

Analyzing the Economic Efficiency of eBay-like Online Reputation Reporting Mechanisms

Chrysanthos Dellarocas
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA 02139, USA
dell@mit.edu

ABSTRACT

This paper introduces a model for analyzing marketplaces, such as eBay, which rely on binary reputation mechanisms for quality signaling and quality control. In our model sellers keep their actual quality private and choose what quality to advertise. The reputation mechanism is primarily used to determine whether sellers advertise truthfully. Buyers may exercise some leniency when rating sellers, which needs to be compensated by corresponding strictness when judging sellers' feedback profiles. It is shown that, the more lenient buyers are when rating sellers, the more likely it is that sellers will find it optimal to settle down to steady-state quality levels, as opposed to oscillating between good quality and bad quality. Furthermore, the fairness of the market outcome is determined by the relationship between rating leniency and strictness when assessing a seller's feedback profile. If buyers judge sellers too strictly (relative to how leniently they rate) then, at steady state, sellers will be forced to understate their true quality. On the other hand, if buyers judge too leniently then sellers can get away with consistently overstating their true quality. An optimal judgment rule, which results in outcomes where at steady state buyers accurately estimate the true quality of sellers, is analytically derived. However, it is argued that this optimal rule depends on several system parameters, which are difficult to estimate from the information that marketplaces, such as eBay, currently make available to their members. It is therefore questionable to what extent unsophisticated buyers are capable of deriving and applying it correctly in actual settings.

Keywords

Electronic commerce, reputation systems, trust in electronic markets.

1. INTRODUCTION

Online reputation reporting systems are emerging as an important quality signaling and quality control mechanism in online trading communities (Kollock 1999; Resnick et. al. 2000). Reputation systems collect feedback from members of an online community regarding past transactions with other members of that

community. Feedback is analyzed, aggregated and made publicly available to the community in the form of member *feedback profiles*. If one accepts that past behavior is a relatively reliable predictor of future behavior, then these profiles can act as a powerful quality signaling and quality control mechanism, essentially acting as the digital equivalent of a member's *reputation*.

eBay relies on its reputation mechanism almost exclusively in order to both produce trust and induce good behavior on the part of its members. eBay buyers and sellers are encouraged to rate one another at the end of each transaction. A rating can be a designation of "praise", "complaint" or "neutral", together with a short text comment. eBay makes the sums of praise, complaint and neutral ratings submitted for each member, as well as all individual comments, publicly available to all its users. Anecdotal and empirical results seem to demonstrate that eBay's reputation system has managed to provide remarkable stability in an otherwise very risky trading environment (Dewan and Hsu 2001; Resnick and Zeckhauser 2001).

The rising practical importance of online reputation systems not only invites but rather necessitates rigorous research on their functioning and consequences. Are such mechanisms truly reliable? Do they promote efficient market outcomes? To what extent are they manipulable by strategic buyers and sellers? What is the best way to design them? How should buyers (and sellers) use the information provided by such mechanisms in their decision-making process? This is just a small subset of unanswered questions, which invite exciting and valuable research.

The study of reputation as a mechanism for inducing good behavior in markets with asymmetric information is certainly not new. Several economists have published important works analyzing its properties (Rogerson, 1983; Schmalensee, 1978; Shapiro, 1982; Smallwood and Conlisk, 1979; Wilson, 1985, just to name a few).

Nevertheless, although past work in economics has studied some of the overall effects of reputation, it has paid very little attention to the analysis of *specific* mechanisms for forming and communicating reputation, in part because in traditional brick and mortar societies such mechanisms are largely informal (they are often referred to as "word-of-mouth advertising") and defy detailed modeling. The few published results focusing on the effects of specific properties of reputation mechanisms clearly make the point that such properties can have significant effects on the market outcome. For example, Rogerson (1983) shows that

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'01, October 14-17, 2001, Tampa, Florida, USA.

Copyright 2001 ACM 1-58113-387-1/01/0010...\$5.00.

reputation based on subjective binary ratings (e.g. good/bad, praise/complaint) creates an externality, which affects the entire market. Shapiro (1982) shows that, unless the mechanism by which reputation is formed satisfies certain properties, sellers may find it optimal to continuously oscillate in quality, periodically building good reputation and subsequently milking it.

On the other hand, the design and implementation of online reputation systems has so far been the research domain of computer scientists (see Bresee et. al., 1998; Sarwar et. al., 2000; Schafer et. al. 2001 for overviews of past work). The emphasis of past work in the area has been on developing algorithms and systems for collecting, aggregating and extracting useful information from sets of user ratings, drawing from work in information retrieval, data mining and collaborative filtering. The analysis and evaluation of the proposed algorithms is typically done in terms of computational complexity and statistical metrics, such as their running time, memory requirements, average recall and precision, average bias, etc.

We believe that there is a need for work that bridges the two disciplines: research, which takes into account the algorithmic details of specific reputation systems but also models how these systems are embedded inside trading communities and investigates their effectiveness and impact, not only in terms of computational and statistical properties, but rather in terms of their overall impact in the efficiency of the market and the welfare of the various classes of market participants. Given that reputation systems were conceived in order to assist choice in environments of imperfect information, their impact in those latter market dimensions should be the ultimate determinant of success of any new proposed new algorithm and system.

This paper contributes in this direction by proposing a model for analyzing the economic efficiency of binary reputation systems, such as the one used by eBay. Section 2 introduces the model and its underlying assumptions. Following that, Section 3 defines a number of properties that reputation mechanisms should satisfy in such settings, in order to be considered well functioning. Section 4 applies our model in order to determine under what circumstances binary reputation systems can indeed be well functioning. Our main conclusion is that binary reputation mechanisms can, in theory, be well functioning, provided that buyers strike the right balance between rating leniency and seller assessment strictness. However, getting this balance right is difficult without additional information, which is not currently provided by eBay. Section 5 considers the implications of relaxing some of the simplifying assumptions on which our analysis is based. Finally, Section 6 summarizes the contributions and conclusions of the paper.

2. BINARY REPUTATION MECHANISMS

This section introduces a model for analyzing marketplaces, which rely exclusively on a binary reputation mechanism for quality signaling and quality control. A *binary reputation mechanism* is a mechanism where raters are given the opportunity to rate past transactions using one of two values, commonly interpreted as “positive” (i.e. satisfactory) and “negative” (i.e. unsatisfactory, problematic). Our intention is to use this model in order to study the economic impact of reputation mechanisms

similar to the one used by eBay (see Resnick and Zeckhauser 2001 for a detailed description)¹.

In our model, qualities are non-negative real-valued quantities, which subsume aspects of both product quality and service quality. We assume that each seller produces items, whose *real quality* q_r is unknown to buyers and can only be determined with accuracy after consumption. We further assume that all buyers prefer higher quality to lower quality, although they might differ in the extent to which they are prepared to pay for an extra unit of quality. Finally, we assume that although the real quality of items is not communicated to buyers, sellers do inform buyers by advertising. On eBay, advertising corresponds to the seller-supplied description, which accompanies all items. The *advertised quality* q_a of an item is completely controlled by the seller (i.e. there is no validation of any kind by a third party) and may or may not correspond to its real quality.

Sellers aim to maximize the present value of their payoff function $\pi(x, q_r, q_a) = G(x, q_r, q_a) - c(x, q_r)$ where x is the volume of sales, $G(\cdot)$ is the gross revenue function and $c(\cdot)$ is the cost function. We assume that $\partial c / \partial q_r \geq 0$ and $\partial \pi / \partial q_a \geq 0$ for all sellers.

Under the above assumptions, sellers have an incentive to over-advertise quality. The market would then degenerate to a “market for lemons” (Akerlof 1970). In order to avoid this from happening, buyers are given the option to rate each transaction using a “positive” or “negative” rating. A reputation system, operated by a trustworthy third party, accumulates all ratings into a feedback profile $\mathbf{R} = (\Sigma_+, \Sigma_-, \Sigma_{\text{no rating}})$ for each seller, where Σ_+ is the sum of all positive ratings received for that seller during the most recent time window, Σ_- is the sum of all negative ratings received during the same period and $\Sigma_{\text{no rating}}$ is the number of transactions for which no rating was submitted². Time windowing is used in order to address the possibility that sellers may improve or deteriorate their behavior over time. For example, on eBay, feedback profiles display the sums of ratings received during the past 6 months only.

Buyer utility from purchase of a single item is modeled by $U = \theta \cdot q - p$, where p is the price, θ is a buyer’s quality sensitivity and q is the level of quality perceived by the buyer *after* consumption. When considering a purchase, buyers combine all the information that is available to them, i.e. an item’s

¹ In addition to “positive” (praise) and “negative” (complaint) ratings, eBay’s reputation mechanism also supports “neutral” ratings (which, however, are rarely used in actual practice). As will become apparent below, our model subsumes raters who would submit “neutral” ratings on eBay into the set of raters who don’t submit any rating at all.

² eBay does *not* currently publish $\Sigma_{\text{no rating}}$. The results of this paper make a strong case that they should.

advertised quality and a seller's feedback profile, in order to form a subjective assessment of an item's *estimated quality* q_e , where:

$$q_e = f(q_a, \mathbf{R}) \quad (1)$$

Armed with knowledge of prices and estimated qualities, buyers proceed to purchase one of the available items, presumably the one which maximizes their expected utility $U_e = \theta \cdot q_e - p$. Following a purchase, buyers observe the item's perceived quality $q = q_r + \varepsilon$, where ε is a normally distributed error term with standard deviation σ . The introduction of an error term is intended to collectively model a number of phenomena, which occur in actual practice. For example:

- buyers may misinterpret a seller's advertised quality (this should be modeled as $q_e = q_a + \varepsilon$, however, our analysis is identical if we add the error term to q_r instead)
- sellers may exhibit small variations in actual quality from one transaction to another
- buyers may have small differences in their perception of quality based, say, on their moods that day
- some aspects of perceived quality depend on factors beyond a seller's control (e.g. post-office delays)

Finally, buyers decide whether to rate a transaction as well as what rating to give. Our model assumes that ratings are a function of a buyer's satisfaction relative to her expectations. We define a buyer's satisfaction from a given transaction to be the difference between perceived and expected utility. That is, $S = U - U_e = \theta \cdot (q_r - q_e + \varepsilon)$. Under the above assumptions, S is a normally distributed random variable with mean $\theta \cdot (q_r - q_e)$ and standard deviation $\theta \cdot \sigma$.

One interesting property of eBay, which has been reported on several empirical studies, is that most buyers give very few negative ratings to sellers. Resnick and Zeckhauser (2001) have speculated that there are several reasons for this: eBay allows reciprocal ratings (that is, sellers also rate buyers) and buyers are often afraid that posting a negative rating for a seller will lead to retaliatory bad ratings, harassing emails etc. eBay does not provide mechanisms to prevent or assist such situations. Furthermore, it has been reported that sellers often communicate with buyers via email and negotiate settlements to transaction problems, while explicitly pleading with them to not post negative ratings. Finally, eBay has created a "culture of praise", where the vast majority of ratings and comments are extremely positive. In such a setting, most buyers feel a moral obligation to conform to the prevailing social norms and be nice and relatively forgiving to their trading partners.

Our model uses a rating function $r(S)$, which attempts to model the above empirically observed behavior. More specifically, we are assuming that buyers rate a transaction as "positive" if their actual utility from the transaction exceeds their expected utility (i.e. if $S > 0$). On the other hand, buyers only rate a transaction as negative if their actual utility falls short of their expected utility by more than a *leniency factor* λ , that is, if $S < -\lambda$. Finally, for transactions, which end up being "slightly bad but not too bad"

(i.e. where $-\lambda < S \leq 0$), we are assuming that buyers prefer to simply refrain from rating at all³. To summarize:

$$r(S) = \begin{cases} "+" & \text{if } S > 0 \\ "-" & \text{if } S \leq -\lambda \\ \text{no rating} & \text{if } -\lambda < S \leq 0 \end{cases} \quad (2)$$

where $S = U - U_e = \theta \cdot (q_r - q_e + \varepsilon)$, $\varepsilon \sim N(0, \sigma)$,

To simplify the initial analysis, we are making the assumption that θ, σ and λ are constant across the entire population of buyers and sellers. In Section 5, we will relax those assumptions and study how they impact the results derived in Sections 3 and 4.

3. WELL FUNCTIONING REPUTATION MECHANISMS

The following sections will use the model developed in Section 2, in order to explore under what circumstances binary reputation mechanisms can be *well functioning*. Before doing that, however, in this section we will discuss what it means for a reputation mechanism to be well functioning in marketplaces with private quality information. We *define* a well-functioning reputation mechanism to be one, which satisfies the following two properties:

WF1: If there exists an equilibrium of prices and qualities under perfect information (i.e. in settings where $q_e = q_a = q_r$) then, in environments where q_r is private to sellers, the existence of the reputation mechanism makes it optimal for sellers to settle down to a steady-state pair of real and advertised qualities, rather than to oscillate, successively building up and milking their reputation.

WF2: Assuming WF1 holds, under all steady-state seller strategies (q_r, q_a) the quality of sellers as estimated by buyers *before* transactions take place, is equal to their true quality (i.e. $q_e = q_r$).

Before we proceed, let us justify the above definition by providing a brief rationale for the desirability of properties WF1 and WF2.

First, the value of reputation mechanisms in general relies on the assumption that past behavior is a reliable predictor of future behavior (Wilson 1985). If oscillations were optimal, the predictive value of cumulative functions of past ratings, such as Σ_+, Σ_- , would be greatly diminished. In environments where the primary (or only) mechanism for certifying and controlling seller quality is based on reputation, it is, therefore, desirable that sellers find it optimal to settle down to a steady-state behavior rather than to oscillate.

Second, according to our model, buyers make purchase decisions based on knowledge of prices and estimated qualities. If it were possible for sellers to settle down into a steady-state strategy that would consistently deceive buyers into estimating $q_e > q_r$, then

³ On eBay, some buyers would post a "neutral" rating in this case.

this would allow sellers to earn additional profits at the expense of buyers. In the presence of competitive marketplaces, buyers would then eventually leave the marketplace in favor of other markets with better information. On the other hand, if, under all possible steady-state seller strategies the effect of the reputation mechanism was such, so that buyers estimated $q_e < q_r$, then the opposite effect would take place: buyers would realize extra surplus at the expense of sellers. Once again, we would then expect that sellers would desert the marketplace in favor of other, more transparent markets. The only fair steady-state strategy, therefore, is one where $q_e = q_r$.

4. CAN BINARY REPUTATION MECHANISMS BE WELL FUNCTIONING?

This section will demonstrate that, given a rating function, which has the general form given by (2), whether an eBay-like binary reputation system satisfies property WF2 depends on the relationship between (2) and the quality estimation function $q_e = f(q_a, \mathbf{R})$. Furthermore, we will show that, if buyers are lenient enough when they rate and correspondingly strict when they judge seller profiles, sellers will find it optimal to settle down to steady-state true and advertised quality levels if such an equilibrium exists under perfect information.

4.1 Estimated vs. real qualities in steady state

Let us first focus our attention on the circumstances under which a binary reputation mechanism satisfies condition WF2. We are assuming that WF1 holds. Therefore, there exists at least one steady-state strategy (q_r, q_a) for each seller⁴. A steady-state strategy is a strategy that optimizes a seller's payoff function, while at the same time resulting in an estimated quality $q_e = f(q_a, \mathbf{R})$, which is stable over time. Denote $q_e = q_r + \xi$, where ξ is the *deception factor*, that is, the distortion between estimated and real quality at steady state. If $\xi > 0$ then buyers overestimate a seller's true quality, whereas if $\xi < 0$ then buyers underestimate true quality. Let N be the total number of sales transactions of a given seller in the most recent time window. It is easy to see that $N = \Sigma_+ + \Sigma_- + \Sigma_{no\ rating}$. Assuming that buyers rate according to (2), for large N at steady state the following will hold:

$$\begin{aligned} \Sigma_+ &= N \cdot \text{Prob}[S > 0] = N \cdot \Phi[(q_r - q_e)/\sigma] = N \cdot \Phi[-\xi/\sigma] \\ \Sigma_- &= N \cdot \text{Prob}[S \leq -\lambda] = N \cdot \Phi[(q_e - q_r)/\sigma - \lambda/(\theta \cdot \sigma)] \\ &= N \cdot \Phi[\xi/\sigma - \lambda/(\theta \cdot \sigma)] \end{aligned} \quad (3)$$

where $\Phi(\cdot)$ is the standard normal CDF.

Given that $q_e = f(q_a, \mathbf{R})$, satisfaction of condition WF2 depends on the quality assessment function f . More specifically, f

must be chosen so that, for all steady-state strategies (q_r, q_a) the equation:

$$q_e = q_r + \xi = f(q_a, \mathbf{R}(\xi)) \quad (4)$$

has a unique solution at $\xi = 0$.

eBay does not specify, or even recommend, a specific quality assessment function f . It simply publishes the quantities Σ_+ and Σ_- for each seller and allows buyers to use any assessment rule they see fit. It is important to note at this point that eBay does not currently publish the quantity $\Sigma_{no\ rating}$ (and therefore N) for a seller. As we will show below, knowledge of N is essential for constructing reliable quality assessment functions. The results of this paper, therefore, make a strong argument that the number of transactions that have received no rating should be added to the profile information published by eBay and similar systems.

Our objective in this paper is to explore whether *in principle*, binary reputation systems can be well functioning. Therefore, our aim is to explore the *existence* of a quality assessment function f which, when used in conjunction with rating rule (2), satisfies WF2.

In the rest of the paper, we will explore the suitability of the following family of quality assessment functions:

$$q_e = f(q_a, \mathbf{R}) = \begin{cases} q_a & \text{if } \hat{\xi}(\mathbf{R}) \leq 0 \\ 0 & \text{if } \hat{\xi}(\mathbf{R}) > 0 \end{cases} \quad (5)$$

where $\hat{\xi}(\mathbf{R})$ is some estimate of the seller's deception factor based on information contained in the seller's feedback profile.

Through function (5), buyers assess the quality of an item to be equal to that advertised by the seller, if, based on the seller's profile, they conclude that the seller *does not over-advertise*. Otherwise, buyers assume that the seller lies and assess minimum quality. Function (5) therefore uses the information provided by the reputation mechanism in order to derive a (binary) assessment of *truthfulness in advertising*.

It is easy to see that, if buyers have a way of reliably estimating the sign of ξ from feedback profile information and assess seller quality through function (5), sellers who over-advertise their quality will quickly see their estimated quality fall to zero. Therefore, if f is given by (5), equation (4) has no solution for $\xi > 0$.

Note that function (5) does not prevent sellers from under-advertising their quality because for $q_a = q_r + \xi$, all $\xi \leq 0$ are also solutions of equation (4). However, given that we have assumed that $\partial\pi/\partial q_a \geq 0$, we would not expect any profit-maximizing seller to under-advertise. Therefore, the only steady-state seller strategy for sellers would be to truthfully advertise their real quality. In that case, buyers would estimate $q_e = q_a = q_r$, a desirable outcome, which satisfies WF2.

⁴ Section 4.2 will explore the conditions under which sellers will indeed find it optimal to settle down to a steady state strategy.

In conclusion, a binary reputation system where buyers

- rate according to (2)
- assess item quality according to (5) and
- have a reliable rule for calculating $\hat{\xi}(\mathbf{R})$ for a given seller

satisfies WF2. Let us now explore three different ways in which buyers can use Σ_+ , Σ_- and N in order to estimate the sign of ξ .

Estimation based on the number of positives

One way to estimate seller honesty is to require that the fraction of positive ratings of good sellers exceed a threshold. From (3) we can see that $\hat{\eta} \equiv \Sigma_+ / N$ can be interpreted as a point estimator of $\Phi[-\xi / \sigma]$. Given that $\Phi[-\xi / \sigma] < 0.5$ for all $\xi > 0$, assessment of the sign of ξ reduces to testing the statistical hypothesis $H_0 : \eta \geq 0.5$ given $\hat{\eta}$. The corresponding quality assessment function then becomes:

$$q_e = \begin{cases} q_a & \text{if } H_0 \text{ accepted} \\ 0 & \text{if } H_0 \text{ rejected} \end{cases} \quad (6)$$

where $H_0 : \eta \geq 0.5$ given $\hat{\eta} \equiv \Sigma_+ / N$

Hypothesis H_0 can be tested using one of the known techniques for computing confidence intervals of proportions following binomial distributions (e.g. Blyth and Still 1983).

Function (6) is an appealing method for assessing seller quality because of its relative simplicity. Note that its computation does not require knowledge of the model parameters λ, θ and σ . However, (6) is difficult to compute reliably without knowledge of N , the total number of rated plus unrated transactions of a seller. As was mentioned, eBay does not make N known to its members. Taking $\hat{\eta} \equiv \Sigma_+ / (\Sigma_+ + \Sigma_-)$ would result in large overestimation of $\Phi[-\xi / \sigma]$, especially because of the rating leniency factor. For that reason, one would infer that quality assessment based on the number of positive ratings is not (and *should not* be) widely used on eBay. This hypothesis is consistent with empirical observations (Dewan and Hsu 2001).

Section 4.2 will discuss another disadvantage of function (6), which is that it makes it easier for sellers to oscillate between periods where they milk their good reputation by overstating their quality and deceiving buyers and periods where they restore their reputation by offering better quality than what buyers expect.

Estimation based on the number of negatives

In an analogous manner, we expect good sellers to have few negative ratings. Therefore, another way to estimate seller honesty is to require that the fraction of negative ratings of good sellers stay below a threshold. From (3) we can see that $\hat{\zeta} \equiv \Sigma_- / N$ can be interpreted as a point estimator of $\Phi[\xi / \sigma - \lambda / (\theta \cdot \sigma)]$. Given that $\Phi[\xi / \sigma - \lambda / (\theta \cdot \sigma)] > \Phi[-\lambda / (\theta \cdot \sigma)]$ for all $\xi > 0$, assessment of the sign of ξ reduces to testing the

statistical hypothesis $H'_0 : \zeta \leq \Phi[-\lambda / (\theta \cdot \sigma)]$ given $\hat{\zeta}$. The corresponding quality assessment function then becomes:

$$q_e = \begin{cases} q_a & \text{if } H'_0 \text{ accepted} \\ 0 & \text{if } H'_0 \text{ rejected} \end{cases} \quad (7)$$

where $H'_0 : \zeta \leq k^* \equiv \Phi[-\lambda / (\theta \cdot \sigma)]$ given $\hat{\zeta} \equiv \Sigma_- / N$

Let us call $k^* \equiv \Phi[-\lambda / (\theta \cdot \sigma)]$ the optimum trustworthiness threshold. k^* is a monotonically decreasing function of the leniency factor λ . Therefore, the more lenient buyers are when they rate, the lower the threshold of negative ratings to transactions above which they should not trust sellers, and vice versa. This is a result that corresponds well to documented empirical findings: most eBay buyers weigh negative ratings much more heavily than positive ratings when assessing the trustworthiness of a prospective seller (Dewan and Hsu 2001). Given that they seem to be rather lenient when they rate those sellers, according to (7), we would expect them to be strict when assessing the quality of sellers, and therefore to be relatively intolerant of negative ratings.

From equation (7) we can also see that, *in theory*, buyers will complete knowledge of the system parameters λ, θ and σ can derive an optimum k^* for every λ ⁵. One way of interpreting this result is that satisfaction of WF2 is always possible no matter how lenient (or strict) buyers are when they rate, *provided that they strike the right balance between rating leniency and quality assessment strictness*. In the next section, we shall prove that, more lenient rating (and correspondingly strict assessment) increases the likelihood that sellers will find it optimal to settle down to a steady-state behavior. Some degree of leniency, therefore, can be beneficial to the stability of the marketplace.

It is also important to point out that, unless buyers use the right threshold k^* when evaluating the number of negative ratings of a seller, WF2 will not be satisfied. If buyers use a threshold $k > k^*$ then there will be some $\xi > 0$ for which H'_0 will be satisfied and sellers will be able to consistently deceive buyers by over-advertising their quality. In contrast, if $k < k^*$, there will be $\xi_0 < 0$ such that H'_0 will be rejected for all $\xi > \xi_0$. In the latter case, to prevent their estimated quality from dropping to zero, sellers will be forced to under-advertise and, therefore, be consistently under-appreciated by $|\xi_0|$.

We see, therefore, that the choice of the “right” k^* is crucial to the well functioning of the reputation mechanism, and of the marketplace in general. It is important to ask whether buyers can

⁵ Although this conclusion establishes the *theoretical* well functioning of binary reputation mechanisms, in practice, typical buyers will not have knowledge of the parameters needed to estimate k^* . Please read on for a discussion of the implications of this.

be reasonably expected to be able to correctly derive it. From equation (7), calculation of k^* requires knowledge of the model parameters λ, θ and σ . It is unlikely that buyers would have accurate understanding and knowledge of those parameters (especially σ , which partly reflects properties of the seller). Nevertheless, even if the model parameters are not known, it is possible to estimate the value of $\Phi[-\lambda/(\theta \cdot \sigma)]$ from Σ_+, Σ_- and N . From (3):

$$\left. \begin{aligned} -\xi/\sigma &= \Phi^{-1}(\Sigma_+/N) \\ -\lambda/(\theta \cdot \sigma) + \xi/\sigma &= \Phi^{-1}(\Sigma_-/N) \end{aligned} \right\} \Rightarrow \quad (8)$$

$$k^* = \Phi[-\lambda/(\theta \cdot \sigma)] = \Phi[\Phi^{-1}(\Sigma_+/N) + \Phi^{-1}(\Sigma_-/N)]$$

If N is small then a confidence interval should be constructed for k^* .

Even with the help of equation (8), buyers still need to know N in order to properly compute function (7). Overall, function (7) defines a rather fragile rule for assessing seller quality efficiently. Given that a lot of eBay buyers are heavily basing their seller quality assessments on the number of negative ratings on the sellers' feedback profile, it is very interesting to ask what methods they use to compute their trustworthiness thresholds and, even more important, whether their trustworthiness thresholds do indeed come close to satisfying WF2. Clearly, these are important questions, which invite further empirical and experimental results to complement the results of this work.

Estimation based on the ratio between negatives and positives

In both previous cases, correct implementation of the quality assessment function required knowledge of N . One might think that, by basing quality assessment on the ratio between negative and positive one may be able to derive an optimal assessment function from Σ_+ and Σ_- only. We will show that this is not possible. From (3) we get:

$$\rho(\xi) \equiv \frac{\Sigma_-(\xi)}{\Sigma_+(\xi)} = \frac{\Phi[\xi/\sigma - \lambda/(\theta \cdot \sigma)]}{1 - \Phi[\xi/\sigma]} \quad (9)$$

Function $\rho(\xi)$ is non-negative and monotonically increasing in ξ . Furthermore $\rho(0) = 2 \cdot \Phi[-\lambda/(\theta \cdot \sigma)]$. Since $\rho(\xi) > \rho(0)$ for all $\xi > 0$, assessment of the sign of ξ reduces to testing the statistical hypothesis $H_0^* : \rho \leq 2 \cdot \Phi[-\lambda/(\theta \cdot \sigma)]$ given $\hat{\rho} \equiv \Sigma_- / \Sigma_+$. The corresponding quality assessment function then becomes:

$$q_e = \begin{cases} q_a & \text{if } H_0^* \text{ accepted} \\ 0 & \text{if } H_0^* \text{ rejected} \end{cases} \quad (10)$$

where $H_0^* : \rho \leq 2 \cdot \Phi[-\lambda/(\theta \cdot \sigma)]$ given $\hat{\rho} \equiv \Sigma_- / \Sigma_+$

Unless buyers have knowledge of the model parameters λ, θ and σ , calculation of $2 \cdot \Phi[-\lambda/(\theta \cdot \sigma)]$ from (8) requires knowledge of N . Therefore, using the ratio of negatives to positives is very similar to using the fraction of negatives and is equally tricky to "get right" without knowledge of N .

4.2 Existence of steady-state behavior

The analysis of Section 4.1 has been based on the assumption that sellers settle down to steady-state real and advertised quality levels. This section will investigate the conditions under which sellers will indeed find it optimal to do so. The alternative is to oscillate between building a good reputation and then milking it by over-advertising real quality. As we argued in Section 3, reputation-mediated marketplaces should be designed in order to induce sellers to settle down to steady state behavior (otherwise information about past behavior will not be very helpful as a way of predicting the future).

The principal result of this section is that when quality assessment is based on functions (7) or (10), which involve negative ratings, then, if the rating leniency factor λ is large enough, sellers will find it optimal to settle down to steady state behavior. In contrast, there is no such guarantee when quality assessment is based on function (6), which only involves positive ratings. This result shows that more lenient rating (coupled with more strict quality assessment) supports stability in the system. For the same reason, although more fragile and difficult to "get right", functions (7) and (10), i.e. functions which base seller quality assessment on the number of negative ratings, are preferred to function (6), which only looks at the seller's positive ratings.

In order to derive our result, let us consider ways in which sellers may attempt to realize additional profits through oscillating behavior. Assume that a seller is able to perform N_1 transactions before ratings of those transactions are posted to her feedback profile. This number depends on the frequency of transactions and the delay between transactions and the posting of ratings by buyers (on eBay, this delay is typically 2-3 weeks).

Let us consider a seller who, at the end of period 0, has completed N transactions in the current time window and has accumulated a good reputation, by producing and advertising items of quality q^* , the quality that optimizes profits assuming steady-state behavior. Let us further assume that buyers assess quality based on function (7). At the end of period 0:

$$\left. \frac{\Sigma_-}{N} \right|_{p_0} = \frac{\Sigma_-}{N} (\xi = 0) = \Phi[-\lambda/(\theta \cdot \sigma)] = k^* \quad (11)$$

At the beginning of period 1 the seller decides to milk her reputation by choosing a real quality q' and then over-advertising her quality by ξ_1 so that her profit is maximized relative to the steady state case. Given the seller's good past reputation, initially buyers will be deceived. However, after they purchase the seller's items, they will realize their inferior quality and will post proportionally more negative ratings. Therefore, at the end of period 1 (after N_1 "deceiving" transactions):

$$\left. \frac{\Sigma_-}{N} \right|_{p_1} = \frac{N \cdot \Phi[-\lambda/(\theta \cdot \sigma)] + N_1 \cdot \Phi[\xi_1/\sigma - \lambda/(\theta \cdot \sigma)]}{N + N_1} > k^* \quad (12)$$

and the seller's subsequent estimated quality will fall to zero. Assuming that some buyers are willing to buy from somebody with zero quality if the price is low enough, our seller will stay in business. In order to increase her reputation once again, she needs

to reduce the ratio Σ_- / N to below the threshold k^* . The only way she can achieve this is to go through a period where she produces higher quality items but receives lower prices, surpassing buyers' expectations (who now expect $q_e = 0$) by ξ_2 . Let us assume that it would take N_2 "redeeming" transactions before $\Sigma_- / N \leq k^*$. At the end of period 2:

$$\frac{\Sigma_-}{N} \Big|_{P2} = \frac{N \cdot \Phi[-\lambda / (\theta \cdot \sigma)] + N_1 \cdot \Phi[\xi_1 / \sigma - \lambda / (\theta \cdot \sigma)] + N_2 \cdot \Phi[-\xi_2 / \sigma - \lambda / (\theta \cdot \sigma)]}{N + N_1 + N_2} = k^* = \Phi[-\lambda / (\theta \cdot \sigma)] \quad (13)$$

A profit-maximizing seller will choose to oscillate if the profit from the "deceiving" transactions relative to the steady-state profit exceeds the loss from the "redeeming" transactions relative to the steady-state profit. If these two quantities have a finite ratio, then, provided that the number N_2 of "redeeming" transactions that are necessary in order to "undo" the reputation effects of N_1 "deceiving" transactions is high enough, sellers will not find it profitable to oscillate and will settle down to steady-state real and advertised quality levels.

From (13) after some algebraic manipulation, we get:

$$\frac{N_2}{N_1} = \frac{\Phi[\xi_1 / \sigma - \lambda / (\theta \cdot \sigma)] - \Phi[-\lambda / (\theta \cdot \sigma)]}{\Phi[-\lambda / (\theta \cdot \sigma)] - \Phi[-\xi_2 / \sigma - \lambda / (\theta \cdot \sigma)]} = g(\lambda, \xi_1, \xi_2) \quad (14)$$

After some manipulation we get $\partial g / \partial \lambda > 0$ and $\partial^2 g / \partial \lambda^2 > 0$. In fact $g(\cdot)$ grows exponentially with λ ⁶. Figure 1 plots $g(\lambda)$ for $\theta = \sigma = 1$ and some representative values of $\xi = \xi_1 = \xi_2$.




Figure 1. Minimum ratio of "redeeming" to "deceiving" transactions needed in order to restore one's good reputation following a period of quality over-reporting.

⁶ Furthermore, for a given λ , $g(\cdot)$ grows rapidly with ξ_1 and decreases very slowly with ξ_2 . This means that the minimum necessary ratio of redeeming to deceiving transactions grows with the amount of initial deception (ξ_1) and cannot be significantly brought down by increasing the amount of redemption (ξ_2).

From Figure 1 it is evident that in marketplaces where buyers rate leniently (and assess quality strictly), sellers need many more "redeeming transactions" in order to restore their good reputation following a few "deceiving transactions". The relative number of redeeming transactions increases exponentially with the leniency factor. Otherwise said, the larger the λ , the more difficult it is for sellers to restore their reputation once they lose it. Consequently, if λ is sufficiently large, sellers will find it optimal to settle down to steady-state real and advertised quality levels. Q.E.D.

A similar result can be derived if buyers base quality assessment on the ratio Σ_- / Σ_+ . In contrast, if buyers base quality assessment on function (7), our analysis gives:

$$\frac{N_2}{N_1} = \frac{1 - 2 \cdot \Phi[\xi_1 / \sigma]}{1 - 2 \cdot \Phi[\xi_2 / \sigma]} \quad (15)$$

Equation (15) gives $N_2 = N_1$ for $\xi_1 = \xi_2$ and N_2 slightly less than N_1 for $\xi_1 < \xi_2$. Otherwise said, following a set of deceiving transactions, it takes the same number of (or fewer) redeeming transactions in order to restore one's good reputation. In such a setting, it is more likely that some sellers will have profit functions for which it will be optimal to oscillate. Therefore, one expects that in reputation-mediated marketplaces where buyers use (6) to assess seller quality, there will be less stability than in marketplaces where sellers use (7) or (10).

The results of this section provide some interesting arguments for both rating leniency as well as for basing the quality assessment of sellers on their negative, rather than their positive ratings.

5. REALITY CHECKS AND SOME RECOMMENDATIONS

The results of the previous sections have been derived by making a number of simplifying assumptions about buyer behavior. More specifically, we have assumed that all buyers have the same quality sensitivity θ and leniency factor λ . Furthermore, we have assumed that buyers *always* submit ratings whenever their satisfaction rises above zero or falls below $-\lambda$. Both assumptions are not likely to hold in a real marketplace. Buyers have different personalities, and therefore, are expected to have different quality sensitivities, as well as leniency parameters. Furthermore, ratings do incur a cost (time to log on and submit them) and some buyers do not bother rating, even when transactions turn out really good or very bad. In this section, we will inject a bit of reality to our model and will explore how our results change if we take into account the above considerations.

Reality Check #1: Some buyers never rate

We need to modify our rating function $r(S)$ in (2), so that when $S > 0$, $r(S) = "+"$ with probability β and $r(S) = \text{no rating}$ with probability $(1 - \beta)$. Similarly, when $S \leq -\lambda$, $r(S) = "-"$ with probability γ and $r(S) = \text{no rating}$ with probability $(1 - \gamma)$. Under this new rating function, the statistical hypothesis in (6) becomes $H_0 : \eta \geq \beta \cdot 0.5$, while the hypothesis in (7) becomes $H'_0 : \zeta \leq \gamma \cdot \Phi[-\lambda / (\theta \cdot \sigma)]$. We see

that our new assumption introduces two additional parameters to our model. The parameters need to be reliably estimated in order for property WF2 to be satisfied.

Reality Check #2: Buyers differ in quality sensitivity and leniency

Let's define $\omega \equiv \lambda/\theta$ and let's call $p(\omega)$ the probability distribution of ω among buyers. Then (3) must be modified as:

$$\begin{aligned}\Sigma_+ &= N \cdot \text{Pr ob}[S > 0] = N \cdot \Phi[-\xi/\sigma] \\ \Sigma_- &= N \cdot \text{Pr ob}[S \leq -\lambda] = N \cdot \int \Phi[(\xi - \omega)/\sigma] \cdot p(\omega) \cdot d\omega\end{aligned}\quad (16)$$

If quality assessment is based on the fraction of positive ratings using (5), then reality check #2 does not introduce additional complications. However, if quality assessment is based on the fraction of negative ratings, which, in the presence of lenient ratings is the rule most likely to result in stable seller behavior, then things do become considerably more complicated. More specifically, it is easy to see that the statistical hypothesis to be tested in (6) must become

$H'_0: \zeta \leq k^* \equiv \int \Phi[-\omega/\sigma] \cdot p(\omega) \cdot d\omega$. In order to calculate the "right" k^* , one needs knowledge of $p(\omega)$.

Things become even more complicated if we combine reality checks #1 and #2, which would be the situation that most closely corresponds to actual reality. Of course, one can begin to think of ways in which individual buyers might be able to estimate, maybe with some degree of error, the additional model parameters β, γ and $p(\omega)$ from Σ_+, Σ_- and N .

However, instead of embarking in this direction, at this stage we believe that we are provided enough arguments to make one of the main points of this paper: Binary reputation mechanisms can *in theory* be well functioning under the assumption of simple rating and assessment rules, but only if buyers use the right thresholds when judging seller trustworthiness. Calculating the right threshold from Σ_+, Σ_- alone, the only information currently provided by eBay, is very difficult. Calculating the right threshold from Σ_+, Σ_- and N is possible under the simplifying assumptions of Section 2 but becomes more and more difficult as our models approach reality. In realistic cases, the correct assessment rule depends not only on the feedback profile of a seller but also on properties of the rater population. Given that the efficiency of the marketplace crucially depends on the selection of correct assessment thresholds on the part of the buyer, the most sensible course of research therefore should be to think of additional information that the reputation mechanism can provide to raters, in order to make this calculation easier.

6. CONCLUSIONS

The objective of this paper was to explore to what extent binary reputation mechanisms, such as the one used at eBay, are capable of inducing efficient market outcomes in marketplaces where (a) true quality information is unknown to buyers, (b) advertised quality is completely under the control of the seller and (c) the only information available to buyers is an item's advertised quality plus the seller's feedback profile.

The first contribution of the paper is the definition of a set of conditions for evaluating the well functioning of a reputation mechanism in such settings. We consider a reputation mechanism to be well-functioning if it (a) induces sellers to settle down to a steady-state behavior assuming it is optimal for them to do so under perfect quality information and (b) at steady-state, seller quality as estimated by buyers *before* transactions take place is equal to their true quality.

The second contribution of the paper is an analysis of whether binary reputation mechanisms can be well-functioning under the assumptions that (a) ratings are based on the difference between buyers' true utility following a transaction and their expectations before the transaction and (b) buyers are relatively lenient when they rate and correspondingly strict when they assess a seller's feedback profile.

The first conclusion is that if binary feedback profiles are used to decide whether a seller advertises truthfully (in which case buyers assess quality equal to the advertised quality) or not (in which case buyers assess quality equal to the minimum quality), then, *in theory*, binary reputation systems can be well functioning, provided that buyers strike the right balance between rating leniency and quality assessment strictness. Furthermore, assuming that buyers base their judgment on the ratio of negative ratings received by a seller, if buyers are lenient enough when they rate and correspondingly strict when they judge seller profiles, we have shown that sellers will find it optimal to settle down to steady-state quality levels if such an equilibrium exists under perfect information. This is an interesting way in which (a) judging seller trustworthiness based on their negative ratings is preferable to basing it on their positive ratings and (b) some degree of rating leniency helps bring stability to the system.

The second conclusion is that, unless buyers use the "right" threshold parameters when they judge seller profiles, binary reputation mechanisms will not function well and the resulting market outcome will be unfair for either the buyers or the sellers. In that sense, although binary reputation mechanism can be well functioning in theory, they are expected to be quite fragile in practice. The crucial question therefore becomes whether binary feedback profiles provide sellers (esp. relatively unsophisticated ones) with enough information to derive the "right" seller judgment rules.

It was shown that the "right" judgment rule is difficult to infer correctly from knowledge of the sum of positive and negative ratings alone, which is the only information currently provided by eBay to its members. If knowledge of the sum of unrated transactions is added to feedback profiles, then, under a number of simplifying assumptions, it is possible to derive non-obvious but relatively simple "optimal" judgment rules which result in well functioning reputation mechanisms. However, if the simplifying assumptions are dropped, calculation of the right judgment rule from Σ_+, Σ_- and N once again becomes difficult, as it requires knowledge not only of seller ratings but of the rater population as well. These findings lead to the recommendation that more information should be provided to assist raters of such marketplaces use feedback profiles in the "right" way.

The theoretical results of this paper raise some intriguing questions related to the efficiency, fairness and stability of eBay-like electronic marketplaces. The author would welcome

experimental and empirical evidence that will shed more light into the questions raised and would validate the conclusions drawn from his models.

7. ACKNOWLEDGMENTS

This research was supported by NSF CAREER GRANT IIS-9984147 and an MIT eBusiness Vision Fund Award.

8. REFERENCES

Akerlof, G. (1970) The market for 'lemons': Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84, pp. 488-500.

Blyth, C.R. and Still H.A. (1983) Binomial Confidence Intervals. *Journal of the American Statistical Association* 78, pp. 108-116.

Breese, J.S., Heckerman, D., and Kadie, C. (1998) Empirical Analysis of Predictive Algorithms for Collaborative Filtering. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pp. 43-52, San Francisco, July 24-26, 1998.

Dewan, S. and Hsu, V. (2001) Trust in Electronic Markets: Price Discovery in Generalist Versus Specialty Online Auctions. Working Paper. January 31, 2001.

Kollock, P. (1999) The Production of Trust in Online Markets. In *Advances in Group Processes* (Vol. 16), eds. E.J. Lawler, M. Macy, S. Thyne, and H.A. Walker, Greenwich, CT: JAI Press.

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

Resnick, P. and Zeckhauser, R. (2001) Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. Working Paper for the NBER workshop on empirical studies of electronic commerce. January 2001.

Rogerson, W.P. (1983) Reputation and product quality. *Bell Journal of Economics*, Vol. 14 (2), pp. 508-16.

Sarwar, B. M., Karypis, G., Konstan, J. A., and Riedl, J. (2000). Analysis of Recommender Algorithms for E-Commerce. *ACM E-Commerce 2000 Conference*, Minneapolis, MN, Oct. 17-20, 2000.

Schmalensee, R. (1978). Advertising and Product Quality. *Journal of Political Economy*, Vol. 86, pp. 485-503.

Schafer, J.B., Konstan, J., and Riedl, J., (2001) Electronic Commerce Recommender Applications. *Journal of Data Mining and Knowledge Discovery*. January, 2001.

Shapiro, C. (1982) Consumer Information, Product Quality, and Seller Reputation. *Bell Journal of Economics* 13 (1), pp 20-35, Spring 1982.

Smallwood, D. and Conlisk, J. (1979). Product Quality in Markets Where Consumers Are Imperfectly Informed. *Quarterly Journal of Economics*. Vol. 93, pp. 1-23.

Wilson, Robert (1985). Reputations in Games and Markets. In *Game-Theoretic Models of Bargaining*, edited by Alvin Roth, Cambridge University Press, pp. 27-62.