The Dynamics of Seller Reputation: Theory and Evidence from eBay

Luís Cabral
New York University

Ali Hortaçsu
University of Chicago

Preliminary draft: June 2003

Abstract

We test specific implications of existing models of reputation using eBay data on auction prices and seller reputation. Previous papers have looked at the relation between reputation and sale price. In addition to this relation, we look at various other testable implications, including the evolution of reputation over time.

Our evidence is mixed. We find some support for an adverse selection model where sellers can secretly change their identity. We also find support for a model where, in the spirit of Diamond (1989), a perfect record is very important and is worth investing on. Finally, our results reject implications of a pure moral hazard model, as well as the Holmström (1999) model of "career concerns."

1 Introduction

Electronic commerce presents the theoretical and the empirical economist with a number of interesting research questions. Traditional markets rely significantly on the trust created by repeated interaction and personal relationships. Electronic markets, by contrast, tend to be rather more anonymous. Can they maintain the same level of trust and efficiency?

One possible solution, exemplified by eBay auctions, is to create reputation mechanisms that allow traders to identify and monitor each other. In this paper, we look at the performance of eBay-type reputation mechanisms, both from a theoretical and from an empirical point of view. Specifically, we organize the theoretical literature that is relevant for situations like eBay; we summarize the relevant results by proposing empirical testable hypotheses; and we test these hypotheses based on data for a series of eBay markets.

If anonymity creates problems to the functioning of markets, it actually simplifies the researcher's work. On eBay, a reasonable assumption is that the entire trader's history is publicly observable both by traders and by the researcher.[1] In particular, the information that one trader has about other traders is the same as the researcher's. Essentially, this information consists of a series of positive and negative feedback comments given by past trading partners. In this context, we can make sharper predictions about agent behavior than in other markets, in particular in markets where buyers and sellers share information that is not observed by the researcher.

We may divide our study into demand and supply analysis. On the demand side, we are interested in studying how a buyer's willingness to pay is a function of a seller's reputation, which in turn is a function of that seller's history. Various authors have demonstrated that, on eBay, prices respond to reputation measures.[2] Our main contribution is to bring a clear modelling framework to the analysis and test between different alternative behavior theories.

Consider, for example, the determinants of buyer's willingness to pay. In a situation of pure adverse selection (and under additional reasonable assumptions), willingness to pay is only a function of the seller's historical percentage of negative feedback comments, which serves as a sufficient statistic for the seller's "type". If however we consider the possibility of sellers secretly changing their identities, then, even controlling for the historical percentage of negative feedback comments, seller age has a positive impact on willingness to

[1] We will explore the consequences of the failure of this assumption in Cabral, Hortaçsu and Yin (2003).

[2] References include Bajari and Hortaçsu (2000, 2003), Melnik and Alm (2002), Katkar and Reiley (2000), Resnick and Zeckhauser (2001), Houser and Wooders (2001).

pay. If instead we consider moral hazard on the seller's side, then the opposite is true: age has a negative impact on willingness to pay.

The whole purpose of reputation mechanisms is to create the appropriate incentives on the supply side. One of the main innovations of our empirical analysis (with respect to other papers on eBay) is precisely to study seller behavior. As with demand, we consider several possible theories regarding the nature of asymmetric information. Under pure adverse selection, seller behavior is trivial, but other theories have specific implications for the dynamics of seller reputation. For example, under adverse selection and unobservable effort provision, we expect the rate of negative feedback to increase over time, in particular after the first negative feedback comment is observed.

The paper is structured as follows. In Section 2, we present an overview of the theoretical literature, followed by a summary of the main models that we propose to test. We start with a model of pure adverse selection. We then consider the possibility of secret identity changes, a model similar to Tadelis (1999) and Cabral (2000). As an alternative departure from pure adverse selection, we assume unobserved seller effort. Specifically, we consider variations of the Diamond (1989) and Holmström (1999) models. Still in the context of adverse selection, we study the case when sellers can "buy" a reputation, a case that has been studied by Tadelis (1999) and Mailath and Samuelson (2001).

Section 3 describes eBay and the data we collected to conduct our analysis. We focused our attention on auctions of (arguably) ex-ante homogenous goods: mint quality collectible coins (specifically, 1/16 oz. 5 dollar gold coins of 2002 vintage—gold American Eagle, and 2001 silver proof sets); laptop computers (specifically, the IBM Thinkpad T23 PIII); and collectible toys (specifically, 1998 Holiday Teddy Beanie Babies). Our data includes information regarding sales (number of bidders, winning bid, etc), as well as the quantitative and qualitative feedback received by sellers and buyers.

Section 4, divided into several subsections, presents a series of empirical tests that follow from the theory implications developed in Section 2. We first look at the relation between reputation, age and sale price. Our results show that both the percentage of negatives, n, and age enter the price function, the latter with a positive coefficient. This is consistent with the secret identity changes model (or pure adverse selection with buyer risk aversion); and is contrary to the Holmström model of adverse selection and moral hazard. Numerically, our estimates suggest that a 1% decrease in n implies an 8% decline in return per auction. This figure seems reasonable, given anecdotal evidence and a back-of-the-envelope calculation of a seller's lifetime earnings potential. We also find that the arrival of a negative comment affects younger sellers' sale prices more than it does older sellers'.

Additional results regarding demand responses to history: We find that the significance of n seems to have increased since eBay started to publish this statistic. Moreover, when we divide histories into different periods, we find that prices respond more to a seller's overall history than her recent history. This result is at odds with predictions from a pure moral hazard model.

Our next subsection looks at the evolution of seller ratings over time. We find that the average arrival time of a second negative is lower than that of the first one. Moreover, when we look at the latter part of a seller's history, the percentage of negatives seems not to differ significantly across time periods, which suggests the stochastic arrival process of negative comments settles down. This evidence is consistent with a Diamond (1989) type model of adverse selection and moral hazard: Initially, sellers exert effort to avoid negative feedback and be perceived as "perfect" sellers; and once a negative feedback is received, separation takes places and effort drops to a lower level.

In the last subsection, we analyze sellers' propensity to "purchase" a good reputation by starting out as buyers. On eBay, users can accumulate reputation by buying and selling. Anecdotal evidence suggests that it is easier and cheaper to receive positive feedback as a buyer than as a seller. We find that many sellers in our data set (as many as 45% of them in some markets) started their eBay career as buyers, and subsequently switched to becoming sellers. The propensity of a seller to start as a buyer has a strong negative correlation with the number of total transaction that this seller has conducted. There is also a weak negative correlation with the percentage of negatives in the seller's record, suggesting that better, or at least, bigger sellers have less incentive to purchase their reputations.

2 The economic theory of reputation

The economic theory of reputation can be broadly divided into two categories: models of adverse selection, where reputation corresponds to a Bayesian posterior regarding the informed player's type; and repeated-game models, where reputation corresponds to a self-enforcing implicit contract between two parties. Seminal papers presenting these notions of reputation include Kreps et al. (1982) and Klein and Leffler (1981), respectively. One may argue that the term "trust" is more appropriate to the second meaning of reputation. The fact is that, in the literature, the term "reputation" takes at least these two very different meanings (see Cabral, 2002).

In this paper, we consider both of the above notions of reputation. The literature on each of these strands is extremely vast. Here, we limit ourselves to some of the main ideas, especially as they relate to our empirical work. The original goal of Kreps et al. (1982) and companion papers was to establish the notion of reputation in a dynamic model with asymmetric information. By repeatedly choosing actions that a certain type of player would choose, one acquires the reputation of being of that type. If such a reputation is valuable, then players will "invest" in their reputation by initially taking a suboptimal action that influences other players' belief in the desired direction. Holmström (1982/1989) and Diamond (1989) propose two extensions of the basic framework. One question that is common to these two papers is how a player's effort to invest in reputation evolves over time and as a function of the current reputation level.

As is implicit in the above discussion, reputation is an asset. This raises a number of issues, for example: can reputation be traded as any other asset? Tadelis (1999) provides conditions under which names must be traded in equilibrium. Mailath and Samuelson (2001) examine what types of players have the most to gain from acquiring a good name / reputation. Finally, Cabral (2000) determines when it's better to start a new reputation instead of extending an existing one. Although these papers do not directly model a market like eBay, their main results have testable implications for seller behavior in on-line auctions.

Let us now turn to the repeated-game notion of reputation. Klein and Leffler (1981) stressed that many market agreements are not backed up by formal contracts but rather by mutual trust and the desire to maintain one's reputation. The idea follows from the folk theorem, namely, that in repeated games and with sufficiently patient players any feasible, individually rational payoff can be attained in a Nash equilibrium.[3] Although there are many different equilibria that attain a given average payoff, typical equilibria are based on "grim" strategies: players move along a "good" phase until one of the players deviates from the agreed-upon sequence of strategies. If a deviation takes place, then players move to a "bad" phase (possibly forever). This "bad" phase may be interpreted as the guilty player's loss of reputation.

A number of extensions of the basic folk theorem and grim-strategy equilibria have been suggested, including the case when actions are imperfectly observable (e.g., Green and Porter, 1984). In particular, Dellarocas (2002) considers a model that is fairly close to the problem faced by eBay sellers.

[3]In this sense, there were a number of precursors to Klein and Leffler's (1981) results, including Friedman (1971), Telser (1980).

2.1 Application to eBay: assumptions and notation

The purpose of our paper is to test the economic theory of reputation with eBay data. Although eBay presents a fairly controlled framework, we need to make some simplifying assumptions regarding the behavior of agents before we adapt the theoretical models. Specifically, we assume that

1. A transaction has two possible outcomes: successful or unsuccessful.[4] Consumer benefit is given by $\overline{\omega}$ and $\underline{\omega}$, respectively. With no further loss of generality, for the remainder of the paper we assume $\overline{\omega} = 1$ and $\underline{\omega} = 0$.

2. A successful transaction is reported with probability π_P. It is reported as a positive with probability $\rho_P > \frac{1}{2}$. An unsuccessful transaction is reported with probability π_N and is reported as a negative with probability $\rho_N > \frac{1}{2}$. For simplicity, in the theoretical part of the paper we will assume $\pi_P = \pi_N = \rho_P = \rho_N = 1$. In the empirical portion of our paper, we will maintain the assumption that π_N and π_P are constant across sellers, and within a seller's lifetime.[5]

3. Consumers are risk neutral: expected utility is a linear function of the probability of a successful transaction.[6] Given our assumption that $\overline{\omega} = 1$ and $\underline{\omega} = 0$, willingness to pay is simply the expected probability of a positive transaction.

4. Transaction price is proportional to buyer's expected utility.

Consistent with these assumptions, we now define the following variables:

[4] A possible extension is that the outcome is continuous and the transaction successful if the outcome is above some critical value.

[5] A more general assumption is to allow π_N and π_P to depend on seller fixed-effects. We will investigate this issue in greater detail in Cabral, Hortaçsu and Yin (2003).

[6] Whenever feasible, we will also discuss the implications of risk aversion regarding our empirical hypotheses.

P	Positive transaction. Also, number of past positives.
N	Negative transaction. Also, number of past negatives.
n	negative feedback ratio: $n \equiv N/(P+N)$.
e_t	seller's effort at time t (in models with moral hazard).
θ	seller's type (in models with adverse selection); in some cases, θ is simply be the probability of a positive.
f_0, f	prior and posterior consumer beliefs regarding θ.
$\mu(P,N)$	posterior expected value of θ given history (P,N); also, willingness to pay.

2.2 Application to eBay: summary of models

Based on the literature on the economics of reputation, we consider a series of stylized models and derive testable empirical implications. The main models we will consider are the following:

■ **Pure adverse selection.** In this model, each seller is characterized by $\theta \in [0,1]$, the probability that it produces a P transaction. The value of θ is the seller's private information. Based on the seller's record, buyers update their beliefs regarding the seller's type. Let n_t be the percentage of negatives up to time t and $\mu(P,N)$ the buyers' willingness to pay based on a t history with P positives and N negatives ($t \equiv P+N, n \equiv N/(P+N)$).

Our first result is that, for very low or very high values of t, n_t is a sufficient statistic of the buyers' willingness to pay; and, moreover, the older the seller is, the more buyers care about the value of n_t:

Proposition 1 *If $t=1$ or $t=\infty$, then $n \equiv N/(P+N)$ is a sufficient statistic of $\mu(P,N)$. Moreover, $\partial\,\mu\,/\,\partial\,n|_{t=1} < 1$ and $\partial\,\mu\,/\,\partial\,n|_{t=\infty} = 1$.*

Proof: Willingness to pay for the $t+1$st unit is given by

$$\mu(P,N) = \int_0^1 x\, f_t(x)\, dx,$$

where $f_t(\theta)$ is the posterior on the value of θ:

$$f_t(\theta) = \frac{\theta^P (1-\theta)^N\, f_0(\theta)}{\int_0^1 x^P (1-x)^N\, f_0(x)\, dx}.$$

and f_0 is the prior distribution of θ. We thus have

$$\mu(P, N) = \frac{\int_0^1 x^{P+1}(1-x)^N f_0(x)\,dx}{\int_0^1 x^P(1-x)^N f_0(x)\,dx}. \tag{1}$$

If $t = 1$, there are only two possible histories (P or N); and n_t can take two values (0 or 1), so n_t is trivially a sufficient statistic. Moreover, from (1) we can see that $0 < \mu(0, 1) < \mu(1, 0) < 1$, which implies $\partial\,\mu \,/\, \partial\, n < 1$. By the Central Limit Theory, $\mu(P, N) \to n_t \to \theta$. This implies that, in the limit, n_t is a sufficient statistic and that $\partial\,\mu \,/\, \partial\, n = 1$. ■

The result suggests that, under pure adverse selection, n_t is a sufficient statistic for seller reputation. In particular, once controlling for n_t, seller age should have no explanatory power on willingness to pay. Moreover, Bayesian updating implies that the derivative of willingness to pay with respect to n_t is increasing over time.

Strictly speaking, the result only applies for $t = 1$ or the limit when $t \to 1$. By "continuity" we expect it to hold for very small or very large values of t. Exactly what "small" and "large" mean, and what happens for intermediate values of t, depends on what the particular prior distribution f_0 is. In the special case when $f_0(\theta)$ is uniform on the unit interval, it can be shown, using properties of beta distributions, that $\mu(P, N) = \frac{P+1}{N+P+2} \approx 1 - n$. In Section 4 we show numerically that n_t is a good approximation for μ for a prior f_0 that seems reasonable in the eBay context.

■ **Adverse selection: changing identities.** Suppose that the seller can change its identity without the buyers' knowledge. Specifically, suppose that each seller's life spans three periods, and that a large number of transactions take place in each period. Sellers have the option of changing their identity at the end of the first period without the buyers noticing it.[7] Let $x(\theta, n_1)$ be the seller's identity strategy, the probability of creating a new username as a function of the seller's type, θ, and the seller's performance in the first period, n_1.

Proposition 2 *The equilibrium strategy consists of a threshold function $\theta^*(n_1)$ such that the seller changes identities if and only if $\theta < \theta^*(n_1)$.*

[7] The assumption that buyers cannot identify sellers who changed their identity is crucial, though it's not a knife-edged assumption. See Tadelis (1999) and Cabral (2000).

Proof: See Proposition 2 in Cabral (2000).[8] ■

This result implies that, in equilibrium, keeping the same name is a positive signal of seller quality. In words, age should enter the willingness to pay function in addition to the historical performance (P, N). Specifically, suppose that n_t is a sufficient statistic for the posterior μ under pure adverse selection (e.g., the prior f_0 is uniform). Then, under pure adverse selection, if $n_1 = n_2$ then $\mu(n_1) = \mu(n_2)$. Under identity changes, however, $\mu(n_1) < \mu(n_2)$.

■ **Adverse selection: "buying" a reputation.** Tadelis (1999), Mailath and Samuelson (2001), and others consider the problem of buying names (and the associated reputation). Name trades do not take place at eBay (to the best of our knowledge). However, there is some anecdotal evidence that many sellers started their reputations by making a series of purchases. In fact, it is easier (and cheapter) to create a good reputation as a buyer than as a seller. In this context, the question addressed by Mailath and Samuelson (2001), "Who wants to buy a reputation?," seems to apply here as well: what seller has an incentive to start off by investing (as a buyer) on an initial reputation history? Is it low-type sellers or high-type sellers?

Consider the following stylized model. Suppose that each seller sells one unit per period and that there is a distribution of seller lives with strictly positive mass for every age.[9] All but a measure zero of sellers are "pure sellers," that is, all of their transactions are as sellers. Finally, a measure zero of sellers has the option of starting a reputation by making P_0 purchases and receiving a P in each transaction with probability one.[10]

Given a history (P, N), the (ex post) value of P_0 positives is $\mu(P_0 + P, N) - \mu(P, N)$. Let $v(\theta)$ be the ex ante expected value of $\mu(P_0 + P, N) - \mu(P, N)$. The theoretical and empirical question is then how v varies with θ.

Proposition 3 *The value of an initial reputation is declining in θ for low values of θ.*

Proof: For a seller living for t periods, the (undiscounted) value of an initial

[8] While Proposition 2 is closest to Cabral's (2000) Proposition 2, some of the basic intuitions are also present in Tadelis (1999).

[9] For this and the following models, we return to the assumption that each seller has a well-defined identity which cannot be changed.

[10] This measure-zero assumption implies that buyers take the initial record as a genuine selling record.

series P_0 of Ps is given by

$$v(\theta) = \sum_{P=0}^{t} \left(\begin{array}{c} t \\ P \end{array} \right) \theta^P (1-\theta)^{t-P} \Big(\mu(P+P_0, t-P) - \mu(P, t-P) \Big) .$$

Taking the derivative at $\theta = 0$, we get

$$\left. \frac{\partial\, v(\theta)}{\partial\, \theta} \right|_{\theta=0} = \left(\begin{array}{c} t \\ 0 \end{array} \right) t(1-\theta)^{t-1}(-1) \Big(\mu(P_0, t) - \mu(0, t) \Big) < 0,$$

since $\mu(P_0, t) > \mu(0, t)$. ■

Empirically, this result implies that sellers who choose the "buying a reputation" strategy should be worse than average. Similarly to Proposition 3, Mailath and Samuelson's (2001) Proposition 4 suggests that v is highest for the lowest seller types. See also Proposition 3 in Tadelis (1999). Note however that these models do not quite map into our framework, which we think is more appropriate in the eBay context.

■ **Adverse selection with moral hazard.** Suppose that each transaction's outcome is a function of the seller's type, θ, as well as the effort the seller makes at time t, e_t, at a cost $c(e_t)$. Suppose moreover that there is a large number of sales in each period. This situation is analogous to Holmström's (1999) model of career concerns. Even though the original application was to managerial effort, the same ideas apply in the context of eBay seller reputation.

A seller's strategy, in this context, is how much effort to put in at time t as a function of time and past performance, $e_t^*(t, n_t)$. The main result is that the seller's effort level declines with age.

Proposition 4 $\lim_{t \to \infty} e_t^* = 0$.

Proof: See Proposition 2 in Holmström's (1999). ■

In words, this implies that, at least after some t, e_t^* is decreasing. This in turn implies that, after some t, the probability of a Negative is increasing.

Proposition 4 is fairly intuitive. At initial stages of the process, buyers are still very uncertain about the seller's type. For this reason, exerting effort has a high marginal value for the seller, as it changes the buyer's beliefs. As the buyers' estimate precision increases, however, the reputational effect of effort declines and so does the level of effort.

■ **Adverse selection with moral hazard: the "perfect seller" case.** This is another variant of the adverse selection and moral hazard model. Suppose that there is a positive probability that the seller is "perfect," which in the eBay context can be defined as a seller who always produces P transactions. Then sellers will put a lot of effort into increasing and maintaining their reputation while they have a perfect record. This situation is similar to Diamond's (1989) model of reputation acquisition. Even though the original application of Diamond's model was to credit markets, the same ideas apply in the context of eBay seller reputation.

Specifically, suppose that there are two types of seller. A good seller always produces a P transaction. A bad seller produces a P transaction with probability α at an effort cost e or with probability $\beta < \alpha$ at no effort cost. Finally, let μ_0 be the buyers' prior that the seller is good.

The following result characterizes the Perfect Bayesian equilibrium of this game. This result is different from Diamond's (1989), who considers a finitely lived seller. However, the basic intuition is the same.

Proposition 5 *Suppose that* $\frac{e}{\beta e+(\alpha-\beta)(1-\beta)} < \delta < \frac{e}{\beta e+(\alpha-\beta)^2}$. *In equilibrium, if the seller's past record includes an* N*, then the seller chooses low effort. If the seller's record is perfect, then the seller chooses high effort if and only if* $t \geq t'$*, where* $1 \leq t' < \infty$.

Proof: Consider first the case when the seller's history includes an N. Bayesian updating implies $\mu = 0$, where μ is the posterior that the seller is good. The only possibility of an equilibrium where the seller chooses high effort is one where an N is punished by never believing the seller to choose high effort again. Such a punishment implies a discounted profit of $\beta/(1-\delta)$, where β is the buyer's willingness to pay for a low-quality product.

Suppose buyers expect the seller to choose high effort. The seller's expected payoff from high and low effort, assuming maximum punishment, is then given by

$$\begin{aligned} V^H &= \alpha - e + \alpha\delta V^H + (1-\alpha)\delta\beta/(1-\delta) \\ V^L &= \alpha + \beta\delta V^H + (1-\beta)\delta\beta/(1-\delta) \end{aligned}$$

Straightforward computation shows that the condition $V^L > V^H$ is equivalent to $\delta < \frac{e}{\beta e+(\alpha-\beta)^2}$. It follows that the only equilibrium following an N is low effort.

Consider now the case of a bad seller with a perfect record. Bayesian updating implies that $\mu \to 1$ as $T \to \infty$. In the limit, the seller's expected payoff from high and low effort is given by

$$\begin{aligned}\tilde{V}^H &= 1 - e + \alpha\delta V^H + (1-\alpha)\delta\beta/(1-\delta) \\ \tilde{V}^L &= 1 + \beta\delta V^H + (1-\beta)\delta\beta/(1-\delta)\end{aligned}$$

Straightforward computation shows that the condition $V^H > V^L$ is equivalent to $\delta > \frac{e}{\beta e + (\alpha-\beta)(1-\beta)}$. ■

Whether $t' = 1$ or $t' > 1$ depends on the prior μ_0. If μ_0 is close to zero, then the situation at $t = 1$ is not very different from the situation after an N, and we would expect low effort in equilibrium. If μ_0 is high, then we would expect high effort from the beginning (and until an N appears).

These results have various empirical implications which we will consider later. In particular, for bad sellers that start with a long perfect record, we would expect the time that it takes to get the first N to be longer than the time it takes to get the second N. In fact, effort is high for some periods before the first N; whereas effort is always low after the first N.

■ **Pure moral hazard with noise.** Finally, suppose that the success of a transaction at time t is a function of the seller's effort at that time, e_t, and of noise. Let $c(e_t)$ be the cost of effort. The optimal equilibrium in this situation is some mechanism whereby buyers reward sellers for good performance and punish them for bad performance. Following the seminal work by Green and Porter (1984), this problem has been examined by Abreu, Pearce and Stacchetti (1986, 1990). A specific application to eBay-like markets is presented by Dellarocas (2003).

One problem with repeated-game type of equilibria is their multiplicity. A common refinement is to impose seller optimality, that is, to restrict to the set of equilibria that give the seller the highest expected discounted payoff feasible in equilibrium. Unfortunately, this refinement still leaves us with multiple equilibria. A second possible refinement is simplicity. One possible definition of simplicity is the length of the histories required in order to sustain the equilibrium. Dellarocas (2002) has shown that, in a binomial outcome model like eBay (P or N), the simplest Pareto optimal equilibrium only requires one-period histories. More generally, we would expect average price to be only a function of the more recent performance history. Notice however that, to some extent, we are assuming the result ("if only recent histories are used then only recent histories should matter").

In the remainder of the paper we proceed as follows. In Section 3, we present the data we will use for testing these models. In Section 4, we present the empirical implications of the above models in greater details, and report our findings using data from eBay.

3 Description of the market and data

Since its launch in 1995, eBay has become one of the leading Internet sites, and the dominant online auction site. Every day, millions of items in thousands of categories are being put up for auction and bid on. The auction formats used on eBay are discussed in detail in the survey articles of Lucking-Reiley (1999) and Bajari and Hortaçsu (2002). We are going to largely ignore the intricacies of the price formation process on eBay in what follows; however, an accurate characterization of the auction environment within a given narrowly defined product category is to view it as a cross-section of temporally staggered ascending auctions (which last anywhere between 1 to 10 days), with buyers who can enter and exit an auction and place a "proxy bid" indicating their maximum willingness to pay for the item.

eBay does not deliver goods; it acts purely as an intermediary through which sellers can post auctions and buyers bid.[11] To regulate trade, eBay uses a feedback system. After an auction is completed, both the buyer and the seller can give the other party a grade of +1 (positive), 0 (neutral), or −1 (negative), along with any textual comments. There have been several changes on eBay regarding how these ratings can be given by the users. Since 1999, each grade/comment has to be linked to a particular transaction on eBay (typically, eBay stores transaction data looking back only 90 days, hence this restricts the amount of hindsight). As we will discuss later in the paper, there has also been a recent change in that eBay also records whether the grade/comment has been left by a seller or a buyer.

eBay then displays several aggregates of the grades received by each seller and buyer. These are:

1. Overall rating: this is the sum of positives minus negatives received by a seller from unique buyers throughout her entire history. Until March 1, 2003, this was the most prominently displayed feedback aggregate on eBay – it appeared next to the sellers' user ID on the auction listing page, as can be seen in the sample eBay page in Figure 2. Here, seller "wsb5" is shown to have 127 net positive ratings from unique buyers.

[11] eBay does offer an escrow service for use with especially valuable goods, though this service is used for only a small fraction of the transactions.

2. Percent of positives: As can be seen from Figure 2, eBay also reports the ratio of positives by the seller during her entire history. We should point out that this information was not reported by eBay prior to March 1st, 2003. We will exploit this temporal variation in a later section.

3. Seller's age: Since March 1st, 2003 eBay also reports when the seller registered on the site. Prior to this date, this information was also not directly available from the site.

4. Summary of most recent reviews: A mouse-click on the seller's ID on the auction listing page leads a potential bidder to a more detailed breakdown of the seller's record, as shown in Figure 3. In this page, eBay breaks down the positive, negative, and neutral ratings received by the seller in the past week, past month and past six months. In addition, this page also provides the exhaustive list of reviews left for the seller (sorted by date), giving information about the score (praise, complaint or neutral), who left the feedback, textual comments, the date when the comment was left and the transaction the review pertains to, and whether the reviewer was seller or a buyer.

Following earlier studies using data from eBay, we used Perl-based "spidering" programs to download data from eBay's website. Our sampling strategy is the following: looking at correlations of transaction prices with various measures of reputation constitutes an important part of our empirical strategy; therefore, we focused our attention on auctions of (arguably) ex-ante homogenous goods. We also wanted to capture possible sources of variation across objects with different intrinsic values, hence we collected transaction level information (the contents of Figure 2, plus the final sale price, number of bidders in the auctions) for the following objects, spanning the period between October 24, 2002 and March 16, 2003:[12]

1. Collectible coins. We chose this category since the collectible coin market is one of the most active segments on eBay and several previous studies of eBay auctions have looked at this market.[13] We looked at two different kinds of coins. The first type of coin we look at are 1/16 oz. 5 dollar gold coins of 2002 vintage (gold American Eagle), produced by the U.S.

[12]We attempted to get data from all completed auctions in this period. Several times our spider program was stalled by network problems when acquiring data. We believe that any data loss from this technical problem can be interpreted as being independent of the nature of the auction.

[13]Bajari and Hortaçsu (2000,2003), Melnik and Alm (2002) and Lucking-Reiley, Prasad and Reeves (2000).

mint. The second type of coin are 2001 silver proof sets, a set of ten coins of different denomination, also produced by the U.S. mint. An important difference between these two types of coins is that there is more room for "quality" differences between the single gold coin than the proof set.[14] There is no grading in proof sets, these are all in "mint" condition, as produced by the U.S. Mint, and are preserved in plastic container. Aside from this difference, the markets for both of these coins appear to be very similar. As displayed in the tables below, the average sale price for the gold coin in our data set was $50, and the proof sets sold on average for $78.

2. IBM Thinkpad T23 PIII notebook computers. We chose this category since, according to FBI's online fraud investigation unit, most customer complaints regarding online auction fraud come from laptop auctions. We further chose this object since notebook computers tend to come in many different configurations (regarding memory, disk space, peripherals and screen size), but this particular IBM model seemed to have relatively minor differences in configuration compared to notebooks of other manufacturers. The average sale price of the Thinkpad T23's in our dataset was $580.

3. 1998 Holiday Teddy Beanie Babies, produced by Ty toy company. Beanie babies are another hugely popular collectors' item on eBay, and according to FBI's Internet Fraud unit, comprise the second largest source of fraud complaints on online auctions. As can be seen from the summary statistics on Table 4, this is the least expensive item in our data set, with an average sale price of $10.7.

Along with transaction-level data, we also spidered each seller's "feedback summary" page, shown in Figure 2. We recorded the information contained regarding feedback aggregates from the last week, month and six months. We also recorded the entire sequence of reviews recorded for the seller, as it appears in Figure 2.

3.1 Description of the sample markets

Tables 1 through 4 report summary statistics of the "transaction level" data we collected regarding the four item classes considered above.

[14]In the data, we found that the gold coins came in three different "grades" - MS–70, MS–69 and MS–67, MS–70 being the highest.

We start with the description of the coin market, given in Tables 1 and 2. We found that of 216 gold coin auctions, 90% resulted in sale; similarly, of the 298 mint set auctions, 84% ended in a sale. The average minimum bid set by the gold coin sellers was $20, or about 40% of the average sale price; similarly, mint set sellers started their auctions at $38, or about 50% of sale price. On average 6.8 bidders participated in gold coin auctions, whereas 7.5 bidders bid for mint sets. The sellers of these coins appear to be quite experienced/large: the average coin seller had 1500 to 1600 overall feedback points. The bidders seem less experienced, with an average of 120 to 150 feedback points. This suggests that the eBay coin market is populated by "coin dealers" on the sell side, and "coin collectors" on the buy side.

We also collected data on characteristics of the auction listing, as constructed by the seller. 78% of the gold coin sellers and 66% of the mint set sellers wrote that they would accept credit cards for payment; similar proportions (54% and 60%, respectively) indicated their willingness to use PayPal, the popular online payment system favored by eBay users. 40% and 33% respectively of the gold coin and mint set auction listings contained an image of the coin, pointing perhaps to the larger degree of information asymmetry regarding the condition of the gold coin. Consistent with this, gold coin listings contained more words than mint set listings, although what we have measured is a rough count of the number of words within a listing, rather than making any inferences about the content of the listing. Lastly, the modal length of the coin auctions was 7 days, ranging from 1 day to 10 days.

Table 3 reports the summary statistics for the IBM Thinkpad market. Of the 264 auctions, 85% of them resulted in a sale, with one auction conspicuously resulting in a $1 sale (apparently due to a seller not setting his minimum bid high enough — the average minimum bid was $105). On average 21.6 bidders participated in these auctions, much higher than for coin auctions. The average seller in these auctions was quite large, with 12 445 total feedback points, although there was a seller with 0 total feedback points (and one with 25695!). Bidders were on average less experienced than coin buyers, with average overall feedback rating of 68. 80% of the sellers used PayPal and accepted credit cars, and 80% of provided an image of the computer, using on average 683 words to describe the object. These latter two numbers are consistent with the fact that the seller feels obliged to provide more information regarding a big ticket item like a laptop (as opposed to a $50 coin), but it is conceivable that since a laptop is a more complex product, it takes more words to describe it fully. The Thinkpad auctions also appear to be somewhat shorter than the coin auctions — especially the bigger sellers in this market appear to be online computer stores who use eBay as a shopfront.[15]

Table 4 provides a description of the Holiday Teddy market. Only 50% of these auctions end in a sale, and only 1.7 bidders on average attend these auctions (notice that there is a monotonic relationship between sale price of the object and number of bidders, confirming an entry cost story explored in Bajari and Hortaçsu, 2003). However, sellers tend to set their minimum bids quite high, about \$9.8 — which suggests that these sellers have good outside options for these items.[16] The average seller once again appears to be a dealer, and the average bidder a collector. 75% of the sellers declare that they will accept PayPal, however only 39% say they will accept a creditcard, most likely reflecting the transaction charges of Visa (for a \$11 item, it might not be worth paying the credit card fee). 40% of the auctions contain an image, similar to the figure for coins, and a similar number of words, 300, are used to describe the object. The average auction appears to be shorter than coin auctions, but longer than the Thinkpad auctions.

3.2 Seller characteristics across markets

We now take a more careful look at the characteristics of the sellers operating in these markets. There were 72 unique sellers in the golden coin market (translating into an average 3 auctions per seller), 157 sellers in the proof set market (2 auctions per seller), 62 sellers in the notebook market (4 auctions per seller), and 238 unique (2.4 auctions per seller) sellers in the Beanie Baby market. We should also note that one seller conducted 133 of the total 264 auctions in the notebook market — the other markets were much less concentrated. The HHI for the markets were: Thinkpads, 2756, gold coins, 342, mint/proof sets, 112, teddies, 195. This large disparity in concentration across markets can be attributed to scale effects (one of the sellers in the Thinkpad market is "ibm"), and the relatively higher importance of quality concerns in the laptop market.[17]

[15] These sellers might be interested in keeping inventory turnover high, and hence tend to list their items on shorter auctions — in fact, the correlation between auction length and seller size is -0.0931.

[16] This might correspond to alternative resale venues, or just the value from keeping these toys. Compared to coins, the value of minimun bids is rather surprising, since teddy bears require more space to store than coins, and hence one might think that inventory considerations would lead teddy sellers to want to sell their items faster.

[17] We have not yet fully investigated the dynamic implications of the reputational mechanism on the equilibrium market structure, however, it is intuitively not far fetched to think that small differences in seller performance (in terms of delivery probabilities) can be amplified a lot in the market for Thinkpads to result in a very concentrated market.

In Column (1) of Table 5 shows the breakdown of the distribution of total number of reviews (positive, neutral or negative) received by each seller in our sample, pooled over the four markets. Assuming that a constant fraction of transactions are rated by bidders (reported to be about 50% by Resnick and Zeckhauser, 2001), the total number of feedback points is a good proxy for the total number of transactions conducted by the seller, and hence a good measure of size.

We have 819 unique sellers in our sample. The average seller had 1632 total feedback responses. The median seller had 401. The largest seller has 49558 feedback responses, and the smallest had 0 (is yet to be rated, even though she was selling). As indicated in Figure 4, the distribution of sellers size (proxied by number of feedback points they got) is approximately lognormal. Consistent with the concentration measures, the mean of the size distribution is largest for Thinkpads. This is followed by teddies, gold coins and the proof sets. The dispersion of seller size is largest for Thinkpads, followed by teddy beanies, mint/proof sets and gold coins.

3.3 Frequency of bad seller behavior

We now investigate the empirical frequency of "bad" seller behavior that creates the need for a reputation mechanism to be in place on eBay. As described earlier and in previous studies of eBay, users on eBay can enter a rating (+1,0 or -1) regarding a certain buy or sell transaction on the eBay feedback forum. Some representative negative comments for sellers have the following textual content:[18]

- "THIS PERSON RIPPED ME OFF, SENT SHODDY ITEM INSTEAD OF ITEM LISTED"
- "Sold product he didn't have! Will not send refund! I am filing charges! No ansr"
- "Overgraded junk. Does not respond to emails. An irresponsible seller. Avoid him."

Although the mean and median seller in our sample is quite large (in terms of transactions conducted), they seem to have gotten very few negative comments. As can be seen from Column (2) of Table 5, the average seller in our

[18] We found that more than 40% of the positive comments contain the expression "A+". Some more colorful positive comments were: "Heaven must be missing an angel! Transaction couldn't be better! Thank U!!!" and "Mega cool mad phat deal nasty crazy cool even. Thanks."

sample has 5.3 negative feedback points, corresponding to 1% of all comments. The maximum number of negative feedback points is 651, but this seller took part in 39805 transaction. The "worst" seller in our sample, in terms of percentage of negatives and neutrals, has 6% negatives.

eBay users can also give a neutral feedback rating to signal some amount of dissatisfaction with the seller. User comments on eBay chatboards regarding the interpretation of feedback tables suggest the information contained by a neutral rating is perceived by users to be much closer to negative feedback than positive. Column (3) of table 5 shows the distribution of neutral points. Observe that the distributions of neutrals and negatives are quite close, and that the neutrals distribution first-order stochastically dominates the negative distribution. The average seller received 7.2 neutral comments in her lifetime, with a median of again 1. Given that the striking similarity between their distributions, we will lump negative and neutral comments together when talking about "negative" comments.

4 Testing the theory

We now consider the test of various economic theories of reputation based on the dataset presented in the previous section. We do so in a series of subsections that deal with different aspects of the data. In each of these subsections, we consider the theory's testable implications followed by their empirical test.

4.1 Reputation, age and price

Perhaps the most direct way to assess whether eBay's reputation measures have any effect on regulating trade on this market is to look at the variation induced by differences in seller reputation on sale prices of otherwise identical objects. Several papers (mentioned in the introduction) have run regressions of the form:

$$\text{price} = \beta(\text{reputation measure}) + \gamma(\text{other demand factors}) + \epsilon.$$

to estimate the response of prices to reputation measures reported by eBay. Our main contribution in this section is to provide interpretations of the coefficients on various reputation measures in light of the theoretical models we presented in Section 2.

Consider first the pure adverse selection model. We showed in Section 2 that, if the prior distribution $f_0(\theta)$ is uniform, then n_t is a sufficient statistic for $\mu(P, N)$, and thus for price. However, the prior is clearly not uniform. To allow for a more realistic specification for the prior, we assume that sales rates are approximately constant across types, and take the population distribution of n (quantiles of which are reported in Table 5) as the prior f_0. We then evaluate f and $\mu(P, N)$ by computing the posterior numerically.

The result of this exercise is depicted in Figure 1, which shows the relation between n, the percentage of negative feedback, and $1 - \mu(n, T)$, the expected probability of a negative. Notice that a history (P, N) can be equivalently denoted by (n, T), where T is the total number of transactions. Figure 1 depicts $1 - \mu(n, T)$ for three values of $T : 10, 100, 1000$.

Several features are suggested by the figure. First, as expected, for very short histories (e.g., $T = 10$), the posterior is relatively flat on the value of n. Second, the posterior is approximately linear on the value of n. Finally, as T becomes large, n is a good approximation for $1 - \mu(n, T)$. In other words, for long histories, we get a fair approximation by updating under the assumption of a uniform prior.

To summarize: under pure adverse selection and for large T, price is a function of n_t only. In particular, price is not a function of seller age, t. What do other models say about this prediction? Under the Holmström (1999) model, as we have seen in Section 2, the value of e_t decreases over time. This implies that, for a given posterior on θ, buyers should be willing to pay less as t increases; that is, age has a negative impact on sale price. Under the changed-identities model, for a given n_t, price should be greater the greater t is, that is, age has a positive impact on sale price.

We should also point out that the presence of risk averse buyers under the pure adverse selection model also implies a positive impact on sale price. This is due to the fact that the posterior variance of θ is declining with the number of reviews received by a seller.[19] Hence, the "risk premium" required by buyers declines with the decline in the variance of the estimate.

■ **Empirical test.** To test these hypotheses and draw an inference regarding the model governing seller behavior, we run regressions of the form:

$$\text{price} = \beta(\text{reputation measure}) + \alpha(\text{age}) + \gamma(\text{other demand factors}) + \epsilon.$$

■ **Basic price regressions** Under the null hypothesis of the pure adverse

[19]This fact is especially easy to see in the case where the prior on θ and the realization of the quality of transactions are normally distributed, though the intuition carries to more general distributions.

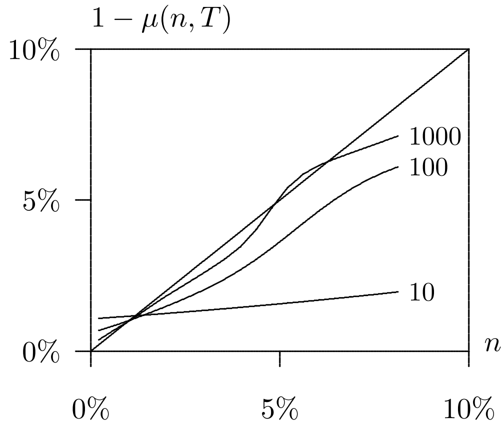

Figure 1: Relation between $n \equiv N/(N+P)$ and the posterior probability of a negative, $1-\mu$, for three values of T : $10, 100, 1000$.

selection model, n is a sufficient statistic for reputation. We measure the seller's age in two different ways. According to our model, the relevant time index is the number of (rated) transactions, since each rated transaction enables buyers to update their prior about the seller's type. We also measure age as the number of days that the seller has been active on eBay.

Table 6(a) reports our first set of results using cross-sectional regressions. In these regressions, the dependent variable is the log of the highest bid registered in the auction (according to eBay rules this is equal to the second highest bid plus the bid increment).

The regression in Column (1) allows for heteroskedasticity across object classes and controls for object fixed effects: the coefficient on n is negative and implies that a 1% increase in n leads to a 9% decline in sale price. The coefficient on the total number of reviews (divided by 1000) received by the seller, is positive (but not significant at conventional levels), and implies that 1000 additional reviews increases sale price by 5%.

Before we interpret this result as being a verification of the "changing identities" model, we should note several caveats regarding the measurement implied by the regressions. In column (2), we adjust the standard errors by allowing for correlation in the error term within a seller. This adjustment leads to the coefficient on n being no longer statistically significant (though the coefficient on total number of reviews becomes significant). In column (3), we include a dummy variable for the auctions run by "hdoutlet", the dominant seller (with close to 50% market share) in the Thinkpad market. This leads to the economic and statistical significance of both n and T to disappear, implying that the comparison of auctions of this seller vis a vis other, much smaller sellers, drives much of the finding in column (1). In column (5), we include the sellers age, measured in days (divided by 100) since her first ever feedback, as the measure of T, instead of the total number of comments. The coefficient on age is significant, implying that a seller who is 100 days older can command 1.5% higher prices.

In column (6), we include eBay's official measure of reputation (number of unique positives minus unique negatives). The coefficient estimate (which is not significant at 10%, but at 12.5%) implies that a 1000 point increase in net positives increases prices by 1.5%.

This first set of results suggest a rather weak connection between sale prices of objects on eBay, and the reputation measures that eBay publishes and those that theories of reputation suggests. The results in columns (3) through (5) suggest that variables correlated with the length of a seller's transaction history (totaltrans, age, and ebayrating) appear to have a more robust relationship with price than n, the percentage of negatives.

Before we investigate the roots of this result further, we briefly note which variables appear to affect prices in these markets the most. Prices were about 80% lower when the word "refurbished" was present in the auction description . When the seller allowed payment by a credit card, prices were higher by 28%. Longer auctions appeared to fetch higher prices (one additional day translates into 4% higher price).

■ **The impact of published statistics.** As we noted earlier, eBay made a modification in its listing format on March 1st, 2003, and began to display the percentage of positive comments that a seller received, as well as the date on which the seller registered on eBay.[20] Before this date, bidders would only see the seller's overall (net positive) feedback points next to his name. To see the fraction of seller's negative comments, the bidder would have to click on the seller's highlighted username, which would take the user to a new "feedback profile" page. There, the bidder could see a breakdown of the seller's overall net feedback into its various components (but only the counts, not the percentages).

In Table 6(b), we analyze whether this change in listing format made a difference in the response of prices to the reputation measures considered in the previous section. In columns (1) and (2), we find that the interaction of n with a dummy variable for the format change implies that the response

[20]We realized this policy change by accident, when our spidering program began to return nonsensical results.

of prices became more negative after the format change.[21] According to the regression results, the economic effect of 1% higher negatives was a positive 6% change in price before the format change (but insignificant), but a -8% return after the format change. Furthermore, the coefficient estimates on totaltrans, age, and ebayrating imply a diminished response of prices on these measures of seller reliability after the change; these coefficients are however not statistically significant.

The results of these regressions suggest that bidders respond to the way statistics are published by eBay, even though in theory bidders can compute any statistic they want using the historical feedback records for the seller. However, given that eBay can display up to 50 comment per page and some of these sellers have more than 10000 comments, this would imply a lot of page-clicking.

In order to provide some anecdotal backing to the results in Table 6(b), we quote a message posted user *elydick*, on April 23, 2003, on the eBay bulletin board on feedback policy:

```
I have one neg. feedback from someone that must of [sic] had a
really bad day. I know it is sofar down the list of over 400
positives it will never be seen. I just hate looking at the 99.8\%.
I will gladly pay the 15.00 to have it removed. I just need
the link to do it. You would think ebay would come up with a
rule that after 1 year of NO negs. the other would drop off or
something. Don't you just hate seeing it in red and on every
auction the 99.8\% because someone jumped the gun or hit the
wrong key or is just down right nasty about things?
Link please.
```

to which he got the response, from user *spiders_up_my_nose*:

```
That thing is buried, isn't it? August 2001. I have one just
like it, my only one. You can file a case with SquareTrade,
but your only hope is that your buyer responds to the email and
says yes. If not, it is too far back for SquareTrade to
continue without a response. BTW, it's $20.00.
```

It appears that there are at least two sellers on eBay who are willing to pay \$15 to \$20 to have their one negative removed, *after* eBay began to publish

[21]This regression corrects standard errors by allowing for heteroskedasticity at the seller level. We also added the dummy for "hdoutlet." Omission of either of these features lead to significance of the coefficient at higher levels.

the proportion of negatives.[22] In the case of *elydick* this would decrease his negative ratio by 0.02%. According to our point estimate in column (1), that's about a 0.16% increase in revenue per auction. When we last checked, user *elydick* sold \$200 items, so 0.16% means a 32 cent revenue increase per auction, i.e., it would take 50 or 60 auctions for him to recoup the investment (he has received a total of 652 feedbacks over the course of 3 years of activity on eBay). We infer from this that our regression based estimate of the value of reputation is not completely unreasonable.

■ **Sale probability regressions.** We also investigate whether n and the seller's age affect the probability that the object is sold, in a manner consistent with the implications for price. Table 7 reports the results of running a probit with an indicator variable for a successful sale as the independent variable. The effect of n on sale probability appears to be negative throughout. Table 7 reports the coefficients as derivatives of the probability of sale, hence our estimates imply that a 1% increase in n leads to a 1% decline in sale probability. However, the differential response of sale probabilities due to eBay's listing format change does not appear here; if anything, sale probability responds to n less than it used to.

The coefficient estimates on totaltrans, age, and ebayrating, regardless of statistical significance levels, imply very small effects on sale probability. Of additional note in these regressions (not reported) is the very strong positive correlation between the probability of sale and whether the seller allows payments to be made using credit cards. Interestingly, the use of PayPal and the presence of an image in the auction listing decrease the sale probability.

■ **Differential effects of negative comments on young vs. old sellers.** Does an additional negative hurt a young seller more than an old seller? Taking the simplest Bayesian updating model of adverse selection with uniform priors, this translates into calculating the derivative of $E(\theta|P,N)$ with respect to N, which yields $-\frac{n}{T}$, that is, given the same proportion of negatives, a younger seller will suffer more than an older seller upon receiving a negative.

To test this prediction cross-sectionally, in Table 8, we introduce N, the total number of negatives, and the interaction $N * T$ into the price regression. The prediction is that, controlling for n, the sign of the first term should be negative, but the second term should be positive.

Columns (1) and (2) of table 8 suggests that the effect of an additional

[22] "SquareTrade" (mentioned in the quotations) is a service to intermediate feedback disputes between buyers and sellers on eBay. Their fee is \$20 per complaint.

negative comment is a 1% decline in price. However, as suggested by the Bayesian updating model, the impact is smaller for large sellers: if the seller has 10000 transactions, the impact of an additional negative is about 0.4% lower (for a seller with 1000 transactions, the impact is 0.04% lower). In column (2), we also investigate the effect of the listing format change. The impact of a negative score seems to be less by 0.7% (though this is not significant). Once again, the impact of n on prices appears to be much stronger after the format change than before.

Observe also that the estimate of the impact of negative comment on price is much larger than what the uniform prior updating model predicts. Since mean T is about 1600 in our sample, the marginal effect of an additional negative predicted by the formula $\frac{\partial E(\theta|P,N)}{N} = -\frac{n}{T}$ is miniscule (about one thousandth of a percent). Our conjecture is that $\frac{\partial E(\theta|P,N)}{N}$ is quite sensitive to the specification of the prior distribution, and we will be exploring this conjecture in future work.

■ **Short vs. long memory in reputation.** In Section 4.1, we focused on the differential impact of age vs. n on sale price. Fortunately, our data set is rich enough to allow us to investigate several other measures of reputation. In particular, as we mentioned previously, eBay allows users to display a seller's feedback record for different time periods: the last six months, the last month, and the last week of activity.

A model of pure moral hazard with no adverse selection, specifically a model based on trigger-type punishment strategies in a repeated-game framework, under the assumptions of seller optimality and simplicity, predicts that only the most recent history of a seller's conduct should figure into the buyers' response, and hence into prices. Specifically, Dellarocas (2002) has shown that, in a binomial outcome model like eBay, the simplest Pareto optimal equilibrium only requires one-period histories. More generally, we would expect average price to be only a function of the more recent performance history.

To see if this prediction has any bite, we introduce the 1 month and 1 week negative ratios into our price regressions. In Table 8, Columns (3) and (4), we find that controlling for overall n, last month's negative percentage actually has a *positive* effect, and that last week's negative percentage has no effect. In column (4), we interacted these variables with the listing change, and did not finding any change. The usual change in the response to n, and the decline in the response to T due to the listing change remain. In Column (5), we look at the impact of these recent vs. overall reputation statistics on sale probabilities. We get that last month's negative percentage once again has a positive (insignificant) effect, but last week's negative percentage has a negative effect (again insignificant). The biggest response is again due to n.

This suggests that the "state variable" summarizing the seller's reputation, on which buyers appear to be conditioning their bids, depends more heavily on lifetime, or longer term performance, as opposed to recent performance. This goes against the prediction of trigger strategy models of repeated moral hazard, and is more consistent with the presence of persistent, but unknown seller "type," as in adverse selection models.

■ **Discussion of cross sectional regression results.** The main results from our cross-sectional price/sale-probability regressions can be summarized as follows:

1. Even though we tried to control for variation across objects by focusing on (relatively) homogenous goods categories, there is a lot of noise in the price regressions. Our R^2s tend to hover around 40%-50%, so there is a lot of variation left to explain. Also, adjusting standard errors to allow for correlation of the unexplained components within a seller has enormous impact on the standard errors of our regressions, a point that has not been discussed by earlier regressions.[23]

2. The length of a seller's history appears to affect prices positively, consistent with models of pure adverse selection (with risk aversion), and models where sellers can exit and reenter the population.

3. What eBay publishes as a summary statistic of reputation appears to affect the response of buyers. Across the board, we found that the positive impact of variables correlated with eBay's original rating (which depends heavily on the length of a seller's transaction history) seemed to matter more before the format change as opposed to after. We also found that the impact of n increases after eBay computes and publishes this statistic.

4. Young sellers appear to suffer more from negative comments than older sellers. However, the estimated impact of an additional negative comment on prices is extremely large, given that the sellers in our data set have very long transaction histories.

[23] Almost all of our coefficient estimates on reputation measures would have been significant at the 5% had we not taken heteroskedasticity into account.

5. Recent performance does not appear to matter more than overall performance, as suggested by grim-trigger strategy equilibria of repeated moral hazard models.

4.2 The evolution of seller ratings over time

The price regressions in the previous sections rely on cross-sectional variation across a population of sellers. However, we also have data on the past behavior of the sellers. In particular, we have access to the entire feedback history for each seller observed in the sample. Therefore we now investigate how we can exploit this interesting dimension of the data.

■ **Theory.** Under pure adverse selection, the process generating Ps and Ns is the same over time. This implies a straightforward testable implication, namely that the relative frequency of negative feedback, n_t, should be constant over time. Consider now the Holmström (1999) model. As we saw in Section 2, equilibrium effort is decreasing over time, which implies that n_t is increasing over time. Finally, consider the changing-identities model. The results presented in Section 2 imply that (a) n is higher for young sellers than for old sellers when they were young and that (b) n is increasing for old sellers.

■ **Empirical test.** We now look back at Table 5. This table displays the distribution of the ratio of negatives and neutrals to the total number of transactions, n. We calculate this ratio for all transaction conducted by the seller (Column (4)), transactions conducted in the last 6 months (Column (5)), transactions conducted in the last month (Column (6)), and transactions conducted in the last week (Column (7)).

We first look at the evolution of the mean of n. Table 6 suggests that the average percentage of negative comments is 1%, and this average appears to be stable across the entire history of the sellers, as well as the most recent history. (There are some selection effects here that we are not controlling for. The next draft will look at this much more carefully.)

We now investigate this a bit more carefully. We divide a seller's lifetime into two periods: the last six months vs. the period preceding the last six months. We then calculate n for each of these periods separately, and test their equality using a t-test. The t-test can not reject the equality of n across the two periods ($t = -0.2767, p = 0.7821$). To account for selection effects, we redo this test using a before/after window of one month. The t-test once again fails to reject the equality of means across the seller's record within the last month vs. his record preceding the last month ($t = 0.8361, p = 0.4035$).

We conducted this test also using information in the feedback tables. We compared the number of negatives received by a seller in the first T reported transactions, as opposed to the number of negatives received in the last T reported transactions, for $T = 20$ and $T = 100$. The results show that there is very little difference between the frequency of negatives between the early and late periods — which is not very surprising since 95% of the sellers did not receive a single negative comment in this window!

In summary, the results on the evolution of n seem broadly consistent with the pure adverse selection model and not with the Holmström model (career concerns) or the Tadelis-Cabral model (changing identities).

4.3 Arrival of negative feedback

As suggested by the descriptive statistics presented in Section 3, the distribution of negative feedback is very skewed: most sellers have no or very few negative comments. This suggests that, in addition to looking at the evolution of n, we also look at the timing of appearance of the first, second, etc., negative comment.

■ **Theory.** As seen in Section 2, the Diamond (1989) model of reputation acquisition implies that, conditional on a perfect record, effort is increasing, up to the highest level; and, once the first negative appears, efforts changes to the lowest level. One empirical implication is that, once the first negative appears, the probability of a second negative instantly increases. Consequently, we would expect the average arrival time of the second negative to be lower than the first one.

■ **Empirical test.** To test this hypothesis, we collected data on the timing of each negative/neutral comment received by the sellers in the data set (excluding those negative/neutral comments that were received as a "buyer" — though there were only 4 instances of this). We measured "time" in two ways: 1) in transactions, and 2) in days; and computed the time elapsed between consecutive negatives. Interestingly, we noticed that many times when an eBay seller receives a negative comment, there is a "war of words" between the seller and the buyer who places the negative. During this "war of words," the two parties can give several negatives to each other within a period of two or three days. Therefore, we did not count the negatives that the sellers received during such episodes, and concentrated on the timing between *de novo* negatives.

Table 9(a) reports the results of regressions of the time elapsed between the first two negative comments for each seller on a dummy variable that turns on for the second negative. Column (1) reports the result for the entire sample of first and second negatives, including controls for the different product markets. The base case is the time elapsed until the first negative; hence we see that the time between the first and second negative, on average, is 200 transactions shorter than the time to the first negative. In column (2), we redo the regression by controlling for seller fixed effects (equivalent to a paired t-test). The average time to the first negative is 475 transactions in the sample. The average time to the second negative after the first negative is 188 transactions shorter. In column (3), we throw out sellers with only 1 negative, and condition on having at least 2 negatives. The result is very similar.

In column (4), we restrict the sample to sellers with only 2 recorded negatives, as one plausible way of accounting for selection (though the sample is still subject to survival bias). Note that this leaves us with only 35 sellers. The statistical significance of the difference between the first and second negatives disappears, but the second negative still comes much faster than the first. In columns (5) and (6), we do the same exercise for sellers with only 3 recorded negatives, and sellers with between 2 and 5 negatives. We get very similar results – the comparison in column (6) yields a statistically significant result.

In column (6), we try to account for survival bias by conditioning our sample on sellers who were born after June 1st, 2002, and who received at least 2 negative comments. This gives us 9 sellers. We started sampling the markets in October 2002, so our hope is that these 9 sellers entered our sample without subject to a substantial survivor bias. The paired t-test once again yields that second negatives come faster than the first negative, however, the result is not significant.

In table 9(b), we repeat the same regressions in table 9(a), but measuring time between negatives in days, rather than in transactions. We find that the result holds up once again, and in a statistically significant fashion for all subsamples considered.

In table 9(c), we investigate whether the arrival of negatives accelerates, by adding in dummy variables for the third, fourth, fifth and sixth negatives. Although the arrival times between these subsequent negatives are significantly less than the time to the first negative, we do not see a very clear trend of acceleration. For example, the number of transactions between the fourth and fifth negative appears to be higher than the time between the first and second negative. When we measure time in days, as opposed to transactions, we get a similar result. Although there is some acceleration between the second and third negative, the trend does not continue to the fourth and higher negatives.

These results seem to support for the prediction of the Diamond model. However, we get much stronger support for the hypothesis if we measure "time" in days, instead of in "transactions." One might argue that in terms of a seller's effort provision decision, the variation across days might be more informative than variation across individual transactions, since the seller might be shipping objects out in discrete batches.

4.3.1 Alternative explanations

It appears from these results that there is something "special" about the very first negative that a seller receives. Once the first negative arrives, the second one arrives faster. We will now investigate three alternative explanations for this phenomenon.

The first alternative explanation is a "scale-up" effect: it might be possible that a seller takes longer to acquaint himself with the market, and does not do that much business in the early days. The fact that when we measured interarrival times in days, rather than in transactions gave us stronger results suggests that this might be what's going on.

To investigate this hypothesis, in Table 9(d), we replicate regression (1) in Table 9(c), but this time also controlling for the seller's age in days. Controlling for the age of the seller, the arrival time between the second and third negatives is still shorter, by 200 transactions, than the arrival time to the first negative. In column (2), we repeat this also controlling for seller fixed effects, and the result remain the same.

The second alternative explanation we consider is one suggested by several eBay users to whom we have presented our preliminary results. As we noted above, many of the negative comments are followed by a war-of-words between the buyer who leaves the negative, and the seller who responds to this comment. Therefore, we first checked, for every negative or neutral comment-giver in our sample, whether their particular negative comment was accompanied by a retaliatory negative left by the seller. The result was striking: of the almost 10000 negative/neutral instances in our data, 2462 resulted in a retaliatory comment by the seller. It is also interesting to note that sellers were less likely to retaliate against neutral comments, as opposed to negatives: we found that a buyer leaving a negative comment has about a 40% chance of getting hit back, while a buyer leaving a neutral comment only has a 10% chance of being retaliated upon by the seller.

Although retaliation by the seller is an important phenomenon that could compromise the integrity of the reputation mechanism installed by eBay, its presence could confound our result regarding the difference in interarrival rates of negatives only if sellers are more likely to retaliate to their first negative comment as opposed to subsequent negatives (provided, of course, that buyers recognize this possibility). We should note that it is not at all clear whether this would play out in an equilibrium setting. However, since eBay users suggested this as an alternative explanation, we turn to the data to investigate this possibility.

To do this, we regress an indicator for retaliation by the seller following a particular negative/neutral comment, on dummy variables for the second through sixth occurrence of such a comment. As displayed in table 9(d), columns (3) and (4), the dummy variables do not enter significantly – the seller is not more likely to retaliate against the first negative comment, as opposed to subsequent negatives. Interestingly, in the regression in column (3), we find that sellers with higher ex-post percentage of negatives are more likely to retaliate (the regression coefficient can be interpreted as saying that a seller with 1% higher n is 4% more likely to retaliate). However, it does not appear that "fear of retaliation" is a significant driver of the difference in interarrival times of negative comments.

The third and final alternative explanation we investigate is to see whether the buyer who places the first negative comment on a seller's record has a more "critical attitude" than buyers who place subsequent negatives. eBay sellers complain regularly about "grouchy" buyers who pepper them with negatives with no apparent reason. And to the extent that buyers might not want to tarnish a seller's "perfect" record, more "critical" buyers will be more likely to cast the first stone.

To construct an empirical test of this alternative hypothesis, we downloaded the string of feedbacks that every negative/neutral comment giver in our data set left *about other users*. We then computed the percentage of negative comments that each of these reviewers left about others, as a measure of each reviewer's "critical attitude."

In table 9(d), columns (5) and (6), we regress the critical attitude of the reviewer leaving a particular negative/neutral comment on dummy variables for the second through sixth occurrence of a negative/neutral. The regression result tells us that buyers who left the first negative were not systematically more "critical" than the buyers who left the subsequent feedbacks. Interestingly, we find that negative comment givers in the markets for coins and Beanie-Babies are much "nicer" than those in the laptop market – the percentage of negatives that laptop users left about others is about 10%, whereas for coins and Beanie Babies, this number is only about 2-3%.

Hence we conclude that scale effects, fear of retaliation, and "niceness" of buyers do not provide alternative explanations for our finding that first negatives arrive slower than subsequent negatives. There might be alternative explanations we have not taken into account, but the prediction of the Diamond model appears quite robust in our data.

4.4 "Buying" a reputation

A casual look at the feedback comment histories of some eBay sellers reveals yet another interesting pattern: some sellers appear to have started out as "buyers," completing a string of purchases before attempting their first sale. As an example of this phenomenon, in Figure 5, we plot the percentage sell vs. buy transactions that user "bearsylvania", an established Beanie Baby dealer on eBay, conducted, as a function of the number of weeks this user has been active on eBay. As can be seen, "bearsylvania" started out as a buyer first, and quickly ramped up to become a seller.

This brings about the possibility that some sellers might try to build up a reputation by being a careful buyer, and then leverage this reputation as sellers. Do sellers indeed do this? If so, what kind of sellers are most prone to do it?

■ **Theory.** In Section 2, we presented a simple adverse selection model where a seller has the option of starting with a string of P_0 positives. We showed that the ex ante value of this initial reputation is decreasing in seller's type, θ. More generally, this suggests that sellers of lower type have a greater incentive to "purchase" an initial reputation.

■ **Test.** To test these hypotheses, we utilize our data on the sellers feedback records. Starting June 26, 2001, feedback scores received by an eBay user indicated whether the feedback was in relation to a purchase or a sale. Prior to this date, users had to read through the comments to figure out where the comment came from.

Since we do not have the buyer/seller classification for feedback comments prior to June 26, 2001, we asked a research assistant to read through and classify the first 20 comments that the sellers in our data set received before June 26, 2001.[24] We then constructed a dummy variable to indicate a change

[24] We checked the accuracy of the classification by giving our research assistant comments that were classified by eBay. The classification accuracy was above 99% for "sells". We

in the seller's profile from being a buyer in his early career, to being a seller. Specifically, we looked at the first and last twenty comments received by a seller, and looked at the difference in the percentage of sale transactions across these "early" and "late" windows in their careers. We then defined a seller as having switched from being a buyer to being a seller if the difference between the percentage of sale transactions is more than 50%, conditional on the user being a seller (with 50% or higher of his transactions being sales).

We found that 38% of the sellers in the beanie baby and silver proof set categories followed the "buy first, sell later" strategy, as defined above. Sellers of gold coins followed this strategy 33% of the time, and Thinkpad sellers followed this strategy 20% of the time. The mean difference in the percentage of sale transactions across the seller's early and late lifetimes was 33%. A paired t-test of the percentage of sells in a seller's early lifetime versus the fraction of sells in the late lifetime yielded a strongly significant increase, with a standard error of 2%.

These results show that "buying first and selling later" is a widespread phenomenon on eBay, and is somewhat more prominent on some object categories than others. This latter result may indicate that forces other than bad types wanting to "buy" a reputation may play a role. For example, one might think that in the market for Beanie Babies, one may start out as a buyer first, and once one has amassed a collection, one may start thinking about making a profit by selling pieces of the collection, gradually becoming a dealer. In fact, in the market for Beanie Babies, 45% of the sample sellers could be classified as being sellers (with 50% or more sell transactions) in their first 20 transactions. On the other hand, in the market for laptops, buying is probably much more expensive, and sellers most likely obtain their inventories from other sources. Therefore, it is not surprising that 68% of the sample sellers were indeed sellers during their first 20 transactions.

Despite this caveat, one may want to see whether Proposition 3 above has any bite in the data. In table 10, column (1), we ran a regression of the "buy first sell later" dummy variable on a seller's current negative percentage, his age in days, and the total number of transactions he has completed until now. Of these variables, excluding the market fixed effects, only the total number of transactions conducted by the seller entered significantly. It appears that the total number of transactions is negatively correlated with "buying first,

had trouble classifying about a third of the "buy" comments, because the textual content is often very ambiguous. We should note that the performance of our research assistant was very similar to that of a naive Bayes text classifier algorithm, which had the drawback of not reporting ambiguous cases as being so.

selling later." The coefficient on n, the percentage of negatives, though not significant at 10%, is also negative, as predicted by the theory.

It is not entirely clear how one could interpret the previous regression. The negative coefficient on the number of transactions conducted by the seller indicates that, controlling for n, and a seller's age in days, bigger sellers are less likely to have undertaken the "buy first sell later" strategy. This might be purely a function of the way a seller does business; one might think that eBay users who are focused on selling might conduct a higher volume of transactions than those who are collector/dealers, who buy and sell at the same time, and are more likely to have started as buyers. On the other hand, one could also interpret the regression as being weakly supportive of Proposition 3, if one interprets being "big" as being good.

The fact that eBay did not explicitly report whether a given feedback was left for a buy or sell transaction previous to June 21, 2001 gives us another opportunity to check whether the "buy first, sell later" phenomenon we observe in the data is a device for sellers to build up a reputation cheaply. In column (2) of Table 10, we repeat the regression in column (1), but also add a dummy variable for whether the seller started selling on eBay before June 21, 2001. One prediction might be that the coefficient on this dummy would be positive; since post June 21,2001, one might expect more buyers to be aware of what the seller is trying to do. However, we do not find a significant coefficient on the regime change; the sign of the estimate is in fact negative.

This result is not all that surprising, given the following two reasons. First, since the sellers in our sample who started selling on eBay are a very select sample of sellers who have survived at least 2 years, one might expect these sellers to be "good" types, who would not attempt to buy a "reputation." One might expect the really bad types to have disappeared after their first attempt to milk their "bought" reputation. Second, given that the textual content of feedback comments give a lot of information about whether it is about a buy or a sell, eBay's reporting policy might not have played a significant role in shaping buyers', and hence sellers' incentives in undertaking this activity.

In conclusion, we have found very strong evidence that many sellers indeed start out as buyers, and later on become sellers. However, we have not found very convincing evidence for the theoretical prediction that worse sellers should engage in this activity. This finding, however, is subject to serious empirical limitations.

5 Conclusion

The economic theory of reputation has developed greatly over the past two decades. However, little empirical evidence has been supplied for the various models' implications. In this paper, we make a first attempt at filling this void. Our analysis is based on a fundamental assumption, namely that buyers offer feedback in a non-strategic way (specifically, according to Assumption 2 in Section 2.1). A natural next step is thus to study the various agents' feedback behavior. This we plan to do in a new empirical project (Cabral, Hortaçsu and Yin, 2003).

References

Abreu, Dilip, David Pearce and Ennio Stacchetti (1986), "Optimal Cartel Equilibrium with Imperfect Monitoring," *Journal of Economic Theory* **39**, 251–269.

Abreu, Dilip, David Pearce and Ennio Stacchetti (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* **58**, 1041–1064.

Bajari, Pat, and Ali Hortaçsu (2003), "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay," *Rand Journal of Economics*, forthcoming.

Cabral, Luís M B (2000), "Stretching Firm and Brand Reputation," *Rand Journal of Economics* **31**, 658-673.

Cabral, Luís M B (2002), "The Economics of Reputation," Lecture Notes, New York University.

Cabral, Luís M B, Ali Hortaçu, and Pai-Ling Yin (2003), "The Strategic Dynamics of eBay Buyer and Seller Feedback," Work in Progress (note: title is provisional).

Dellarocas, Chrysanthos (2003), "Efficiency and Robustness of eBay-like Online Reputation Mechanisms in Environments with Moral Hazard," MIT.

Diamond, Douglas W (1989), "Reputation Acquisition in Debt Markets," *Journal of Political Economy* **97**, 828–862.

Friedman, James (1971), "A Noncooperative Equilibrium for Supergames," *Review of Economic Studies* **28**, 1–12.

Green, Ed and Robert Porter (1984), "Noncooperative Collusion Under Inperfect Price Information," *Econometrica* **52**, 87–100.

Holmström, Bengt (1999), "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies* **66**, 169–182.

Mailath, George J, and Larry Samuelson (2001), "Who Wants a Good Reputation?," *Review of Economic Studies* **68**, 415–441.

Tadelis, S. (1999), "What's in a Name? Reputation as a Tradeable Asset," *American Economic Review* **89**, 548–563.

Telser, L G (1980), "A Theory of Self-enforcing Agreements," *Journal of Business* **53**, 27–44.

Table 1: Golden Coin summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
saleprice	194	50.76	10.74	34.03	107.5
highbid	216	50.04	13.99	.99	119
issold	216	.90	.30	0	1
numbids	216	6.81	4.63	0	22
minbid	216	19.97801	23.05935	.01	125
sellerrating	216	1596.324	1639.755	0	7501
bidderfb	160	147.2688	304.3082	0	3464
usepaypal	216	.5416667	.4994183	0	1
image	216	.412037	.493345	0	1
wordcount	96	271.2292	115.6992	12	795
creditcard	216	.787037	.4103527	0	1
auction le h	191	5.874346	2.234872	0	10

Table 2: Mint set summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
saleprice	250	77.82204	21.62466	2.01	115
highbid	298	75.78856	25.7028	1	115
issold	298	.8389262	.3682174	0	1
numbids	298	7.540268	6.930637	0	28
minbid	298	38.34238	38.16785	.01	115
sellerrating	298	1475.211	2250.246	3	9497
bidderfb	181	118.0221	181.0399	0	1235
usepaypal	298	.6107383	.488403	0	1
image	298	.3355705	.4729838	0	1
wordcount	107	241.3458	140.0917	47	845
creditcard	298	.6610738	.4741409	0	1
auction le h	269	5.542751	2.674338	0	10

Table 3: Thinkpad summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
saleprice	226	578.6746	413.5531	1	999
highbid	264	529.5183	429.897	.01	999.99
issold	264	.8560606	.351695	0	1
numbids	264	21.55303	16.46141	0	60
minbid	263	104.6864	260.6692	.01	999.99
sellerrating	264	12442.13	11628.39	0	25695
bidderfb	211	68.12796	244.782	-3	1711
usepaypal	264	.7840909	.412233	0	1
image	264	.8068182	.3955442	0	1
wordcount	220	683.8455	192.5363	33	1360
creditcard	264	.8257576	.3800383	0	1
auction le h	240	4.6375	1.840346	0	10

Table 4: 1998 Holiday Teddy Beanie Babies summary statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
saleprice	293	10.74406	3.833471	2	30
highbid	555	11.08982	4.257038	.99	30
issold	555	.5279279	.4996698	0	1
numbids	555	1.745946	2.938349	0	15
minbid	555	9.821351	5.040711	.01	30
sellerrating	555	2634.148	4371.484	0	19293
bidderfb	203	154.1823	296.4487	0	2444
usepaypal	555	.7567568	.4294278	0	1
image	555	.3837838	.486745	0	1
wordcount	345	301.0319	164.3157	52	735
creditcard	555	.390991	.4884126	0	1
auction le h	431	5.269142	2.260263	0	10

Table 5: Distribution of Feedback Aggregates Across Sellers

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	No. of total reviews	No. of Negatives	No. of Neutrals	N/(N+P) entire history	N/(N+P) last six months	N/(N+P) last month	N/(N+P) last week
Mean	1613	5.3	7.6	0.01	0.0098	0.0084	0.0082
Std. Dev.	3840	30.66	38.82	0.046	0.0171	0.027	0.0463
Min.	0	0	0	0	0	0	0
Max.	49558	651	654	1	0.125	0.273	0.556
1%	0	0	0	0	0	0	0
5%	6	0	0	0	0	0	0
10%	19	0	0	0	0	0	0
25%	98	0	0	0	0	0	0
50%	463	1	1	0.0029	0	0	0
75%	1464	3	4	0.0093	0.0084	0	0
90%	4094	9	13	0.021	0.0204	0.0217	0
95%	6997	19	25	0.0327	0.033	0.0566	0.0345
99%	15043	49	84	0.0683	0.0909	0.125	0.286
N	519	519	519	506	502	488	449

Table 6(a): Log(price) regressions

	(1) loghighbid	(2) loghighbid	(3) loghighbid	(4) loghighbid	(5) loghighbid
npn_all	**-9.051**	-9.051	-0.346	2.835	-0.400
	(3.115)*	(10.808)	(7.415)	(7.618)	(7.419)
totaltrans	0.056	**0.056**	0.004		
	(0.040)	(0.027)**	(0.003)		
age				**0.015**	
				(0.008)*	
ebayrating					0.012
					(0.009)
logminbid	0.003	0.003	0.004	0.004	0.004
	(0.000)***	(0.001)***	(0.001)***	(0.001)***	(0.001)***
image	-0.219	-0.219	-0.080	-0.084	-0.083
	(0.060)**	(0.147)	(0.107)	(0.129)	(0.107)
refurb	-0.415	-0.415	-2.259	-2.214	-2.263
	(1.135)	(1.079)	(0.736)***	(0.735)***	(0.736)***
usepaypal	0.188	0.188	-0.049	0.034	-0.047
	(0.205)	(0.200)	(0.098)	(0.120)	(0.098)
creditcard	0.365	0.365	0.293	0.281	0.293
	(0.230)	(0.104)***	(0.104)***	(0.110)**	(0.104)***
auction_length	0.039	0.039	0.042	0.038	0.042
	(0.022)	(0.019)**	(0.017)**	(0.019)**	(0.017)**
peakhour	0.242	0.242	0.185	0.223	0.187
	(0.215)	(0.168)	(0.164)	(0.182)	(0.165)
dofweek	-0.028	-0.028	-0.030	-0.031	-0.029
	(0.019)	(0.020)	(0.020)	(0.021)	(0.020)
week	-0.014	-0.014	-0.004	-0.002	-0.004
	(0.013)	(0.017)	(0.015)	(0.016)	(0.015)
eagle	0.398	0.398	0.772	0.941	0.765
	(0.070)**	(0.515)	(0.501)	(0.521)*	(0.501)
mintset	0.725	0.725	1.104	1.327	1.099
	(0.058)***	(0.510)	(0.494)**	(0.503)***	(0.494)**
teddy	-1.069	-1.069	-0.571	-0.411	-0.579
	(0.041)***	(0.525)**	(0.497)	(0.514)	(0.498)
hdoutlet			4.598	4.698	4.482
			(0.543)***	(0.539)***	(0.576)***
Constant	3.554	3.554	2.787	2.352	2.797
	(1.156)*	(1.203)***	(1.139)**	(1.198)*	(1.131)**
Observations	1114	1114	1114	1003	1114
R-squared	0.39	0.39	0.48	0.48	0.48

Robust standard errors in parentheses (Specification (1) clusters on object, other cluster sellerid)
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 6(b): Impact of Format Change

	(1) loghighbid	(2) loghighbid
npn_all	6.603	6.335
	(8.770)	(8.827)
totaltrans	0.004	
	(0.004)	
age	0.018	0.016
	(0.011)*	(0.011)
ebayrating		0.016
		(0.012)
npn_all_new	**-14.764**	**-14.075**
	(8.238)*	(8.450)*
totaltrans_new	-0.008	
	(0.010)	
age_new	-0.011	-0.010
	(0.015)	(0.015)
ebayrating_new		-0.018
		(0.018)
new_format	0.003	-0.005
	(0.379)	(0.382)
(Other auction level regressors omitted)		
Observations	1003	1003
R-squared	0.49	0.49

Robust standard errors, clustered by sellerid, in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 7: Sale Probability Regressions (probit results)

	(1) issold	(2) issold	(3) issold	(4) issold	(5) issold
npn_all	**-1.016**	**-1.218**	**-1.001**	**-1.459**	**-1.404**
	(0.811)**	(0.816)**	(0.798)**	(0.874)**	(0.870)**
totaltrans	0.001			0.001	
	(0.001)			(0.001)	
age		0.002		0.002	0.001
		(0.003)		(0.003)	(0.003)
ebayrating			0.002		0.004
			(0.002)		(0.003)
npn_all_new				0.910	0.919
				(0.679)	(0.675)
totaltrans_new				**-0.007**	
				(0.005)**	
age_new				0.003	0.004
				(0.003)	(0.003)
ebayrating_new					**-0.019**
					(0.014)**
new_format				-0.039	-0.037
				(0.059)	(0.057)

(Other auction level regressors omitted)

Observations	1004	893	1004	893	893

Robust standard errors (clustered by sellerid) in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 8: Additional Comparative Statics

	(1) loghighbid	(2) loghighbid	(3) loghighbid	(4) loghighbid	(5) issold
negatives	-0.009	**-0.011**			
	(0.006)	(0.007)*			
negatives*T	**0.00038**	**0.00044**			
	(0.00012)***	(0.00015)***			
negatives_new		0.007			
		(0.007)			
negatives*T_new		-0.00021			
		(0.00017)			
npn_mon			**7.015**	**6.817**	0.599
			(3.930)*	(4.111)*	(0.627)
npn_wk			1.650	1.195	-0.125
			(1.488)	(1.579)	(0.144)
npn_mon_new				0.872	
				(5.999)	
npn_wk_new				0.679	
				(3.984)	
npn_all	3.148	9.800	-24.473	-18.029	-0.637
	(7.521)	(9.166)	(15.746)	(17.733)	(0.688)
totaltrans	0.005	0.005	**0.048**	**0.050**	0.000
	(0.004)	(0.005)	(0.024)**	(0.025)**	(0.001)
npn_all_new		**-23.724**		**-25.054**	
		(12.868)*		(14.415)*	
totaltrans_new		0.004		-0.011	
		(0.009)		(0.009)	
(Other auction level regressors omitted)					
Observations	1114	1114	1040	1040	1040
R-squared	0.48	0.49	0.40	0.41	

Robust standard errors in parentheses

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

Table 9(a): Timing of First and Second Negatives, Time Measured in Transactions

	(1) Entire sample dneg_trans	(2) Entire sample dneg_trans	(3) >=2 negs dneg_trans	(4) ==2 negs dneg_trans	(5) ==3 negs dneg_trans	(6) [2,5] negs dneg_trans	(7) born 06/01/02 dneg_trans
d2	-202.037	-187.982	-187.982	-189.057	-225.897	-225.337	-240.444
	(45.533)***	(64.805)***	(61.511)***	(278.663)	(220.612)	(123.035)*	(162.422)
eagle	319.481						
	(68.192)***						
mint	319.200		(seller f.e., equivalently paired t-test)				
	(64.122)***						
teddy	316.237						
	(56.169)***						
Constant	208.474	475.486	467.031	609.086	814.931	642.168	372.556
	(29.673)***	(29.175)***	(30.756)***	(139.331)***	(110.306)***	(61.518)***	(81.211)***
Observations	502	502	452	70	58	202	18
R-squared	0.06	0.76	0.69	0.50	0.80	0.67	0.81

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 9(b): Timing of First and Second Negatives, Time Measured in Days

	(1) Entire sample dneg_age	(2) Entire sample dneg_age	(3) >=2 negs dneg_age	(4) ==2 negs dneg_age	(5) ==3 negs dneg_age	(6) [2,5] negs dneg_age	(7) born 06/01/02 dneg_age
d2	-331.651	-285.566	-285.566	-424.057	-520.690	-414.475	-61.000
	(29.956)***	(46.267)***	(43.915)***	(111.699)***	(117.518)***	(59.039)***	(16.957)***
eagle	48.217						
	(53.096)		(seller f.e., equivalently paired t-test)				
mint	123.450						
	(43.997)***						
teddy	39.512						
	(39.106)						
Constant	416.283	452.737	425.699	677.886	716.138	614.050	91.222
	(36.419)***	(20.829)***	(21.958)***	(78.983)***	(83.097)***	(41.747)***	(8.478)***
Observations	502	502	452	70	58	202	18
R-squared	0.19	0.68	0.57	0.48	0.60	0.55	0.87

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 9 c : Timing of Subsequent Negatives

	(1) dneg	(2) dneg_age
d2	-157.800	-181.408
	(39.150)***	(28.858)***
d3	-168.264	-209.320
	(37.746)***	(24.509)***
d4	-168.856	-206.104
	(37.009)***	(26.644)***
d5	-133.984	-206.176
	(38.203)***	(27.707)***
d6	-142.368	-211.232
	(41.767)***	(25.119)***
	(seller f.e.)	
Constant	325.520	273.512
	(29.487)***	(20.914)***
Observations	750	750
R-squared	0.49	0.39

Standard errors in parentheses

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

Table 9(d): Alternative explanations for differences in arrival times

	(1) # trans. b/w negatives	(2) # trans. b/w negatives	(3) Retaliation?	(4) Retaliation?	(5) Reviewer profile	(6) Reviewer profile
Seller's Age in Days	**0.045**	**0.440**				
	(0.019)**	(0.119)***				
2nd Negative	**-199.770**	**-200.311**	0.016	0.025	0.011	0.011
	(35.905)***	(40.973)***	(0.055)	(0.063)	(0.013)	(0.015)
3rd Negative	**-247.481**	**-237.828**	0.030	0.043	0.003	-0.003
	(36.121)***	(44.250)***	(0.059)	(0.068)	(0.015)	(0.016)
4th Negative	**-280.985**	**-269.744**	-0.005	0.000	0.020	0.020
	(37.226)***	(48.429)***	(0.064)	(0.069)	(0.020)	(0.021)
5th Negative	**-274.520**	**-252.938**	0.044	0.118	0.015	0.011
	(37.920)***	(52.847)***	(0.068)	(0.074)	(0.018)	(0.018)
6th Negative	**-337.759**	**-303.253**	0.053	0.107	**0.045**	0.040
	(43.057)***	(59.847)***	(0.071)	(0.073)	(0.023)*	(0.024)
%Negatives	**-916.901**		**4.664**		-0.053	
	(307.585)***		(1.907)**		(0.372)	
# Transactions	**0.201**		0.000		-0.000	
	(0.012)***		(0.000)		(0.000)	
eagle	29.245	(seller f.e.)	0.100	(seller f.e.)	**-0.079**	(seller f.e.)
	(19.760)		(0.120)		(0.038)**	
mint	**27.381**		0.000		**-0.087**	
	(13.746)**		(0.094)		(0.037)**	
teddy	12.066		0.091		**-0.071**	
	(9.449)		(0.089)		(0.039)*	
Constant	**251.071**	**207.052**	0.115	**0.239**	**0.105**	**0.038**
	(33.209)***	(41.952)***	(0.098)	(0.045)***	(0.043)**	(0.012)***
Observations	732	744	558	567	575	584
R-squared	0.44	0.52	0.03	0.38	0.06	0.38

Robust standard errors in parentheses
* significant at 10%; ** significant at 5%; *** significant at 1%

Table 10: Who tries to buy a good reputation?

	(1) Buy first, sell later?	(2) Buy first, sell later?
Pre-6/21/01 seller?		-0.059
		(0.090)
% negatives	-2.206	-2.164
	(1.629)	(1.631)
total transactions	**-0.015**	**-0.015**
	(0.004)***	(0.004)***
age in days	0.036	0.068
	(0.031)	(0.059)
eagle	0.113	0.112
	(0.092)	(0.092)
proof	0.161	0.163
	(0.081)**	(0.081)**
teddy	0.172	0.170
	(0.077)**	(0.077)**
Constant	0.025	-0.154
	(0.221)	(0.351)
Observations	490	490
R-squared	0.05	0.05

Standard errors in parentheses

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

Figure 2: Sample eBay Auction Listing

ebay® item view

2001 US MINT SILVER PROOF SET

Item # 3021093159

Coins:Coins: US:Proof Sets:1999-Now

Current bid	**US $35.25**	Starting bid	**US $29.95**
Quantity	**2**	# of bids	**2** Bid history
Time left	**3 days, 16 hours +**	Location	**EVANS, GEORGIA**
		Country/Region	**United States /Atlanta**
Started	Apr-26-03 11:16:46 PDT		Mail this auction to a friend
Ends	May-03-03 11:16:46 PDT		Watch this item

Description

Bid!

Seller (rating)

wsb5(127)
Feedback rating: **127** with 99.2% positive feedback reviews (Read all reviews)
Member since: Jun-19-99. Registered in United States
View seller's other items | Ask seller a question | Safe Trading Tips

High bidder **See winning bidders list** (include e-mails)

Payment PayPal, or money order/cashiers check.

Shipping Buyer pays for all shipping costs, which are provided in the Payment Details section below. Will ship to United States only.

PayPal: Fast, easy, secure payment. Learn More.

PayPal VISA ECHECK

Seller services Sell similar item

Description

Set has 10 coins. Five state quarters and the penny,nickle,dime,half dollar and golden dollar.

00012 Get Counter Stats

Payment Details		Payment Instructions
United States Shipping and handling	US $4.00	Satisfaction Guaranteed. WILL EXCEPT MONEY ORDERS,CASHIER'S CHECKS OR PAYMENT BY PAYPAL. LET ME KNOW HOW YOU WISH TO PAY. WILL SHIP SAME DAY AS PAYMENT RECEIVED. RETURNS ARE TO BE MAILED WITHIN 7 DAYS.
Additional shipping per item	US $2.00	
Shipping insurance per item (optional)	US $1.30	

Bidding

2001 US MINT SILVER PROOF SET

Item # 3021093159

Current bid:	US $35.25
Bid increment:	US $1.00
Quantity of items desired:	
Your bid per item:	

(Minimum bid: US $36.25)

Place Bid

You will confirm on the next page

This is a Dutch Auction (Multiple Item Auction) - it features multiple quantities of an item. All winning bidders pay the same price - the lowest successful bid at the end of the auction. Dutch Auctions (Multiple Item Auctions) do not use proxy bidding.

Your bid is a contract - Place a bid only if you're serious about buying the item. If you are the winning bidder, you will enter into a legally binding contract to purchase the item from the seller. Seller assumes all responsibility for listing this item. You should contact the seller to resolve any questions before bidding. Auction currency is U.S. dollars (US $) unless otherwise noted.

How to Bid

1. Register to bid - if you haven't already. It's free!
2. Learn about this seller - read feedback comments left by others.
3. Know the details - read the item description and payment & shipping terms closely.
4. If you have questions - contact the seller wsb5 *before* you bid.
5. Place your bid!

eBay purchases are covered by the Fraud Protection Program.

Need help?

Buyers: Bidding and buying tips

Sellers: Manage your listing

Figure 3: Sample Feedback Summary page

Feedback Summary

218 positives. **128** are from unique users.

0 neutrals.

1 negatives. **1** are from unique users.

See all feedback reviews for wsb5.

ID card wsb5(127)

Member since: Saturday, Jun 19, 1999 Location: United States

Summary of Most Recent Reviews

	Past 7 days	Past month	Past 6 mo.
Positive	12	51	116
Neutral	0	0	0
Negative	0	0	0
Total	**12**	**51**	**116**
Bid Retractions	0	0	0

View wsb5 's Items for Sale | ID History | Feedback About Others

Feedback Reviews for wsb5 Feedback Help | FAQ

leave feedback for wsb5 | If you are wsb5 : Respond to comments | wsb5 was the **Seller = S** wsb5 was the **Buyer = B**

Left by	Date	Item#	S/B
rattman50(11)	Apr-29-03 14:05:51 PDT	3019804072	S
Praise : Nice coin! Fast shipment!			
silverpeacedollar(26)	Apr-29-03 09:09:31 PDT	3018674118	S
Praise : hi great job nice coin and good service thanks!!!!!!			
z3forefun(351)	Apr-29-03 06:39:59 PDT	3018676358	S
Praise : very nice coin, accurately represented, fast shipping			
patrag40(161)	Apr-28-03 17:41:37 PDT	3018673349	S
Praise : The coin has been cleaned but a great deal			
bernardtreeman(62)	Apr-25-03 18:11:09 PDT	3014810862	S
Praise : thanks for a nice coin. ++++++++++++AAAAAAA			
kucak1(114)	Apr-25-03 06:07:31 PDT	3013485158	S
Praise : HIGHLY RECOMMEND THIS GENTLEMAN!!! Thanks, Willard!!!			
rdt9819 (73)	Apr-24-03 14:37:12 PDT	3018676926	S
Praise : GOOD TRANSACTION WOULD BUY AGAIN A+++++++			
bfjfkman (24)	Apr-23-03 15:03:21 PDT	3018675234	S
Praise : Fast Delivery, Good Packaging, Great Deal. (Very Nice Coins, too.)			
bfjfkman (24)	Apr-23-03 15:02:00 PDT	3018677589	S

Figure 4: Distribution of Seller Sizes

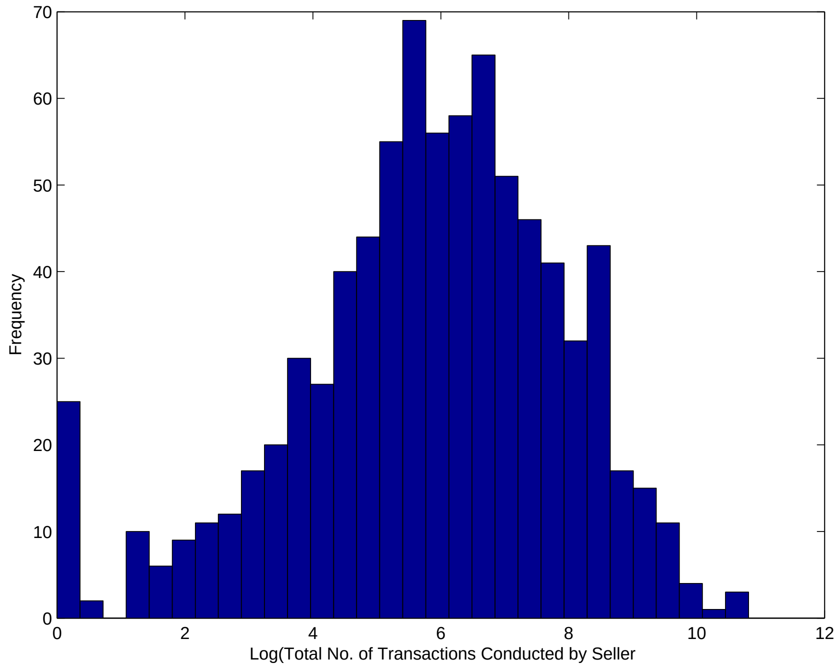

Figure 5: Becoming a seller

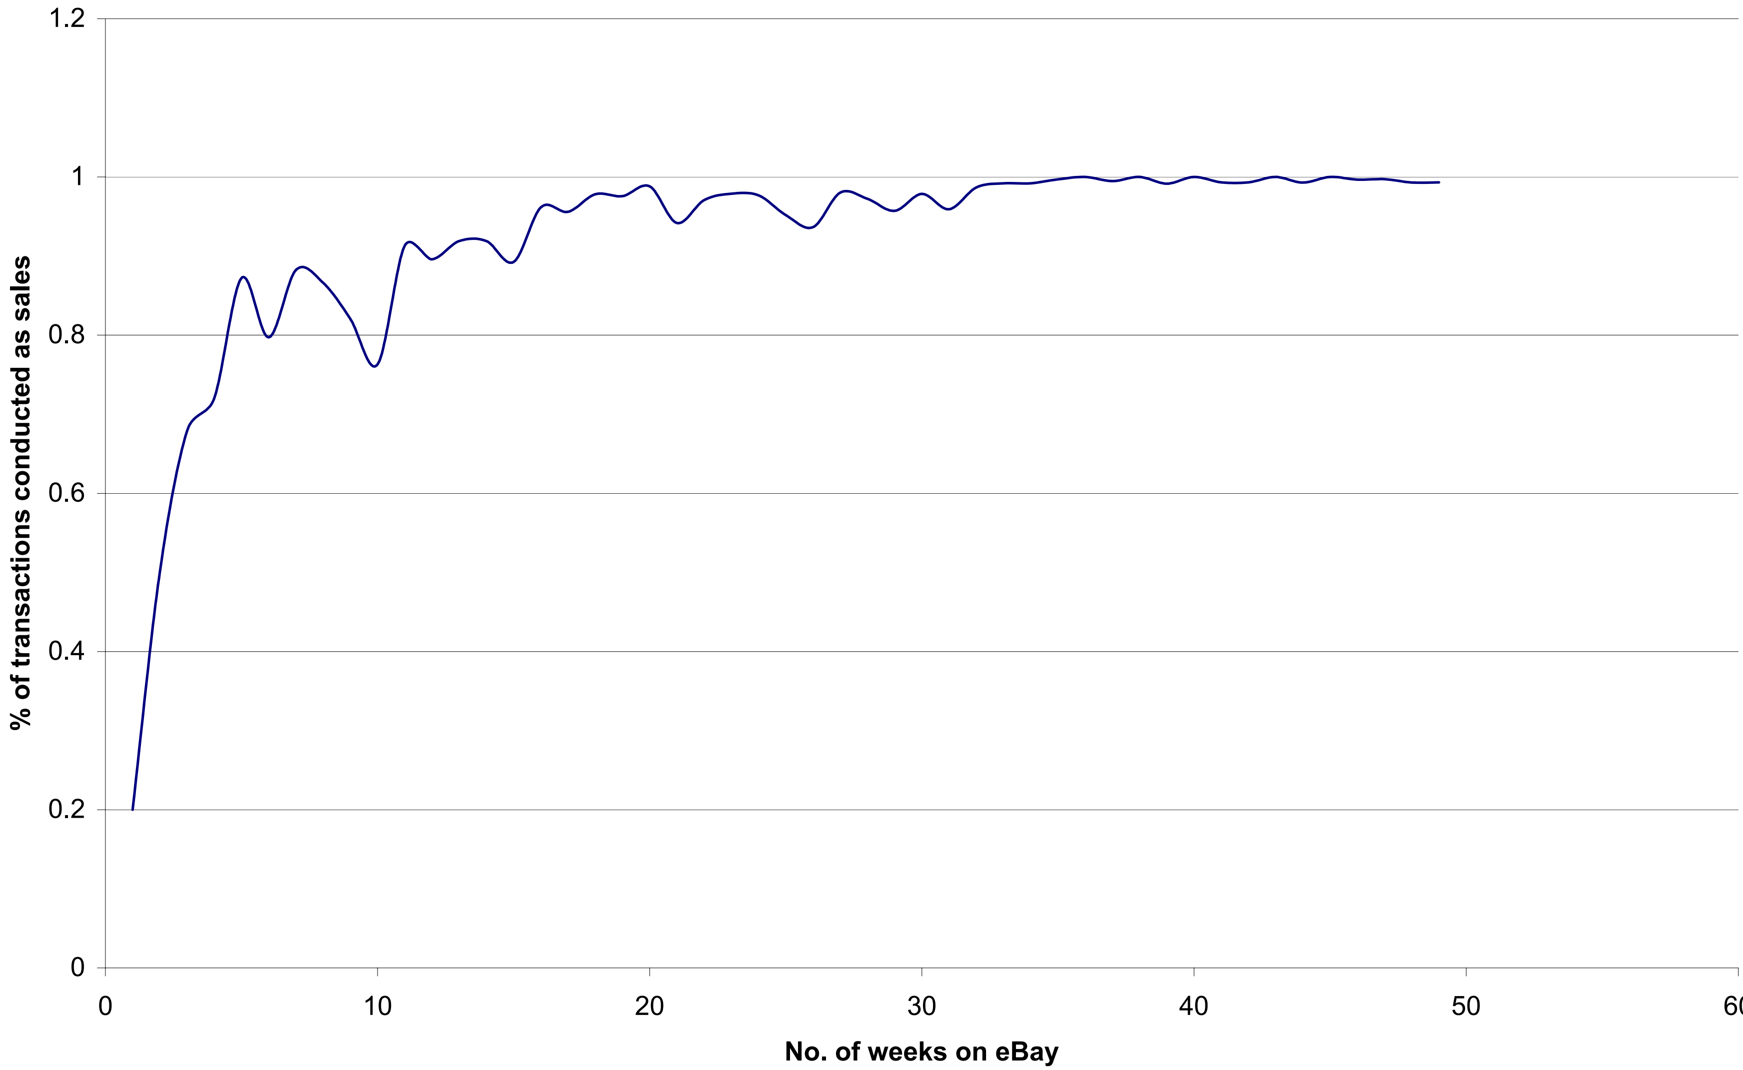

Figure 6. Pictures of auctioned objects in this study.

(a) 1/16 oz 5 dollar gold coin of 2002 vintage (gold American Eagle)

(b) 2001 silver proof set.

(c) IBM Thinkpad T23 PIII

(d) 1198 Holiday Teddy Beanie Baby