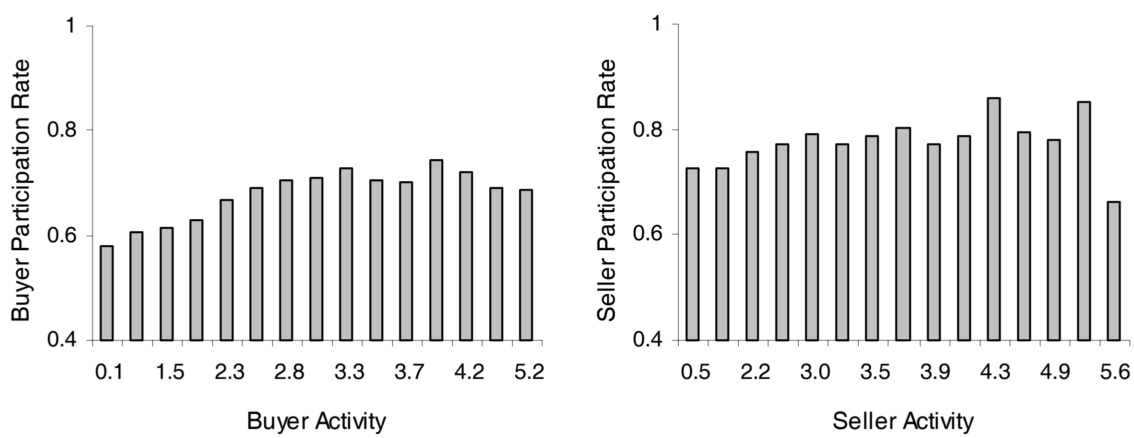

Figure 3. The Effect of the User Activity Level on Participation

Figure 3 shows the effect of user activity on user participation. It indicates a slight increase in participation as user activity increases, as predicted by our model. However, as activity reaches an upper end, we see a decrease in both the buyer and the seller. Based on an optimal investment model, Glaeser et al. (2002) predict that social capital investment would increase with the returns and decline with the opportunity cost of time. Our results could indicate a level of time constraints caused by a large number of activities. In addition, we have reasons to believe that sellers generally incur a higher opportunity cost of time. They might participate in high levels earlier in their tenure.

But as the potential returns decline and cost increases, their participation declines more sharply compared to that of buyers.

Repeat Business

In Figure 4, we show a decrease in participation as repeat business increases. We attribute this to eBay's design, which only counts one comment per partner when calculating a user's score. Because there are immediate diminishing returns to a partner for a positive comment from the same user, we expect a high level of repeat business to result in a low level of participation system participation. Note that this design actually *punishes* sellers who attract the same buyers through good service, since this repeat business causes a reduced participation in the reputation system.

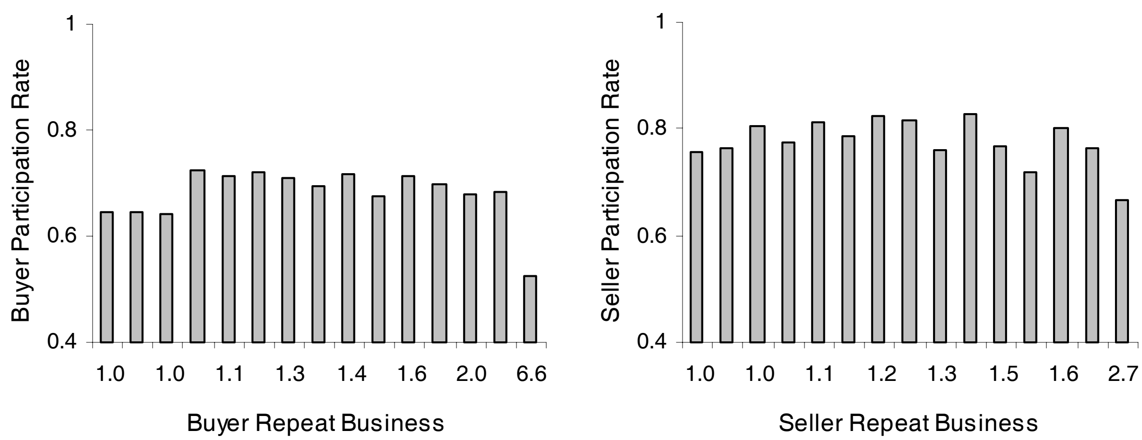

Figure 4. The Effect of Repeat Business on User Participation

6. Concluding Remarks

This paper is the first to study the drivers and dynamics of buyer participation in eBay's feedback system. Through a combination of theory and empirical analysis we demonstrate that the high levels of voluntary participation on eBay's reputation mechanism can be explained through the combined effects of altruism, self-interest, and reciprocation. We further develop a Probit model that sheds light on the drivers of feedback submission in individual transactions. We find that both a user's and a partner's reputation scores positively affect a user's propensity to participate in the reputation system. We also find that activity levels are positively correlated to the probability of participation for both buyers and sellers. Finally, we show that repeat business between the same pairs of buyers and sellers has a negative effect on the probability of participation, probably due to eBay's auction design, which does not count multiple comments from the same user in their reputation score.

An important insight of this study is that voluntary participation in online feedback mechanisms seems to be largely motivated by self-interest. Furthermore, although the participation rate is quite high, it is not robust. Sellers' participation declines when sellers' reputation scores and activity levels get very high. Our findings, thus, suggest that eBay might need to look more closely into the issue of participation and perhaps introduce explicit or implicit incentives that make participation levels more robust. For example, one plausible suggestion would be to discount older comments in the calculation of a user's eBay "reputation score", making the reputation mechanism reflect more recent activity and thus sustaining the motivation to elicit comments from trading partners even for the most active users. Another suggestion would be to make everybody's feedback contribution rate public knowledge and to "reward" frequent contributors with distinctions similar to those that eBay currently reserves for users with high reputation scores.

High levels of participation are crucial to the success of any online community. In future work, we will study the extent to which the behavioral patterns we discovered on eBay are present in other online trading communities that base the elicitation of good behavior on online feedback. We will also look more closely on mechanism enhancements that improve community participation.

REFERENCES

Anderson, E.W. 1998. Customer satisfaction and word of mouth. *Journal of Service* 1 (1), 5-17.

Andreoni, J. 1989. Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of Political Economy* 97, 1447-1458.

Andreoni, J. 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. *The Economic Journal* 100, 464-477.

Ba, S., 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.

Bajari, P. and A. Hortacsu. 2003. Winner's curse, reserve prices and endogenous entry: Empirical insights from eBay Auctions. *Rand Journal of Economics* 34 (2).

Bowles, S. and H. Gintis. 2002. Social capital and community governance. *The Economic Journal* 112, F419-F436.

Becker, G. S. 1974. A Theory of social interactions. *Journal of Political Economy* 82, 1063-1093.

Butler, J.S. and R. Moffitt. 1982. A computationally efficient quadrature procedure for the one-factor multinomial probit model. *Econometrica* 50 (3), 761-764.

Cornes, R. and Sandler, T. 1984. Easy Riders, Joint Production, and Public Goods. The Economic Journal 94, 580-598.

Cornes, R. and Sandler, T. 1994. The Comparative Static Properties of the Impure Public Good Model. Journal of Public Economics 54, 403-421.

Dellarocas, C. 2001. Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms. Proceedings of the 3rd ACM Conference on Electronic Commerce. Association for Computing Machinery, Tampa, FL, 171-179.

Dellarocas, C. 2003a. The Digitization of Word-of-Mouth: Promise and Challenges of Online Feedback Mechanisms. Management Science 49 (10), 1407-1424.

Dellarocas, C. 2003b. Efficiency and Robustness of Binary Feedback Mechanisms in Trading Environments with Moral Hazard. MIT Sloan Working Paper No. 4297-03.

Dufwenberg, M. and Kirchsteiger, G. 2003. A Theory of Sequential Reciprocity. Games and Economic Behavior (forthcoming).

Fehr, E. and S. Gachter. 2000. Fairness and retaliation: the economics of reciprocity. Journal of Economic Perspectives 14, 159-181.

Geanakopolos, J., Pearce, D. and Stacchetti, E. 1989. Psychological Games and Sequential Rationality. Games and Economic Behavior 1 (1), 60–80.

Glaeser, E., D. Laibson, and B. Sacerdote. 2002. An economic approach to social capital. *The Economic Journal* 112, F437-F458.

Glazer, A., Konrad A. K. 1996. A Signaling Explanation for Charity, The American Economic Review, 86 (4) 1019-1028

Greene, W. H. 1999. *Econometric Analysis*, 4th Edition, Englewood Cliffs, NJ: Prentice Hall.

Gulati, R. 1999. Network location and learning: the influence of network resources and firm capabilities on alliance formation. *Strategic Management Journal* 20, 397-420.

Hausman, J., Daniel McFadden, D. 1984. "Specification Tests for the Multinomial Logit Model," Econometrica 52 (5), 1219-1240.

Heckman, J. 1981. Statistical models for discrete panel data. In D. McFadden and C. Manski (eds.), *The Econometrics of Panel Data*. MIT Press, Cambridge, MA, 114-178.

Hocking, R. R. 1996. Methods and Applications of Linear Models. New York: John Wiley & Sons.

Houser, D., J. Wooders. 2000. Reputation in Auctions: Theory and Evidence from eBay. Department of Economics Working Paper 00-01, The University of Arizona, Tucson, AZ

Hsiao, C. 1986. *Analysis of Panel Data*. Cambridge University of Press, New York.

Kennedy, P. 1998. A Guide to Econometrics, 4th Edition. Cambridge, MA: MIT Press.

Ledyard, J. 1995. Public goods: A survey of experimental research. In The Handbook of Experimental Economics. J. Kagel and A. Roth, eds. Princeton, NJ: Princeton University Press, 111-104.

Levine, D. 1998. Modeling Altruism and Spitefulness in Experiments. Review of Economic Dynamics 1, 593-622.

Livingston, J. 2002. How Valuable is a Good Reputation? A Sample Selection Model of Internet Auctions. Working Paper, University of Maryland, College Park, MD (available at: http://www.wam.umd.edu/~kth/reputation1.pdf.)

McDonald, C. G., V. C. Slawson, Jr. 2002. Reputation in an Internet Auction Market. Economic Inquiry 40 (3) 633-650.

Melnik, M. I., J. Alm 2002. Does a Seller's Reputation Matter? Evidence from eBay Auctions. Journal of Industrial Economics 50 (3) 337-349.

Neter, J. Kutner, M. H., Nachtsheim, C. J., Wasserman, W. 1996. Applied Linear Statistical

Models, Fourth Edition, Irwin Press, Chicago, IL.

Putnam, R. 2000. Bowling Alone: The Collapse and Revival of American Community. New York: Simon and Scurster.

Rabin, M. 1993. Incorporating Fairness into Game Theory and Economics. American Economic Review 83 (5), 1281-1302.

Resnick, P., Zeckhauser, R., Friedman, E., Kuwabara, K. 2000. Reputation Systems. Communications of the ACM 43 (12) 45-48.

Resnick, P. and Zeckhauser, R., 2002. Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. In The Economics of the Internet and E-Commerce, Michael R. Baye, editor. Volume 11 of Advances in Applied Microeconomics. Amsterdam, Elsevier Science.

Ribar, D. C. and Wilhelm, M. O. 2002. Altruistic and Joy-of-Giving Motivations in Charitable Behavior. Journal of Political Economy 110 (2). 425-457.

Wood, C., M. Fan, and Y. Tan. 2003. An examination of a reputation system for online auctions. Workshop for Information Systems and Economics, Barcelona, Spain.

APPENDIX A: OVERVIEW OF EBAY'S FEEDBACK MECHANISM

Founded in September 1995, eBay is the leading online marketplace for the sale of goods and services by a diverse community of individuals and businesses. Most items on eBay are sold through English auctions[13]. A typical eBay transaction begins with the seller listing an item he has on sale, providing an item description (including text and optionally photos), a starting bid, an optional reserve price and an auction closing date/time. Buyers then place bids for the item up until the auction closing time. The highest bidder wins the auction. The winning bidder sends payment to the seller. Finally, the seller sends the item to the winning bidder.

The above mechanism incurs significant risks. Sellers can exploit the underlying information asymmetries to their advantage by misrepresenting an item's attributes or by failing to complete the transaction. Buyers can renege on their commitment to buy the items of the auctions they have won.

To address these problems, eBay uses online feedback as its primary trust building mechanism[14]. More specifically, following completion of a transaction, both the seller and the buyer are encouraged to rate one another. A rating designates a transaction as *positive*, *negative*, or *neutral*, together with a short text comment. eBay aggregates all ratings posted for a member into that member's *feedback profile*. An eBay feedback profile consists of four components (Figures 5 and 6):

A. A member's *overall profile makeup*: a listing of the sum of positive, neutral and negative ratings received during that member's entire participation history with eBay.
B. A member's *summary reputation score* equal to the sum of positive ratings received by *unique* users minus the number of negative ratings received by *unique* users during that member's entire participation history with eBay. Repeat ratings from the same users do not count towards a member's summary score.
C. A member's "*eBay ID Card,*" which displays the sum of positive, negative and neutral ratings received during the most recent six month period (further subdivided into ratings received during the past week, month and past six months).

---

[13] eBay also supports Dutch auctions but these are rarely used.

[14] In addition to its feedback mechanism, eBay offers its members the option of using escrow services at extra cost. However, so far the percentage of transactions that opt for the use of those services is very low.

D. The *complete ratings history*, listing each individual rating and associated comment posted for a member in reverse chronological order.

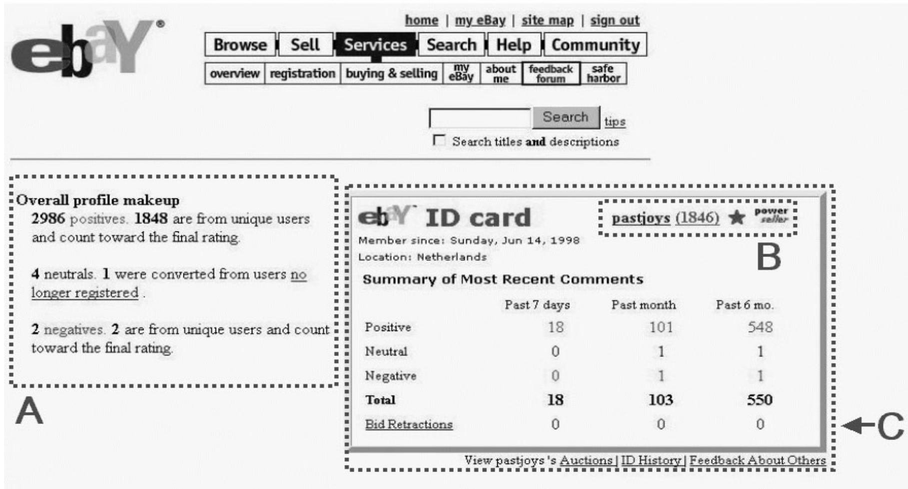

Figure 5. eBay Member Profile Summary

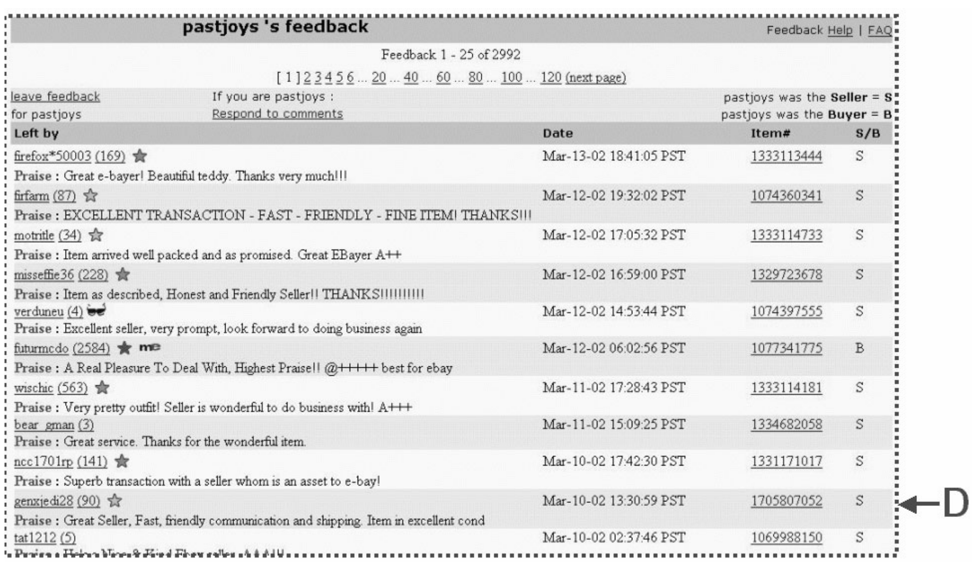

Figure 6 Detailed Feedback History

Seller feedback profiles are easily accessible from within the description page of any item for sale. More specifically, all item descriptions prominently display the seller's eBay ID, followed by his summary reputation score (component B in Figure 5). By clicking on the summary feedback score, prospective buyers can access the seller's full feedback profile (components A, B and C) and can then scroll through the seller's detailed ratings history (component D).

For a more detailed description of eBay's feedback mechanism, see Resnick and Zeckhauser (2002).

APPENDIX B: PROOFS

Proposition 1: Suppose that the fraction of altruists in the population of users is a, whereas the fraction of altruists in the population of partners is b. Suppose, further, that users and partners are randomly matched. A fraction p of altruists (corresponding to a fraction pa of the user population and a fraction pb of the partner population) is expected to submit a rating independent of their partner's action. The probability that an altruist partner is paired with a reciprocator user is (1-a). The expected fraction of reciprocator users who receive a rating from their partners is thus equal to pb(1-a). The probability that a reciprocator who receives a rating will respond is q, leading to an expected fraction qpb(1-a) of reciprocator users who submit a rating. Therefore, if $p<q$, the overall fraction of users who are expected to submit a rating is equal to:

$$\Pr[UserRates] = pa + qpb(1-a) < q(1-a) \tag{B.1}$$

The conditional probability that a user will rate (second) given that her partner has rated first is equal to:

$$\Pr[UserRates \mid PartnerRatedFirst] = \frac{\Pr[UserRatesSecond]}{\Pr[PartnerRatedFirst]}$$

A user rates second either if she is a reciprocator or if she is an altruist whose partner happened to be faster than her. We have previously established that the expected fraction of reciprocators who submit a rating is equal to qpb(1-a). Let λpa, $0 \le \lambda \le 1$ be the fraction of altruist users who rate after their partners have rated them. This gives Pr[$UserRatesSecond$]= λpa + qpb(1-a). Finally, only altruist partners will rate first. Let Pr[$PartnerRatedFirst$]=μpb, $0 \le \mu \le 1$. Therefore:

$$\begin{aligned} &\Pr[UserRates \mid PartnerRatedFirst] = \\ &\qquad \frac{\Pr[UserRatesSecond]}{\Pr[PartnerRatedFirst]} = \frac{\lambda pa + qpb(1-a)}{\mu pb} \ge q(1-a) \end{aligned} \tag{B.2}$$

By combining (B.1) and (B.2) we get the desired result:

$\Pr[UserRates] < \Pr[UserRates \mid PartnerRatedFirst]$. **QED**

Proposition 2: The proof follows similar logic to the proof of Proposition 1. To simplify the notation we assume that the fraction of self-interested traders in both populations is s, and that users and partners are randomly matched[15]. We, finally, assume that the sequence of events is as follows: (i) a fraction λ ($0 \le \lambda \le 1$) of self-interested users and partners (corresponding to a fraction λs of each population) simultaneously rate first, (ii) a fraction q of reciprocators who received a rating in phase (i) respond. The probability that a self-interested partner is paired with a reciprocator user is $(1-s)$. The expected fraction of reciprocator users who receive a rating from their partners is thus equal to $\lambda s(1-s)$, leading to an expected fraction $q\lambda s(1-s)$ of users who rate second because of reciprocation. Therefore, the overall fraction of users who are expected to submit a rating is $\Pr[UserRates]=\lambda s+q\lambda s(1-s)$. The conditional probability of a (reciprocator) user rating second given that the partner has rated first is:

$$\Pr[UserRates \mid PartnerRatedFirst] = \frac{\Pr[UserRatesSecond]}{\Pr[PartnerRatedFirst]} = \frac{q\lambda s(1-s)}{\lambda s} = q(1-s)$$

The relative magnitude of $\lambda s+q\lambda s(1-s)$ and $q(1-s)$ depends on the magnitude of the fraction of self-interested users s. For s sufficiently close to one, $\lambda s+q\lambda s(1-s) > q(1-s)$ and, thus,

$\Pr[UserRates] > \Pr[UserRates \mid PartnerRatedFirst]$. **QED**

[15] The result does not change if we assume that the fraction of self-interested agents is different in each population.