DISCUSSION PAPER SERIES

No. 4345

THE DYNAMICS OF SELLER REPUTATION: THEORY AND EVIDENCE FROM eBAY

Luís M B Cabral and Ali Hortaçsu

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP4345.asp**

ISSN 0265-8003

THE DYNAMICS OF SELLER REPUTATION: THEORY AND EVIDENCE FROM EBAY

Luís M B Cabral, New York University and CEPR
Ali Hortaçsu, University of Chicago

Discussion Paper No. 4345
April 2004

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 4345

April 2004

ABSTRACT

The Dynamics of Seller Reputation: Theory and Evidence from eBay*

We propose a basic theoretical model of eBay's reputation mechanism, derive a series of implications and empirically test their validity. Our theoretical model features both adverse selection and moral hazard. We show that when a seller receives a negative rating for the first time his reputation decreases and so does his effort level. This implies a decline in sales and sale price, and an increase in the rate of arrival of subsequent negative feedback. Our model also suggests that sellers with worse records are more likely to exit (and possibly re-enter under a new identity), whereas better sellers have more to gain from 'buying a reputation' by building up a record of favourable feedback through purchases rather than sales. Our empirical evidence, based on a panel dataset of seller feedback histories and cross-sectional data on transaction prices collected from eBay, is broadly consistent will all of these predictions. An important conclusion of our results is that eBay's reputation system gives way to strategic responses from both buyers and sellers.

JEL Classification: D44, L15 and L86
Keywords: auctions, eBay, moral hazard, quality and reputation

Luís M B Cabral
Stern School of Business
New York University
44 West 4th Street
New York, NY 10012-1126
USA
Tel: (1 212) 998 0858
Fax: (1 212) 995 4218
Email: lcabral@stern.nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=114053

Ali Hortaçsu
Dept of Economics
University of Chicago
1126 E. 59th Street
Chicago, IL 60637
USA
Tel: (1 773) 702 5841
Email: hortacsu@uchicago.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=159294

*We thank Kenny Ballendir, Tim Miller and Jeremy Shapiro for truly outstanding research assistance. We also thank Damien De Walque, Svetlozar Nestorov, Mike Riordan, Anne Rogers, Steve Tadelis, and seminar participants at Chicago, Essex, Copenhagen, SED meetings, NBER Summer Institute, Georgetown, Tufts, Rochester, NYU and Columbia for helpful comments and suggestions. Hortaçsu acknowledges financial support from the National Science Foundation (Grant SES-0242031). The usual disclaimer applies.

Submitted 02 March 2004

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 the same level of trust and efficiency be obtained in these markets?

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 study eBay-type reputation mechanisms, both from a theoretical and from an empirical point of view. Specifically, we propose a basic theoretical model of eBay's reputation mechanism, derive a series of implications and empirically test their validity.

Our focus on eBay's reputation mechanism is justified for two reasons. First, electronic commerce in general and eBay in particular are a significant economic phenomenon: in 2003, more than \$21bn were transacted on eBay by 69 million users. Second, with its well defined rules and available information, eBay presents the researcher with a fairly controlled environment for theory testing. Specifically, a reasonable assumption on eBay is that the information 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.

Our theoretical model features both adverse selection and moral hazard on the seller's side. In the spirit of Diamond (1989), we show that in equilibrium there is a positive correlation between seller reputation and seller effort. Specifically, when a seller receives a negative rating for the first time, his reputation decreases and so does his effort level. This implies a decline in sales and sale price; and, moreover, an increase in the rate of arrival of subsequent negative feedback. Our empirical evidence is broadly consistent with these predictions. Specifically, we find that the growth rate of a seller's transactions drops from about 7% per week to about -7% following the first negative feedback. We also find that the rate of negative feedback arrival increases twofold following this event. Both findings are strongly statistically significant across a variety of empirical specifications. We also find that the sale price for identical goods varies across sellers with differing feedback records: a 1% level increase in the fraction of negative feedback is correlated with a 9% decrease in price

A natural experiment based on a change in eBay's reporting format suggests there is indeed a causal relation between seller reputation and sale price.

We consider two extensions of our basic model. First we allow for the possibility of seller "exit", which we assume corresponds to a secret change in identity. We show theoretically that exit is more likely the worse the seller's record is. Our empirical findings are once again consistent with this result. We find that a tenfold increase in a seller's transaction record length is correlated with a 18 to 27% lower probability of exit within the observation period. Moreover, a 1% level increase in the fraction of negative feedback is correlated with a 1 to 2% increase in probability of exit (however, this coefficient is not statistically significant).

Second, we consider the possibility of sellers building up a record ("buying a reputation") by starting off as buyers and then switching to selling (anecdotal evidence suggests that it is easier and cheaper to accumulate positive feedback as a buyer than as a seller). Our theoretical model suggests that better sellers have more to gain from building such a record, a prediction that is borne out by the data. Specifically, we define a seller as a "switcher" if more than 50% of the first 20 transactions were purchases whereas more than 70% of the last 20 transactions were sales. About 30% of all sellers fall in this category; sellers with 1% lower percentage of negative feedback are 6% more likely to have started out "switchers".

A number of authors have conducted empirical studies of eBay's reputation mechanism. Almost all of these prior studies focus on the buyer response to published feedback aggregates. In particular, a large number of studies estimate cross-sectional regressions of sale prices on seller feedback characteristics: Dewan and Hsu (2001), Eaton (2002), Ederington and Dewally (2003), Houser and Wooders (2003), Kalyanam and McIntyre (2003), Livingston (2002), Lucking-Reiley, Bryan, Prasad and Reeves (2000), McDonald and Slawson (2002), Melnik and Alm (2002), Resnick and Zeckhauser (2001).[1] Resnick, Zeckhauser, Swanson and Lockwood (2003) point out the potential for a significant omitted variable bias in these cross-sectional regressions, and conduct a controlled field experiment in which a seasoned seller sells identical postcards using his real name and an assumed name. They find an 8% premium to having 2000 positive feedbacks and 1 negative over a feedback profile with 10 positive comments and no negatives. Ba and Pavlou (2002) conduct a laboratory experiment in which subjects are asked to declare their valuations for experimenter generated profiles, and find a positive

[1]See Dellarocas (2002), Resnick, Zeckhauser, Swanson and Lockwood (2003), and Bajari and Hortaçsu (2004) for surveys of these results.

response to better profiles. Jin and Kato (2004) assess whether the reputation mechanism is able to combat fraud by purchasing ungraded baseball cards with seller-reported grades, and having them evaluted by the official grading agency. They report that while having a better seller reputation is a positive indicator of honesty, reputation premia or discounts in the market do not fully compensate for expected losses due to seller dishonesty.

Our main contribution to the study of online reputation mechanisms is to devise a number of theory-driven empirical tests to investigate the incentives created by eBay's feedback system. Our focus is on the empirical implications of sellers' equilibrium behavior. By contrast, with the exception of Jin and Kato (2004), previous work has studied buyers' reaction to seller's feedback record. Moreover, our empirical tests are primarily based on panel data, whereas most of the previous work is primarily based on cross-section data. Using panel data allows us to account for seller-level heterogeneity in most of our empirical tests.[2]

In addition to the literature on eBay and its reputation mechanism, our paper also relates to the empirical study of models with adverse selection and moral hazard. In particular, one of the most striking and robust results in our paper is that, once a seller receives negative feedback from buyers, the frequency of such feedback increases dramatically. We show that this is consistent the presence of moral hazard and rejects a pure adverse selection model. Abbring, Chiappori, and Pinquet (2003) suggest a related test for the presence of moral hazard in auto insurance by looking at interarrival times of reported accidents in panel data on claims histories. Their test exploits discontinuous changes in driver incentives created by an exogenously specified experience rating scheme determining insurance premia. They fail to find evidence for moral hazard in a sample of French drivers. In our setting, seller incentives are created endogenously through buyers' expectations of what the seller will do in the future, and hence a discontinuity in incentives in response to an "accident" (i.e. a negative comment) is more difficult to establish. Nevertheless, we succeed in deriving a robust empirical implication that is strongly verified

[2]We believe the difference between panel and cross-section data is important. In fact, our results from panel data are typically very significant, whereas our results from cross-section data, consistently with much of the previous literature, have weak statistical significance.

in the data.[3]

The paper is structured as follows. In Section 2, we briefly describe the institutional setup of eBay, in particular the mechanics of its reputation mechanism. In Section 3, we present our basic model of buyer and seller behavior, as well as a number of extensions. Section 5 tests the implications from our basic model regarding sales rate (Section 5.1), price (Section 5.2), frequency of negative feedback arrival (Section 5.3), exit (Section 5.4), and reputation building (Section 5.5). Section 6 concludes the paper.

2 The eBay reputation mechanism

Since its launch in 1995, eBay has become the dominant online auction site, with millions of items changing hands every day. We will not attempt a detailed account of how eBay has evolved and what its trading rules are; the interested reader may find this in a number of survey articles and in the popular press.[4] Thus we are going to largely ignore the intricacies of the price formation process on eBay in what follows; however, from our modelling purposes it will not be too inaccurate to characterize the auction mechanism as a variant of the second-price auction.[5]

eBay does not deliver goods: it acts purely as an intermediary through which sellers can post auctions and buyers bid. eBay obtains its revenue from seller fees collected upon successfully completed auctions.[6] Most importantly, to enable reputation mechanisms to regulate trade, eBay uses an innovative feedback system.[7] 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),

[3] Abbring, Chiappori, and Pinquet (2003) show that, in the French auto insurance market, an accident increases the cost of future accidents. An implication of moral hazard is that the arrival rate of accidents decreases when an accident takes place. By contrast, our model predicts that the marginal benefit of effort decreases when an "accident" (negative feedback) happens. Therefore, the arrival rate of "accidents" should go up when the first "accident" happens.

[4] See Cohen (2002) for an entertaining historical account of eBay. Survey articles on Internet auctions include Lucking-Reiley (1999), Dellarocas (2003), and Bajari and Hortaçsu (2004).

[5] In reality eBay auctions are dynamic auctions in which bidders place (possibly multiple) "proxy bids" indicating their maximum willingness-to-pay. See Roth and Ockenfels (2002), Ockenfels and Roth (2003), and Bajari and Hortaçsu (2003) for detailed analyzes of dynamic bidding behavior on eBay.

[6] Success is defined as a bid above the minimum bid or a secret reserve price set by the seller. eBay collects its fee even if the physical transaction does not take place.

[7] 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.

along with any textual comments.[8]

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 1. (Here, seller *wsb5* is shown to have 127 net positive ratings from unique buyers.)

2. Percent of positives: As can be seen from Figure 1, eBay also reports the ratio of positives received 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 Section 5.2.

3. Seller's age: Since March 1st, 2003 eBay also reports the date when the seller registered on the site. Prior to March 1st, 2003, this information was 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 2. 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.

5. The entire feedback record: 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 a seller or a buyer. (The latter piece of information was not available prior to June 16, 2001, a fact that we will take into account.)

As can be seen, eBay provides a complete record of the comments received by each seller, along with various summary statistics, and this information is

[8] 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 (in particular price) data looking back only 90 days, hence this restricts the extent of "historical research" that a buyer can conduct.

publicly available. Hence, as claimed in the introduction, this is an environment where the economic analyst has the same information that a new buyer has about a seller.[9] We will thus take this informational equivalence as given when formulating our theoretical model and its empirical implications.

3 Basic model and empirical predictions

Although, as described in the previous section, eBay presents the economic modeler with a fairly structured and controlled framework that is harder to replicate in other real-world empirical settings, we need to make some simplifying assumptions regarding the behavior of agents before developing a theoretical model.

Assumption 1 *A transaction has two possible outcomes: successful or unsuccessful, with consumer benefit equal to 1 and 0, respectively.*

More generally, we could assume that consumer benefit is given by $\overline{\omega}$ and $\underline{\omega}$, respectively. However, for the remainder of the paper and with no loss of generality we assume $\overline{\omega} = 1$ and $\underline{\omega} = 0$. Another possible extension is that the outcome is continuous and the transaction considered successful if the outcome is above some critical value; see Section 5.

Assumption 2 *A successful transaction is reported with probability one as a successful transaction. An unsuccessful transaction is reported with probability one as an unsuccessful transaction.*

All of the results in the paper can be extended to the case when there is a small probability of error in feedback or a less than 100% feedback rate. However, the analysis becomes substantially more complicated. A more crucial assumption we need is that the probability and accuracy the feedback be independent of the seller and of the seller's history, an assumption that we will attempt to test empirically in section 5.3. Following eBay's terminology, we will refer to a successful transaction as a "positive," or simply P; and an unsuccessful transaction as "negative," or simply N.

Assumption 3 *Buyers are risk neutral.*

[9] Of course, "old" buyers may know about private transactions that they did not comment on.

Given Assumption 1, Assumption 3 implies that willingness to pay is simply the expected probability of a P transaction.[10]

Our basic model combines elements of adverse selection and moral hazard. It is similar in structure to Diamond's (1989) model of reputation acquisition in credit markets.[11] Although the context in which we apply it is quite different, the basic mechanism is the same. In his model, the informed party is a firm who knows its type and must choose effort level. The uninformed parties are lenders, who must determine the interest rate. In our formulation, the informed party is a seller who knows her type and must choose effort level. The uninformed parties are the buyers, who must determine whether to bid and, if so, how much to bid.

Specifically, we assume that each seller can be of two types. A good seller always produces P transactions.[12] A bad seller produces a P transaction with probability $\alpha < 1$ at an effort cost e or with probability $\beta < \alpha$ at no effort cost. Let μ_0 be the buyers' prior belief that the seller is good. Each seller lives for an infinite number of periods and discounts the future according to the discount factor δ. In each period, the seller auctions one unit with a second price auction with no secret reserve price or minimum bid.[13]

On the buyer's side, we assume that, in each period, there are B potential identical bidders who live for one period. Each bidder has a valuation given by

$$v(\mu, \rho) = \mu + (1 - \mu)\left(\rho\alpha + (1 - \rho)\beta\right),$$

where μ is the posterior belief that the seller is good and ρ is the belief that the seller, being bad, will make an effort to improve transaction quality. Basically, $v(\mu, \rho)$ is the buyers' expected probability of a P transaction: with probability μ, the seller is good, in which case P happens with probability one; with probability $1 - \mu$, the seller is bad, in which case the outcome is P with probability α or β, depending on whether the seller exerts effort (probability ρ) or not (probability $1 - \rho$).

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

[11] Diamond's model, in turn, builds on the earlier work of Kreps, Milgrom, Roberts and Wilson (1982), Kreps and Wilson (1982), Milgrom and Roberts (1982). See also Hölmstrom (1999) for a related model featuring similar dynamics.

[12] Later in the paper we consider the case when a good seller produces a P with probability less than one.

[13] A straightforward extension is to assume the seller puts an object up for auction at an exogenously given rate, independent of its type. At the end of the section we consider an extension where the decision to auction an object is endogenously determined.

Each bidder must pay a cost c in order to enter an auction.[14] Each time an object is put up for auction, all B potential bidders simultaneously decide wether to enter the auction, paying a cost c if they decide to enter and bid. We assume bidders play the unique symmetric entry equilibrium. In this equilibrium, each bidder enters the auction with probability p, where p is determined by the indifference condition between entering and not entering the auction. Finally, the bidders that decided to enter simultaneously set their bids and payoffs are paid.

There are three relevant possible outcomes of the bidder entry game. If two or more bidders enter, then the seller gets v and each bidder gets zero. If one or zero bidders enter, then the seller gets zero and the bidder (if there is one) gets v.[15] The entry probability p is thus given by the indifference condition $(1-p)^{B-1}\,v = c$, or simply

$$p(v) = 1 - \sqrt[B-1]{\frac{c}{v}}\,, \tag{1}$$

Note that p is increasing in v. The seller's expected payoff is given by

$$\pi(v) = \Big(1 - B p(v)\,\big(1-p(v)\big)^{B-1} - \big(1-p(v)\big)^{B}\Big)\,v. \tag{2}$$

(The expression in parentheses on the right-hand side is the probability that there is more than one bidder, the only case when the seller makes a profit.) Both $p(v)$ and $\pi(v)$ are increasing in v. Finally, since v is increasing in μ, it also follows that p and π are increasing in μ.

We now turn to the characterization of the seller's equilibrium strategy. We do so in the context of the following important assumption, which we will maintain throughout:

Assumption 4 $\frac{e}{\beta e + (\alpha-\beta)\big(\pi(1)-\pi(\beta)\big)} < \delta < \frac{e}{\beta e + (\alpha-\beta)\big(\pi(\alpha)-\pi(\beta)\big)}$.

In words, we assume that the value of the discount factor, δ, is intermediate. A very high value of the δ implies that there is a multiplicity of equilibria. In fact, for δ sufficiently close to one any feasible, individually rational payoff profile is attainable as a perfect Bayesian equilibrium of the game. A very low value of δ, in turn, implies that there is only one equilibrium, one where the (bad) seller never exerts effort.

[14] See Levin and Smith (1994), Bajari and Hortaçsu (2003).

[15] We implicitly assume that the seller's object is perishable. A possible extension is to assume that an unsold object has value v_U to the seller.

The following result characterizes a 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, namely, the idea that reputation and effort are "correlated" in equilibrium.

Proposition 1 *In a perfect Bayesian equilibrium,*

1. *After the first N, the buyers' willingness to pay decreases.*
2. *After the first N, the seller chooses low effort.*
3. *There exists a t' such that the seller chooses high effort if he has a perfect record longer than t'.*

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 will choose high effort again, $\rho = 0$. Such a punishment implies a discounted profit of $\pi(\beta)/(1-\delta)$, where β is the buyer's willingness to pay a bad seller who does not exert effort.

If instead buyers expect the seller to choose high effort, that is $\rho = 1$, then the seller's expected payoff from high and low effort, assuming maximal punishment, is given by

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

Straightforward computation shows that the condition $V^L > V^H$ is equivalent to $\delta < \frac{e}{\beta e + (\alpha-\beta)\big(\pi(\alpha)-\pi(\beta)\big)}$. 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 $\mu \to 1$, and thus $v \to 1$, as $T \to \infty$. In the limit, the seller's expected payoff from high and low effort is given by

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

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

The above calculations imply that ρ declines (weakly) as the first N appears. Moreover, Bayesian updating implies that μ drops from a positive value to zero. We thus conclude that v decreases as the first N is given. ■

Notice that Proposition 1 says nothing about uniqueness of equilibrium. What it does state is that, for the particular interval of values of δ, *any* Perfect Bayesian equilibrium calls for the seller to choose low effort following the arrival of an N. If μ_0 is close to one (and for the same interval of values of δ), then there exists a unique perfect Bayesian equilibrium: high effort until the first negative arrives, low effort thereafter. For lower values of μ_0, multiple equilibria are possible, some with initial high effort, some with initial low effort. However, *any* perfect Bayesian equilibrium has the property that, if the string of initial Ps is long enough, then the seller chooses high effort.

We should note that Proposition 1 is not a knife-edged result: following the steps of the proof, one can see that continuity arguments apply if we assume that a good type produces a P with probability γ lower than, but close to, one. In fact, below we consider an extension of the basic model where γ is strictly less than one.

Having said that, we should restate that the result depends crucially on the particular values of δ we consider. If δ is very high, then the folk theorem applies: any equilibrium path that is feasible and individually rational is the result of a perfect Bayesian equilibrium for a high enough value of the discount factor δ. In other words, if the discount factor is high enough, then equilibrium theory has no predictive power. At the other extreme, if δ is very low then there is a unique equilibrium where the seller chooses low effort in every period. Points 1 and 2 in Proposition 1 still hold true, but not Point 3.

The results above have various empirical implications which we now consider.

Corollary 1 *Let $R(P, N)$ be the rate of transactions per period for a seller with history (P, N). Then*

$$R(P, 0) > R(P, 1) = R(P, i) \ \ \forall i > 1.$$

Proof: From (2), we see that $p(v)$ is increasing in v. The result then follows from the first part of Proposition 1. ■

Corollary 2 *Let $\Pi(P, N)$ be the average sale price for a seller with history (P, N). Then*

$$\Pi(P, 0) > \Pi(P, 1) = \Pi(P, i) \ \ \forall i > 1.$$

Proof: From (2), we see that $\pi(v)$ is increasing in v. The result then follows from the first part of Proposition 1. ■

In words, Corollaries 1 and 2 state that, as the seller receives his first negative, both the sales rate and price go down. This is because the buyers' willingness to pay, v, and the buyers' entry probability, $p(v)$, are increasing in v. Although we are considering a particular auction mechanism, the above results are valid for more general settings. The important feature is that both $p(v)$ and $\pi(v)$ be increasing.

Note that Corollaries 1 and 2 follow from our model with adverse selection and moral hazard, they are also consistent with a pure adverse selection model, the case when $\alpha = \beta$. Our next result, however, is specific to the case of moral hazard:

Corollary 3 *Let $T(P, N)$ be the expected number of transactions between the Nth and the $N+1$st Negative. Then*

$$T(P, 0) > T(P, 1) = T(P, i) \;\; \forall i > 1.$$

Proof: In equilibrium, the seller chooses low effort after the first N and a new N arrives at the rate $1 - \beta$. Moreover, for any μ_0 the seller chooses high effort after a sufficiently long string of initial Ps, a positive probability event. It follows that expected average effort during the perfect record phase is greater than minimal effort, whereas average effort after the first N is minimal. ■

In words, Corollary 3 implies that, once an N appears, additional Ns will appear at a higher rate. Notice that, for low μ_0, there may be equilibria realizations such that the seller chooses low effort in every period. However, there are positive probability equilibrium realizations such that effort goes down after the first negative. In expected terms, therefore, the arrival rate of Ns goes up.

■ **Extension: endogenous entry.** We now extend the basic model by making endogenous the seller's decision of whether to auction a given object.[16] Suppose, as before, that there is an exogenous stream of one object per period that the seller has available. At the beginning of the period, the seller learns the cost of auctioning the object, a. We assume a is i.i.d. across periods with c.d.f. $F(a)$. This cost includes the monetary cost that sellers must pay eBay in addition to the opportunity value of keeping the object (for future sale at

[16] See Bar-Isaac (2002) for a related result on the endogenous selling decision.

eBay or outside of eBay, or for personal use). The seller must then determine whether or not to auction the object.

Let $a^*(P, N)$ be the critical value of a below which the seller will auction the object. Let $A(P, N)$ be the probability (before learning a) that the seller decides to auction an object. We thus have $A(P, N) = F\left(a^*(P, N)\right)$. Our first result relates to the evolution of $A(P, N)$. We show that the higher the seller's reputation, the more he has to gain from putting an object up for auction, and thus the more often he will do so.[17]

Proposition 2 *Let $A(P, N)$ be the rate at which the seller with history (P, N) places objects for auction. If α is sufficiently large, then $A(P, 0) > A(P, 1) = A(P, i)$, for all $i > 1$.*

Proof: Before the first negative is received, the value of a^* solves

$$\pi(P, 0) - a^*(P, 0) + \delta\alpha V(P+1, 0) + \delta(1-\alpha)V(P, 1) = \delta V(P, 0).$$

If $\alpha \approx 1$, then we get

$$a^*(P, 0) \approx \pi(P, 0) + \delta\left(V(P+1, 0) - V(P, 0)\right). \tag{3}$$

After the first negative is received, the value of a^* solves

$$\pi(P, 1) - a^*(P, 1) + \delta\alpha V(P+1, 1) + \delta(1-\alpha)V(P, 2) = \delta V(P, 1).$$

Since $V(P, 1) = V(P, 2)$, it follows that

$$a^*(P, 1) = \pi(P, 1). \tag{4}$$

Comparing (3) and (4), and noting that $\pi(P, 0) > \pi(P, 1)$ and $V(P+1, 0) > V(P, 0)$, we conclude that $a^*(P, 0) > a^*(P, 1)$. ■

■ **Extension: endogenous exit.** We now consider the possibility of a seller "exiting," by which we mean secretly changing his identity and starting a new reputation history.[18] Intuitively, we would expect the seller's tendency

[17] In the following result, we make the implicit assumption that consumers do not observe the calendar date at which previous transactions took place.

[18] If identity changes are unobservable, than an identity change amounts to an exit in terms of our model. A number of theoretical papers have analyzed the phenomenon of creation and trade of names: Tadelis (1999), Cabral (2000), Friedman and Resnick (2001), Mailath and Samuelson (2001).

to do so to be decreasing in the seller's reputation. In order to prove a result along these lines, we consider the following extension of our basic model. First, we assume that good sellers produce positive transactions with probability γ close to one but strictly less than one. This is important as we want the value function $V(P,N)$ to be strictly increasing in P when $N > 0$, whereas $\gamma = 1$ implies that $V(P,N)$ is constant for $N > 0$. Second, we assume that, in each period, the seller has a cost x of changing identity (excluding the opportunity cost of abandoning a reputation). We assume x is i.i.d. according to the c.d.f. $F(x)$. Our next result shows that sellers with a better reputation are less likely to change their identity.

Proposition 3 *Let $X(P,N)$ be the probability of identity change after history (P,N). Suppose that γ and μ_0 are close to, but lower than, one. If $N > 0$, then $X(P,N)$ is decreasing in P and increasing in N.*

Proof: First notice that the value function, $V(P,N)$, is increasing in P and decreasing in N. Following the argument in the proof of Proposition 1, we see that, if μ_0 is sufficiently close to one, then there is a unique perfect Bayesian equilibrium. It follows that the value function is entirely determined by μ. Bayesian updating implies that μ is increasing in P and decreasing in N.

Second, notice that the value of starting a new history is independent of the current history. Taken together, these facts imply that the incremental value of changing one's identity is decreasing in P and increasing in N. ■

■ **Extension: free-entry equilibrium.** The natural next step after looking at the possibility of entry and exit is to look for the existence of a free-entry equilibrium. Suppose there is a measure of potential sellers, a fraction $\bar{\mu}$ of which is of high type. Suppose also that sellers don't know their type until they pay the entry cost, which we denote by k. Finally, suppose that an existing seller can change his name at a low cost. Given the equilibrium effort strategies described in Proposition 1 and the exit strategies described in Proposition 3, there is an equilibrium belief μ_0 that takes into account both the equilibrium entry strategies and the equilibrium exit strategies. Furthermore, μ_0 is decreasing in the probability of entry and exit by low types, whereas the expected value from entering is increasing in μ_0. Together, these facts imply that there is a unique value of μ_0.

Notice that Proposition 1 is valid for *any* value of μ_0. Following a sufficiently long streak of Ps, the first N must lead to lower sales and an increase in the hazard rate of future Ns. What happens at the initial stages, however, does depend on the nature of the free-entry equilibrium. In particular, rational buyers will factor in the fact that "new" sellers can either be genuinely new sellers, taken from the pool with a fraction of $\bar{\mu}$ good sellers; or they can be bad sellers who secretely changed their identity. We would therefore expect a particularly negative premium on new sellers: a new name partly signals low type.[19]

■ **Extension: "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 on 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 cheaper) 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?

Suppose that 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.[20] If $V^i(P,N)$ $(i = H, L)$ is the seller's value given a history (P,N), then the value of an initial reputation is given by $V^i(P_0,0) - V^i(0,0)$. The answer to our question is then given by the difference of differences

$$\Delta(P_0) \equiv \left(V^H(P_0,0) - V^H(0,0)\right) - \left(V^L(P_0,0) - V^L(0,0)\right).$$

In an appendix, we derive the value of $\Delta(P_0)$. It is given by a complicated expression, one we have not been able to sign analytically. However, for reasonable values of the main parameters we find that $\Delta(P_0) > 0$, that is, a good-type seller is willing to pay more for an initial reputation than a bad-type seller.[21]

[19] See Tadelis (1999), Cabral (2000). Although we do not test for this prediction, anecdotal evidence seems broadly consistent with this prediction.

[20] This measure-zero assumption implies that buyers take the initial record as a genuine selling record. Although it is not a necessary assumption for the results we derive here, it greatly simplifies the analysis.

[21] Assuming that the seller chooses high effort while $N = 0$, we have $T(P,0) = \frac{1}{1-\alpha}$ and $T(P,N) = \frac{1}{1-\beta}$ for $N > 0$. Table 6 (to be discussed below) suggests that $T(P,0) = 350$, whereas $T(P,0) = 175$. Solving for the relevant parameters, we get $\alpha \approx 0.997$ and $\beta \approx 0.994$. On the other hand, Table 1 suggests a value of μ_0 between 0.25 and 0.5. For these values and for $\delta = .9$ (which we set somewhat arbitrarily), we find that $\Delta(P_0) > 0$. Moreover, for all other parameter value constellations we tried we also obtained $\Delta > 0$, suggesting that

■ **Other extensions.** There are other possible extensions of the basic model, two of which we mention here. First, we could consider variations in the auction mechanism, for example, reserve prices. However, as long as the equilibrium probability of entry, $p(v)$, and the expected price, $\pi(v)$, are increasing in v, then the sign predictions of Corollaries 1–3 remain valid. Second, we assumed that the seller offers one object per period. Alternatively, and more realistically, we could assume that there is an increasing trend in the number of objects offered per period. In this case, Corollary 1 applies to the de-trended sales rate (or the sales growth rate, as we will consider in Section 5.1).

The rest of the paper proceeds as follows. In Section 5, we test the empirical implications of Proposition 1, specifically Corollaries 1, 2, and 3. We also consider alternative explanations for the prediction in the latter result. Before that, in the next section, we describe the data sources we used.

4 Data description

We used Perl-based "spider" programs to download data directly from eBay's website at monthly intervals between October 24, 2002 and March 16, 2003. We focused our attention on auctions of (arguably) ex-ante homogenous goods to minimize the impact of object-level heterogeneity, but we also wanted to capture possible sources of variation across objects with different characteristics. Hence we collected transaction level information on the following objects (displayed in Figure 3):[22]

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.[23] We looked at two different kinds of coins. The first type of coin we look at are 1/16 oz. 5 dollar gold

there may be a general result for our basic model.

In a model with more than two types, Mailath and Samuelson's (2001) Proposition 4 suggests that $V(P_0, 0) - V(0, 0)$ is highest for the intermediate 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.

[22] eBay stores data on completed auctions for 30 days. We attempted to get data from all completed auctions in the above period. Several times our spider program was stalled by network problems. We believe that any data loss from this technical problem is independent of the nature of the auction.

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

coins of 2002 vintage (gold American Eagle), produced by the U.S. mint. The second type of coin are 2001 silver proof sets, a set of ten coins of different denominations, also produced by the U.S. mint. An important difference between these two types of coins is that, while the proof set is in mint condition, the gold coin may come in various grades.[24] There is no grading in proof sets, these are all in "mint" condition, as 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 because, according to the FBI's online fraud investigation unit, most customer complaints regarding online auction fraud arise from laptop auctions. We further chose this object because, while notebook computers tend to come in many different configurations (regarding memory, disk space, peripherals, screen size), this particular IBM model seemed to have relatively minor differences in configuration compared to other manufacturers. The average sale price of the Thinkpad T23's in our data set was \$580.

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

The Data Appendix and Table 12 present various summary statistics from the transaction level data.

Along with transaction-level data, we also downloaded each seller's "feedback summary" page, as shown in Figure 2. We recorded the information regarding feedback in the most recent week, month and six months. We also recorded the seller's entire sequence of reviews.

We should note that the feedback record of the seller and the transaction-level data can be linked only for the particular transactions we sampled. That is, a seller may be selling many different kinds of laptops other than IBM Thinkpads, or different coins and Beanie Babies; or she may be operating on many other markets as well. However, the only transaction level information (i.e., prices, object descriptions, number of bidders) we have are for the

[24] In the data, we found that the gold coins came in three different "grades:" MS–70, MS–69 and MS–67, in decreasing order of value. By contrast, the proof set is produced by the U.S. Mint and preserved in plastic container.

Table 1: Distribution of feedback aggregates across sellers.

	Number of Positives	Number of Negatives	Number of Neutrals	$N/(N+P)$ (entire history)
Mean	1,625	4.9	7.2	0.009
Std. Dev.	3,840	25.1	33.5	0.038
Min.	0	0	0	0
Max.	52,298	651	654	1
1%	0	0	0	0
5%	5	0	0	0
10%	18	0	0	0
25%	99	0	0	0
50%	397	1	1	0.0028
75%	1,458	3	4	0.0092
90%	4,361	9	13	0.021
95%	7,134	19	29	0.034
99%	15,005	52	86	0.068
N	819	819	819	795

particular categories for which we sampled this data. Unfortunately, the construction of entire transaction histories for many of the sellers in our sample is infeasible, since eBay does not allow users to access transaction level information that is more than 30 days old, and many of the sellers in our sample have been on eBay for much longer than that.[25]

■ **Seller characteristics.** We now use the feedback summary data to report some characteristics of the cross-section of sellers operating in these markets. Table 1 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.

The average seller in our sample had 1625 total feedback responses. The median seller had 397. The largest seller has 52,298 feedback responses, and

[25]In principle, one could construct forward-looking transaction and feedback histories for a sample of "young" sellers; however, such a sample would not necessarily represent the cross section of sellers operating in these product markets.

the smallest had 0 (is yet to be rated, even though she sold something). We found the distribution of seller sizes (proxied by number of feedback points they got) to be approximately lognormal. Sellers were largest in the market for Thinkpads, followed by teddies, gold coins and the proof sets.

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 1, the average seller in our sample has 4.9 negative feedback points, corresponding to 0.9% of all comments. The maximum number of negative feedbacks received by a seller is 819, but this seller took part in 52298 transactions. Also observe that the median seller in our sample has only one negative, and more than a quarter of the sellers have no negatives.[26]

One issue regarding the interpretation of comments is whether neutral comments are closer to positives or to negatives (our model did not allow for neutral comments). Our subjective impression, after browsing through eBay community chatboards where users discuss issues regarding the feedback system, is that the information contained by a neutral rating is perceived by users to be much closer to negative feedback than positive. Indeed, observe that in Table 1, the distributions of neutrals and negatives across sellers are extremely similar. The average seller received 7.2 neutral comments in her lifetime, with a median of again 1. Given this striking similarity, we will henceforth lump negative and neutral comments together when talking about "negative" comments.

5 Testing the model's empirical implications

We now use our data to test the empirical predictions of the model presented in Section 3. First, we use panel data on sellers' feedback records to examine the impact of negative feedback on the sales rate (cf Corollary 1 of Proposition 1, Proposition 2). Second, we use cross-section data to examine the impact of reputation on sales price (cf Corollary 2 of Proposition 1). Third, we use panel data to analyze the interarrival times of the first and subsequent negatives (cf Corollary 3 of Proposition 1). Fourth, we inquire whether and how sellers in

[26]Some negative comments for sellers have the following textual content: "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." On the other hand, 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."

our sample exit (cf Proposition 3). Finally, we report a series of interesting findings on how eBay users choose between buying and selling activities (see the concluding discussion in Section 3).

5.1 Negative feedback and sales

Corollary 1 implies that, after the first Negative is received, the fraction of objects offered by the seller that are actually sold decreases. Moreover, Proposition 2 implies that the seller offers fewer objects upon the first Negative. Together, these results imply lower sales after the first Negative. In this section, we test this implication.

Our typical seller receives his first negative during the early stages of his career. During this period, sales rates are typically increasing over time. Our theoretical results and empirical tests should therefore be adjusted to this fact. Accordingly, we test the implications of Corollary 1 and Proposition 2 by looking at the impact of the first Negative on the seller's growth rate.

It is difficult to construct entire transaction histories for eBay sellers, especially those that existed before a 30 day window preceding the sampling period.[27] Moreover, since negatives are rare occurrences, this type of data is quite costly to acquire by Web-spidering. However, one can utilize eBay's comprehensive feedback comments database to construct retrospective sales histories for a cross section of sellers. As was discussed in Section 2, eBay displays every single feedback comments received by a user over their entire lifetime. Unfortunately, these comments do not yield price information, but under Assumption 2 (a constant fraction of sales are accompanied by feedback comments), we can use the number of feedback comments received by a seller as a (lagged) measure of sales. Specifically, we constructed a proxy for weekly sales totals by summing the total number of sales-related feedback comments received by a seller in a given week. We then marked the weeks in which a seller received her first, second, and third Negatives.

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. We did not count the negatives that the sellers received during such episodes, and concentrated on the timing between *de novo* Negatives.

We then averaged the weekly sales rates over a four week "window" be-

[27] This is the period during which the complete data regarding a particular transaction is available.

fore and after the week in which the seller got his first (or second, or third) negative.[28] We also calculated the sellers' "before" and "after" weekly growth rates by averaging growth rates over these two four-week windows. Finally, we conducted paired t-tests of the null hypothesis of equality of growth rates before and after the first Negative.

The results, reported in Table 2, are striking: For all four object categories, the impact of the first Negative is to slow growth by 14% a week, and this difference is highly statistically significant. The difference in growth rates before an after the second Negative is positive. However, except for Golden American Eagle, the difference is not statistically significant. The impact of the third Negative also does not appear to be statistically significant.

Two notes are in order. First, our exercise depends crucially on the assumption that the probability of feedback is the same before and after a Negative. However, this is only a problem if buyers are somewhat reluctant to give *positive* comments about a seller after the seller has received her first (or second or third) Negative. Intuition suggests that the opposite is more likely to be true.

Second, our strategy for collecting seller histories retrospectively may imply a sample bias (we only have data for surviving sellers). In particular, there may be sellers who exited after receiving the first Negative and are thus excluded from our sample. But intuition suggests that, if anything, this reinforces the point that the first Negative has a negative impact on sales.

In summary, there is significant evidence that the first Negative has a strong negative impact on the seller's growth rate; and that subsequent Negatives have lower or no impact on the sales rate.[29]

5.2 Reputation and price

The theoretical model in Section 3 implies that differences in feedback histories lead to differences in the sale price of otherwise identical objects *across* sellers with different feedback aggregates. To investigate the empirical nature of these differences, several papers in the prior empirical literature on eBay have run regressions of the form:[30]

[28]For many sellers, longer evaluation periods would include subsequent Negatives. We believe a four-week window is a good balance between avoiding loss of data and statistical significance.

[29]As Footnote 2 of Table 2 states, we computed growth rates as differences in logs. When computed as the ratio $(x_{t+1} - x_t)/x_t$, we obtained different values but the same qualitative patterns.

[30]For surveys of these papers, see Bajari and Hortaçsu (2003), Resnick et al. (2003)

Table 2: Impact of negatives on sales growth (%).

	Avg. Week. Growth R.	Object Thinkpad	Proof set	G. Eagle	B. Baby
First Negat.	Before	7.12	6.85	9.04	14.19
	After	−6.76	−7.51	−3.89	−4.28
	Difference	−13.88***	−14.36***	−12.92***	−18.47***
	Std. Error	4.88	3.45	3.58	3.69
	N	66	130	95	136
Second Negat.	Before	3.96	4.50	−0.22	7.68
	After	9.93	8.00	9.47	8.03
	Difference	+5.97	+3.50	+9.69**	+0.36
	Std. Error	5.00	5.96	4.82	6.12
	N	37	78	70	83
Third Negat.	Before	9.19	3.80	3.58	2.00
	After	5.28	2.48	−2.09	10.25
	Difference	−3.90	−1.32	−5.68	+8.24
	Std. Error	6.14	3.22	7.44	6.23
	N	28	57	52	64

Notes:

1. Standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).
2. Weekly growth rates are based on the number of sales-related feedbacks received by the seller.
3. Growth rate in week t = ln(no. feedbacks in week t) - ln(no. feedbacks in week $t-1$).
4. Weekly growth rates are averaged over 4 week periods taken before and after the reception of a negative.

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

Since we have data for a series of auctions across four homogeneous product categories, we follow the literature by running similar cross-sectional regressions.

Table 3 reports our first set of results. In these regressions, the dependent variable is the log of the highest bid registered in the auction.[31] Hence the coefficient estimates can be interpreted (loosely) as percentage changes in price. The regression in column (1) allows for heteroskedasticity across object classes and controls for object dummies. The coefficient on the percentage of negatives in a seller's feedback history is negative and implies that a one point increase in this percentage (at the mean value, from 1% to 2%) leads to a 9% decline in sale price. The coefficient on the total number of transaction 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%.

Observe that the magnitude of this estimate is close to the findings of several other cross-sectional studies. In particular, the 5% price premium implied by 1000 additional reviews is comparable to an 8% premium found by the field experiment of Resnick et al. (2003), which compared sales prices obtained by a seller ID with 2000 positive comments (and 1 negative), and a seller with about 15 positive comments (and zero negatives).

However, as first pointed out by Resnick et al. (2003), several unobservable confounding factors may render a "causal" interpretation of the reputation measure difficult. For example, sellers with better reputation measures may also be much better at providing accurate and clear descriptions of the items they are selling; hence their writing ability, and not their reputation, may be underlying the higher prices they are receiving.

The next set of results reported in Table 3 enable us to get a feel for the importance of such confounding factors in cross-sectional price 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 the percentage of negatives being no longer statistically significant (though the coefficient on total number of reviews becomes significant). Column (3) provides even more clear evidence that unobservable factors may be at work. In this regression, we include a dummy variable for the auctions run by *hdoutlet*,

[31] According to eBay rules this is equal to the second highest bid plus the bid increment.

Table 3: Cross-sectional regressions. Dependent variable: logarithm of price.

	(1)	(2)	(3)	(4)	(5)
Percentage of negatives	-9.051 (3.115)***	-9.051 (10.808)	-0.346 (7.415)	2.835 (7.618)	-0.400 (7.419)
Number of transactions	0.056 (0.040)	0.056 (0.027)**	0.004 (0.003)		
Age (in days)				0.015 (0.008)*	
eBay rating					0.012 (0.009)
Indicator for user *hdoutlet*			4.598 (0.543)***	4.698 (0.539)***	4.482 (0.576)***
logarithm of minimum bid	0.003 (0.000)***	0.003 (0.001)***	0.004 (0.001)***	0.004 (0.001)***	0.004 (0.001)***
Listing includes photo	-0.219 (0.060)**	-0.219 (0.147)	-0.080 (0.107)	-0.084 (0.129)	-0.083 (0.107)
Refurbished object	-0.415 (1.135)	-0.415 (1.079)	-2.259 (0.736)***	-2.214 (0.735)***	-2.263 (0.736)***
Paypal accepted	0.188 (0.205)	0.188 (0.200)	-0.049 (0.098)	0.034 (0.120)	-0.047 (0.098)
Credit cards accepted	0.365 (0.230)	0.365 (0.104)***	0.293 (0.104)***	0.281 (0.110)**	0.293 (0.104)***
Auction duration (days)	0.039 (0.022)	0.039 (0.019)**	0.042 (0.017)**	0.038 (0.019)**	0.042 (0.017)**
Peak hour	0.242 (0.215)	0.242 (0.168)	0.185 (0.164)	0.223 (0.182)	0.187 (0.165)
Day of week	-0.028 (0.019)	-0.028 (0.020)	-0.030 (0.020)	-0.031 (0.021)	-0.029 (0.020)
Weeks since start of sample	-0.014 (0.013)	-0.014 (0.017)	-0.004 (0.015)	-0.002 (0.016)	-0.004 (0.015)
American Eagle	0.398 (0.070)**	0.398 (0.515)	0.772 (0.501)	0.941 (0.521)*	0.765 (0.501)
Mint set	0.725 (0.058)***	0.725 (0.510)	1.104 (0.494)**	1.327 (0.503)***	1.099 (0.494)**
Beanie Baby	-1.069 (0.041)***	-1.069 (0.525)**	-0.571 (0.497)	-0.411 (0.514)	-0.579 (0.498)
Constant	3.554 (1.156)*	3.554 (1.203)***	2.787 (1.139)**	2.352 (1.198)*	2.797 (1.131)**
Observations	1114	1114	1114	1003	1114
R-squared	0.39	0.39	0.48	0.48	0.48

Notes:
1. Significance levels: 10, 5, 1 percent (one to three stars).
2. Robust standard errors (clustered by sellerid) in parentheses in columns (2)-(5).

the dominant seller (with close to 50% market share) in the Thinkpad market. This leads to the economic and statistical significance of the percentage of negatives and the length of the transaction record to disappear entirely, implying that the comparison of auctions of this seller vis-a-vis other, much smaller sellers, drives much of the finding in column (1).

The results in column (2) and column (3) suggest that factors other than differences across sellers transaction histories may affect the cross-sectional variation in prices; and it may be difficult for an econometrician to account for these factors since the econometrician is typically not a very knowledgeable buyer in these markets. In fact, a few of the other coefficient estimates in Table 3 also suggest that factors other than reputation scores play a larger role in the cross-sectional variation of prices. For example, 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%. Finally, longer auctions appeared to fetch higher prices (one additional day translates into 4% increase in price).

However, it may also be the case that the weakened results in columns (2) and (3) are due to a misspecification of how reputation measures should enter the regression. Hence, in column (4), we include the sellers age, measured in days (divided by 100) since her first ever feedback 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. Finally, in column (5) 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%.

In summary, the results in Table 3 suggest, at best, a rather weak connection between sale price and the reputation measures that eBay publishes. The results in columns (3) through (5) suggest that variables correlated with the length of a seller's transaction history (total number of reviews, age, and eBay's rating) appear to have a more robust relationship with price than the percentage of negatives.

■ **Impact of a change in website design.** One way to strengthen the case for a causal connection between cross-sectional variation in reputation and sale price is to exploit an exogenous change in reputation measures which is not correlated with the way sellers prepare their listings. We exploit the following exogenous change in eBay's website format: before March 1st, 2003, bidders would only see the seller's overall (net positive) feedback points next to the seller's name. On March 1st, 2003, eBay began to display the percentage of positive comments received by the seller, as well as the date when the seller registered on eBay (see Figure 2).[32]

In columns (1) and (2) of Table 4, we find that the interaction of the percentage of negatives with a dummy variable for the format change implies that the response of prices became more negative after the format change.[33] According to the regression results, the economic effect of a 1% increase in negative feedback was a 6% change in price before the format change (but insignificant), and a −8% change after the format change. Furthermore, the coefficient estimates on eBay's own reputation rating (ebayrating), which was the only reported reputation measure on the listing page before March 1, 2003, and variables that are highly correlated with this rating (such as the total number of transactions conducted by a seller, and the seller age measured in days) are lower after the change.

The results of these regressions suggest two things: that bidders respond to the reputation statistics published by eBay (the March 2003 dummy is significant), and that there might be costs to information acquisition and processing (the same information was available before and after March 2003).

5.3 Frequency of arrival of negative feedback

To test Corollary 3 of Proposition 1, we once again utilize seller feedback histories to construct a data set containing the timing of each negative/neutral feedback.[34] We measured "time" in two ways: number of sales transactions and days. As mentioned above, negative comments often came in the context of a "war of words" between seller and buyer. To prevent such incidents from biasing our results, we excluded consecutive negative comments by the same buyer. We also excluded any Negatives that were left within a two-day period after another negative.[35]

[32] We found out about this policy change by accident. We should point out that before March 1st, 2003, the information shown in Figure 2 was already available to bidders. However, in order 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 bidder to a new "feedback profile" page) and manually compute the ratio $N/(N+P)$.

[33] This regression corrects standard errors by allowing for heteroskedasticity at the seller level. We also added a dummy variable for a particularly large seller in the laptop market. Omission of either of these features lead to significance of the coefficient at higher levels.

[34] We excluded those negative/neutral comments that were received as a "buyer." There were only four instances of this in our sample.

[35] We also experimented with 1 day and 5 day periods. Our results are robust to the choice of window length.

Table 4: Impact of change in eBay's site design. Dependent variable: logarithm of highest bid.

	(1)	(2)
Percentage of negatives	6.603	6.335
	(8.770)	(8.827)
Total number of transactions	0.004	
	(0.004)	
Age in days	0.018	0.016
	(0.011)*	(0.011)
eBay Rating		0.016
		(0.012)
Percentage of negatives after format change	−14.764	−14.075
	(8.238)*	(8.450)*
Total number of transactions after format change	−0.008	
	(0.010)	
Age in days after format change	−0.011	−0.010
	(0.015)	(0.015)
eBay rating after format change		−0.018
		(0.018)
Indicator for auctions after format change	0.003	−0.005
	(0.379)	(0.382)
(Other auction level regressors omitted)		
Observations	1003	1003
R-squared	0.49	0.49

Notes:

1. Standard errors in parentheses. Significance levels: 10, 5, 1 percent (one to three stars).
2. Robust standard errors (clustered by seller id) in parentheses.

Table 5: Timing of first and second Negatives.

		Object			
		Th'pad	P. set	Eagle	Beanie
Time measured in number of	Time to first negative	129 (22)	441 (52)	459 (77)	399 (55)
sales-related feedbacks	Time 1st-2nd negative	60 (12)	220 (36)	210 (44)	278 (63)
(All sellers with 2+ negatives)	Difference	69*** (19)	221*** (52)	249*** (85)	121** (66)
	N	57	90	83	110
Time measured in number of	Time to 1st negative	117 (22)	460 (144)	296 (67)	418 (122)
sales-related feedbacks	Time 1st-2nd negative	43 (16)	130 (39)	94 (24)	135 (26)
(Sellers born after June 16, 2001)	Difference	74*** (19)	330*** (135)	201*** (53)	283** (113)
	N	28	11	17	23
Time measured in number of	Time to 1st negative	300 (36)	420 (16)	407 (42)	415 (43)
days	Time 1st-2nd negative	66 (13)	118 (16)	117 (16)	152 (24)
(All sellers with 2+ negatives)	Difference	233*** (37)	302*** (37)	302*** (37)	263** (49)
	N	57	90	83	110

Notes:
1. Standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).
2. Seller samples restricted to sellers with at least 2 negative feedbacks.
3. Retaliation feedbacks are excluded from consideration using procedure described in Section 5.1.
4. Sale vs. purchase nature of transactions are reported by eBay after June 16, 2001. Transactions preceding this period are classified by the algorithm described in Section 5.5.

Table 5 reports comparisons of the number of sales-related comments received by a seller until her first negative comment vs. the number of comments received between the first and second negative comments. The results are obtained by a regression of interarrival times of first and second negatives on a dummy variable that turns on for the second negative, controlling for seller fixed effects (equivalent to a within-seller paired t test) . The results in the first panel indicate that, for the Thinkpad, it takes on average 129 transactions before a seller receives his first Negative, but only 60 additional transactions before the second Negative arrives. The difference is significant at the 1% level. Similar results are obtained for the other three objects.

The second panel of Table 5 replicates the analysis on a subsample of sellers born after June 16, 2001. Prior to this date, feedback comments do not specify if the comment giver is a buyer or a seller. We used the actual textual remarks to guess the nature of the feedback. As we describe in detail in Section 5.5, this is not a perfectly accurate process. However, Table 5 suggests that our results are not biased by the noise introduced by our classification.

The third panel of Table 5 replicates the analysis with time measured in days. The difference between the interarrival times of the first vs. the second negative is again quite striking: in the Thinkpad market, for example, it takes on average 300 days for the first Negative to arrive, but only 66 days for the second one. Once again, the result is robust to restricting attention to subsamples with smaller/younger sellers.

In Table 6, we investigate whether the interarrival times of the third, fourth, fifth and sixth negatives are different from the second or the first. In column (1), we run a regression of interarrival times of negatives (measured in terms of sales-related comments received by the seller) on dummy variables for the first and third to sixth negative comments — i.e., the second negative is treated as the base case of the regression. The regression indicates that the first negative is indeed the slowest one to arrive. In fact, the interarrival times for the third to sixth negatives are not significantly different for the interarrival time leading to the second negative. Very similar results are obtained when we measured interarrival times in days (second column).

■ **Alternative explanations.** The results reported in the previous section support Corollary 3. At the very least, it appears from these results that there is something "special" about the 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 "scaling-up" effect: it might be pos-

Table 6: Timing of subsequent Negatives.

	Number of rated sales till Negative	Number of days till Negative
1st Negative	173.44	189.46
	(31.62)***	(25.82)***
3rd Negative	-17.05	-27.80
	(31.60)	(13.24)**
4th Negative	21.72	-18.38
	(29.96)	(12.85)
5th Negative	40.15	-15.30
	(37.14)	(15.21)
6th Negative	19.13	-26.73
	(30.49)	(12.67)**
	(seller f.e. not reported)	(seller f.e. not reported)
Constant (2nd Negative)	174.89	87.24
	(20.45)***	(10.45)***
Observations	1014	1014
Adj. R-squared	0.38	0.27

Notes: 1. Standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).

sible that a seller takes longer to acquaint himself with the market, and does not do that much business in the early days, implying that it takes a long time for the first negative to arrive. However, our results also hold true when we measure time in number of transactions.

The second alternative explanation is that buyers have a threshold of dissatisfaction above which they give a negative; and this threshold drops after the first negative. There are several behavioral mechanisms through which this can happen, and we consider these in turn.

One way in which such a "threshold decline" may occur is through a decrease in the cost of writing a negative comment. As we noted above, many negative comments are followed by a "war of words" between buyer and seller. Seller retaliation might impose an economic cost on the complaining buyer, especially if the buyer is also a seller. Such an effect would confound our results if the probability of retaliation by a seller in reaction to her first negative is higher than retaliation to her second negative, an explanation proposed by several eBay users we talked to.[36]

To investigate this possibility, 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 10,000 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 a 40% chance of being hit back, while a buyer leaving a neutral comment only has a 10% chance of being retaliated upon by the seller.

However, our data indicates that sellers are not more likely to retaliate upon their first negative, as opposed to subsequent negatives. In Table 7, 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 columns (1) and (2), 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 first regression, 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

[36]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 decided to evaluate its merits.

Table 7: Alternative explanations for differences in arrival times.

	Dependent variable			
	(1) Retaliation	(2) Retaliation	(3) Profile	(4) Profile
2nd Negative	0.016 (0.055)	0.025 (0.063)	0.011 (0.013)	0.011 (0.015)
3rd Negative	0.030 (0.059)	0.043 (0.068)	0.003 (0.015)	-0.003 (0.016)
4th Negative	-0.005 (0.064)	0.000 (0.069)	0.020 (0.020)	0.020 (0.021)
5th Negative	0.044 (0.068)	0.118 (0.074)	0.015 (0.018)	0.011 (0.018)
6th Negative	0.053 (0.071)	0.107 (0.073)	0.045 (0.023)*	0.040 (0.024)
Percentage of Negatives	4.664 (1.907)**		-0.053 (0.372)	
Number of transactions	0.000 (0.000)		-0.000 (0.000)	
eagle dummy	0.100 (0.120)	(seller f.e.)	-0.079 (0.038)**	(seller f.e.)
mint dummy	0.000 (0.094)		-0.087 (0.037)**	
teddy dummy	0.091 (0.089)		-0.071 (0.039)*	
Constant	0.115 (0.098)	0.239 (0.045)***	0.105 (0.043)**	0.038 (0.012)***
Observations	558	567	575	584
R-squared	0.03	0.38	0.06	0.38

Notes:
1. Robust standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).
2. Dependent variable Retaliation = 1 if buyer's negative comment is followed by seller's negative comment.
3. Dependent variable Profile = frequency of negative comments by the buyer who gave a particular negative comment.

in interarrival times of negative comments.

A second variation on the "threshold" story is that, in addition to time variation, there is also buyer variation in propensity to give negative feedback. So first negatives would primarily be given by negative-prone buyers, whereas subsequent negatives would originate in a wider set of buyers. To test this possibility, we downloaded the string of feedbacks that were left by every negative/neutral comment giver in our data set.[37] 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 7, columns (3) and (4), 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 subsequent negative feedback.[38]

To conclude our test of the "threshold" story, we directly tested the hypothesis that second negatives have a lower threshold than first negatives. We constructed a series of pairs of first and second negative comments. We then asked a third party (a student) to make a subjective evaluation as to which of the two remarks was more negative.[39] The results show that 51% of the second negatives were considered "nastier" then the corresponding first negative, a split that is not statistically different from 50/50.

Finally, we consider the possibility that buyers are influenced by other buyers' behavior (herding, conformism, etc).[40] imply that the events leading to the first negative are different than those leading to subsequent negatives. In particular, faced with poor performance by a seller with a perfect record, a buyer might be inclined to think that there is no ground for a negative feedback. For example, if there is a communication problem between buyer and seller, the former may attribute this to a problem with him or herself, not

[37]On eBay one can also observe what each user wrote about each other.

[38]Interestingly, our data suggests a lower critical threshold for giving negatives in the Beanie Babies market than in the laptop market: the average negative comment-giver in the laptop market gave negatives 10% of the time, whereas the average complainant in the Beanie Babies market complained only 3% of the time. We speculate that this result may very loosely be attributed to our observation that the Beanie Babies market on eBay can be seen as a "community of collectors" with frequent repeated interactions, where wrong doings are less tolerated, whereas transactions in the laptop market are not typically repeated.

[39]We randomly mixed the order of the comments so that the student could not tell which was the first, which was the second negative. We also allowed for the following possibilities: "repeat" (remarks are literally identical), "mistake" (remarks are clearly positive even though a negative was given), and "difficult to tell."

[40]There is an extensive psychology literature on this, including Asch (1946), Snyder and Canto (1979) and Hoch and Ha (1986).

Table 8: Reasons for negative feedback (%).

	First Negative	Second Negative
Misrepresented item	22	16
Bad communication	19	20
Item damaged	15	17
Item not received	10	13
Backed out	7	4
Angry / upset	7	7
Overcharged shipping	6	4
Slow shipping	6	10
Bad packaging	4	6
Feedback issues	3	3
Bid on own item	1	1
Total	100	100

with the seller. However, if the seller has already received a negative feedback, especially regarding the same problem that the buyer is now facing, then the buyer may have a greater inclination to attribute this to a problem with the seller and give negative feedback. This is especially true for aspects of the transaction that are more subjective and difficult to input (e.g., communication problems).

To consider this possibility we classified the first and second negative remarks according to their nature. The breakdown of the reasons for negative feedback is presented in Table 8. The buyer influence story should imply an increase in the relative importance of "subjective" problems in second negatives. However, the results suggest a very similar pattern for first and second negative (correlation greater than 0.92). Moreover, "item never sent," arguably the most objective reason for negative feedback, actually increases in relative importance (though by a small amount). At the opposite extreme, "bad communication," arguably the most subjective reason for negative feedback, also increases in importance (though by an even smaller amount).

In sum, the empirical evidence suggests that the behavioral change from the first to the second negative is on the seller side, not on the buyer side; i.e., Corollary 3 is confirmed by the data. Of course, there might be alternative explanations we might not have taken into account, but we believe it would be difficult to test for other behavioral hypotheses using the data available.

5.4 Reputation and exit

In order to test Proposition 3, we supplemented our data set by revisiting our sample of sellers in the first week of January, 2004, and checking whether they were still in business. There was considerable attrition in our sample: of the 819 sellers originally sampled in our sweep of the transaction-level data, we found that 152 had not conducted any transactions within the last 45 days (pre- and post-Christmas are the busiest seasons on eBay), and 61 sellers had not sold anything within the last 45 days, but had bought an item. We also could not locate the feedback records for 104 sellers in our sample, since eBay's database claimed that these seller ID's were no longer valid. Hence, .

We then ran probit regressions of an "exit" outcome on seller's observable reputational statistics as of May 2003 (at the end of our initial sampling period). As explanatory variables, we consider (a) the fraction of negatives and neutrals and (b) the total number of positives.[41] We ran the probit using different definitions of what constitutes an "exit," and using different subsamples of the data. In the first column of Table 9, we define an exit as any one of the three events mentioned above (no transactions in the last 45 days, no sales in the last 45 days, invalid ID). The second column classifies sellers who are still making purchases as still being in the sample. The third column assumes sellers with invalid IDs are still in the sample, and the fourth column combines the exclusions in the second and third columns.

The regression results, reported in Table 9 as marginal probit effects, imply that a tenfold increase in the total number of positives (as of May 2003) translates into a decline in exit probability (in January 2004) of between 14 to 21%. This effect is highly significant for all specifications in Table 9, and is in concordance with the prediction in Proposition 3. Also, a 1% level increase in the percentage of negatives in a seller's record (i.e., from a sample average of 1% to 2%) translates into an increase in exit probability of 1.6 to 2.1%. This effect appears economically smaller and less significant statistically; but it also has the right sign as predicted by Proposition 3.

In the next four columns of Table 9, we investigate whether the marginal probit effects show any differences across the different objects (using the "exit" definition in the first column). As can be seen, the "seller history length" effect is quite significant for all object categories (declines in exit probabilities between 18% and 27% for a tenfold increase in the length of the seller's history). Notice, however, that the correlation of percentage of negatives in May 2003

[41] We also ran specifications with the number of negatives on the right hand side; these did not lead to significant coefficients on this variable.

Table 9: Can reputational variables predict seller exits? Dependent variable: seller has exited by January 4, 2004.

	Subsample			
	All exit events	Still- -sellers excluded	Invalid ID excluded	Invalid ID and still- -sellers excluded
Log. number negat. May 03	0.066 (0.049)	0.085 (0.046)*	0.071 (0.045)	0.090 (0.039)**
Log. number posit. May 03	-0.170 (0.024)***	-0.136 (0.022)***	-0.181 (0.021)***	-0.143 (0.018)***
Observations	818	818	818	818
	Subsample			
	Laptop sellers	Golden sellers	Silver sellers	Beanie sellers
Log. number negat. May 03	0.026 (0.105)	0.131 (0.092)	0.037 (0.150)	0.157 (0.095)*
Log. number posit. May 03	-0.164 (0.049)***	-0.151 (0.044)***	-0.304 (0.093)***	-0.200 (0.045)***
Observations	199	255	115	249

Notes:
1. Robust standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).
2. Probit marginal effects are reported.

with subsequent exits is not very significant for objects other than Beanie Babies. For Beanie Babies, the magnitude of the coefficient estimate implies that an increase from 1% to 2% of negatives in a seller's record translates into 12.5% higher exit probability.

Next, we investigate whether the "exits" we see in our data set are accompanied by opportunistic profit-taking by sellers, and whether reputational variables can predict such behavior.[42]1 In order to do this, we collected data on the last 25 sale transactions conducted by exiting sellers, and counted the number of negative comments for these last 25 sale transactions. Some of the examples were quite striking: one of the sellers in our sample, who had 22755 positives, racked up 11 negatives in her last 25 transactions; whereas she had a total of 54 negatives in her previous transactions (the percentage of negatives and neutrals over her overall history was 0.6%, versus 44% in the last 25 transactions). On average, the percentage of negatives in the last 25 comments of exiting sellers (excluding those who remained as buyers and those sellers whose ID's became invalid, and thus we could not get data) was 4.38%, as opposed to an average 1.61% over their entire histories. This difference is statistically significant at the 1% level.

To see if reputational statistics as of May 2003 have any predictive power over such "opportunistic" exits, we ran probit regressions of an indicator variable for the occurrence of an "opportunistic" exit on the reputational statistics. We defined the indicator variable on the left-hand side to be equal to 1 if the percentage of negatives within the last 25 transactions of a seller was more than twice the percentage of negatives within the seller's entire history.

The results of these regressions, reported in Table 10, indicate that, for the entire sample of sellers, a ten-fold increase in a seller's count of negatives is correlated with a 5% increase in "opportunistic" exit as defined above. The coefficient estimate on the log number of positives is smaller: a 2% decrease in "opportunistic" exits.

Overall, the results are consistent with Proposition 3. Moreover, the "end-of-life" increase in the number of negatives suggests that continuation incentives play an important role in sellers' behavior.

5.5 Who tries to buy a reputation?

Casual observation of feedback histories suggests that many sellers appear to start out as "buyers," completing a string of purchases before attempting

[42]For a model of opportunistic use of reputation, see Phelan (2001). See also Gale and Rosenthal (1994).

Table 10: Opportunistic exits. Dependent variable: seller has exited by January 4, 2004 with an abnormal number of Negatives.

	All sellers	Laptop sellers	Golden sellers	Silver sellers	Beanie sellers
Log. number negat. May 03	0.050 (0.019)***	0.048 (0.026)*	0.072 (0.025)***	-0.076 (0.076)	-0.008 (0.045)
Log. number posit. May 03	-0.017 (0.010)*	-0.026 (0.013)**	-0.024 (0.011)**	0.030 (0.045)	0.018 (0.022)
Observations	818	199	255	115	250

Notes:
1. Robust standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).
2. Probit marginal effects are reported.

their first sale. As an example, Figure 4 plots the percentage of sell vs. buy transactions by user *bearsylvania*, an established Beanie Baby dealer, as a function of the number of weeks he has been active on eBay. As can be seen, *bearsylvania* started out as a buyer first, and quickly changed the pattern of his transactions from purchases to sales.

To estimate the prevalence of this phenomenon, we looked at the first and last twenty comments received by each seller. We then defined a seller as having switched from being a buyer to being a seller if more than 50% of the first 20 comments referred to purchases, and more than 70% of the last 20 comments referred to sales.[43]

An important difficulty with implementing the above coding scheme with our data is that eBay does not report a buyer/seller classification for feedback comments received prior to June 16, 2001. Since about two-thirds of our sample sellers began their eBay careers prior to this date, we made our own assignment as buyer vs. seller based on the textual content of the comments.[44]

[43] To check the robustness of this definition of a "switch," we defined a second indicator with thresholds 40% and 80%, respectively.

[44] We automated the classification procedure by first calculating the empirical frequencies of word-stems like "buy," "sell," "pay," "ship," "pack" across buyer/seller categories in a subsample of the post-June 16,2001 data. We then compared the likelihood of a given comment to be a "buyer" or "seller" comment based on the presence of these keywords. The accuracy of our classification of "seller" comments was remarkable: for the post-June 16,2001 data (for which we have eBay's classifications) we were able to classify all but 117 of

Given the assumptions that go into our classification scheme, we will report some of our results for these two subsamples of sellers separately.

We found that 38% of Beanie Baby sellers, 22% of laptop seller, 31% of gold coin sellers, and 31% of proof set sellers followed the "buy first, sell later" strategy (as defined above). We also found that, on average, 81% of a seller's last 20 transactions were sales, compared to 46% of the first 20 transactions. A paired t-test of equality of the two percentages revealed a strongly statistically significant increase in the percentage of sales (t-statistic equal to 25).[45]

These results show that "buying first and selling later" is a widespread phenomenon on eBay, and is somewhat more prominent in some object categories than others. For example, eBay is widely known as one of the main trading venues for Beanie Babies. It is conceivable that Beanie Baby enthusiasts first start out as buyers in this market, and switch to selling once they accumulate an inventory. On the other hand, laptop sellers are more likely to sell items they have acquired through other channels.

Next, to investigate the prediction of our theoretical model, we investigated the correlation of the "buy first sell later" indicator variable with the percentage of negatives in a seller's record, and the length of the seller's record. Column (1) of Table 11 reports the marginal effects of a probit regression using the sample of sellers who joined eBay after June 16, 2001 (i.e., the set of sellers for whom we have direct data from eBay). This regression suggests that a 1% level increase from the mean value of 0.7% of negatives to 1.7% negatives is correlated with a 6.4% decrease in the probability that the seller "switched" from being a buyer to a seller. The length of the seller's record does not have significant correlation with switching behavior.

In column (2) of this table, we repeat the same probit regression for sellers who started their career before eBay began to report buyer/seller classifications of received feedback. The results appear very similar in sign and magnitude

12952 comments correctly. Our classification of "buyer" comments was less accurate, since most of these buyer comments contain very little information (we checked to see if human classification performed better in a subsample of comments; the improvement was marginal, precisely due to the lack of informative content). In particular, we classified 1934 of 5035 "buyer" comments as a "seller" comment, a 60% accuracy rate. Hence, our classification scheme is biased towards finding "sellers" as opposed to "buyers." To address this problem, when computing the percentage of a sales-related comments that a user gets within a given time period, we add 17% (the average bias in the control sample) on top of the percentage computed using our classification scheme.

[45] To make sure that these results were not driven by the assumptions needed to construct the buyer/seller classification for sellers with missing data, we repeated the same exercise using the post-June 16, 2001 sample of sellers. We found that, on average, 77% of last 20 transactions were sales, as opposed to 46% of the first 20 transactions. Once again the paired t-test strongly rejects equality.

Table 11: Who tries to "buy" a reputation?

	(1)	(2)	(3)	(4)	(5)	(6)
No. of comm./1000 (by May 2003)	-0.006 (0.034)	-0.008 (0.006)	-0.004 (0.005)	-0.002 (0.004)	-0.011 (0.005)**	0.001 (0.002)
Percent. negatives (by May 2003)	-6.372 (3.068)**	-6.093 (2.967)**	-5.987 (2.138)***	-4.774 (2.024)**	2.803 (1.272)**	1.582 (0.676)**
Gold coins	0.019 (0.099)	0.022 (0.106)	0.050 (0.073)	0.043 (0.071)	0.048 (0.060)	-0.042 (0.022)*
Silver proof sets	0.066 (0.098)	-0.011 (0.103)	0.047 (0.070)	0.087 (0.070)	0.062 (0.058)	-0.039 (0.024)
Beanie Babies	0.088 (0.096)	0.105 (0.105)	0.127 (0.071)*	0.151 (0.071)**	0.141 (0.062)**	-0.009 (0.026)
Seller switch from buying to selling					-0.008 (0.032)	0.016 (0.022)
Observations	234	384	618	618	618	618

Dependent variables:

1. Seller switched from buying to selling. Sample of sellers who joined eBay after June 16, 2001. All transactions classified as "buy" or "sell" by eBay.
2. Seller switched from buying to selling. Sample of sellers who joined eBay before June 16, 2001. Transactions were classified using the method described in Section 5.5.
3. Seller switched from buying to selling: pooled sample.
4. Seller switched from buying to selling: stricter definition of "switch" indicator.
5. Seller exited (according to definition in second column of Table 9).
6. Seller exited opportunistically (according to definition in Table 10).

Notes:

1. Robust standard errors in parentheses. Significance levels 10, 5, 1 percent (one to three stars).
2. Probit marginal effects are reported.

to the post-June 16, 2001 sample. In column (3), we pool the two samples together and find once again that a 1% level increase in percentage of negatives is correlated with a 6% decrease in "switching" probability. In column (4), we conduct a robustness check on our dependent variable by changing the threshold of being a "buyer" to having less than 40% of transactions as sales, and the threshold of being a "seller" to having more than 80% of transactions as sales. This modification does not appear to have an important effect on the coefficient estimates.

The coefficient estimates reported in Table 11, columns (1)-(4) suggest that "better" sellers, i.e. the ones with a lower percentage of negatives, are more likely to have switched from being buyers in their early career to becoming sellers later on, a fact that is consistent with our theoretical prediction. Note, however, that the sellers in our sample do not literally "buy" their reputations; they purchase objects from other eBay users, and pay promptly to get positive feedback. Some of these users may not have strategic motives in mind when doing this — they may simply start out as collector/enthusiasts who first buy objects, and then discover that they can make money by selling them. Moreover, in principle, an interested buyer can freely inspect, just like we did, whether a seller accumulated her feedback record by buying or selling. One might think that strategic incentives for "purchasing one's reputation" will be curbed by this reporting activity. In fact, we suspect that eBay's motive in reporting this information is exactly this reason; but we haven't been able to confirm this suspicion. However, the comparison of the results in Column (1) and (2) does not indicate that the reporting policy of eBay has had any effect on who tries to buy a reputation. One would have expected the magnitude of the coefficient estimate in Column (2) to be larger.[46] A possible explanation of this last result may be that, just like in Section 5.2, buyers may not bother to check the details of feedback tables (and that sellers know this).

A last sanity check on our strategic interpretation of the buying/selling patterns is to see whether sellers who switched from buying to selling were more or less likely to exit eBay. One might expect that if better sellers are more likely to "switch," they might be less likely to exit. The probit regressions in Columns (5) and (6) use indicator variables for seller exits, and "opportunistic exits" (as defined in Section 5.4) as dependent variables. Unfortunately, we do not see an economic or statistically significant difference between the exit patterns of switching sellers as opposed non-switching sellers.

[46]We conducted a regime-shift test by conducting the pooled regression with a dummy interaction term for the reporting policy change. We did not find a significant difference in the coefficient estimates.

In sum, although we do find indisputable evidence for the existence of "switching" behavior on eBay, our evidence for a clear economic incentive to do so ("reputation building") is somewhat indirect. In particular, we do find some evidence supporting the prediction of our theoretical model, that (ex-post) better sellers are more likely to undertake such costly reputation building activities early in their careers. However, this evidence is not corroborated by a clearly identified economic incentive to "purchase" a reputation. We hope future research can improve upon our methodology to further investigate this theoretical prediction.

6 Conclusion

Our analysis points to various empirical features of eBay seller dynamics. In particular, we find that, upon receiving their first negative comment, the sales growth rate drops dramatically; moreover, the arrival rate of subsequent negative comments increases considerably. We propose a theoretical model, featuring adverse selection and moral hazard, that explains these facts as follows: the arrival of the first negative comment damages the seller's reputation to the point that his incentives to exert effort are minimal. Lower effort and worse reputation explain the drop in sales rate. Lower effort also explains the increase in the hazard rate of subsequent negatives.

We considered alternative theories of buyer behavior as explanations for the increase in negative comment hazard rates, but none seems to do the job. What about alternative theories of seller behavior? One possibility is a pure adverse selection model with changing types.[47] Specifically, suppose that the seller's type evolves according to a Markov process. After an initial period in the high-type state, a shift to low type would increase the likelihood of the first and subsequent negative comments, consistently with our empirical finding. Fishman and Rob (2002) propose a model with these features, assuming that bad type is an absorbing state. Their model does imply the stylized facts described above. Note however that it is essential that good types may become bad but not vice-versa. Otherwise, our evidence that the arrival rate of subsequent negatives remains flat seems to reject the pure adverse selection story.[48]

[47]The literature on firm growth and industry evolution frequently considers the possibility of firm efficiency evolving according to a Markov process. See Jovanovic (1982), Hopenhayn (1992), Ericson and Pakes (1995). More recently, some reputation models have explicitly considered the possibility of changing types. See Mailath and Samuelson (1998), Phelan (2001).

[48]Fishman and Rob's (2002) model also has the implication that sales rate (and price)

The opposite alternative theory, pure moral hazard, is more difficult to evaluate. In fact, for a wide range of values of the discount factor, there exist multiple equilibria, potentially with very different patterns. In the limit when the discount factor approaches one, "almost any" sequence of actions belongs to some equilibrium path; that is, theory has very little bite. We may, however, consider some particular patterns suggested in the literature. Dellarocas (2003), following earlier work on collusion by Green and Porter (1984) and Porter (1983), suggests a stationary mechanism where poor performance is "punished" by buyers for a period of time. However, the empirical evidence seems at odds with the stationary prediction.

Closer to the patterns we find in the data are the pure moral hazard models of Klein and Leffler (1981) and Shapiro (1983).[49] In these models, seller reputation is part of a bootstrap equilibrium: buyers pay a high price and expect high seller effort. If a signal of low effort is observed, then buyers "boycott" sellers by not purchasing again. Reputation is therefore a valuable asset. To be consistent with a free-entry equilibrium, we must impose some endogenous cost to acquiring a reputation. Money burning in the form of advertising or low introductory prices do the job. Advertising does not seem to play a big role on eBay. Moreover, prices are typically set by buyers, who submit bids, not by sellers.[50] However, we can exploit the bootstrap nature of the equilibrium to create an endogenous cost of acquiring a reputation: suppose that, for the first t periods, sellers exert low effort and, consistently with this expectation, buyers bid low values. We can then find a t that satisfies the zero-discounted-profit entry condition.[51] However, the empirical evidence does not seem consistent with this: Given an equilibrium path of low effort, then high effort, then low effort, we would expect the distribution of first negative arrival time to be bimodal. However, a simple non-parametric analysis suggests a distribution close to log-normal, certainly not bimodal.

Regardless of which theoretical model best explains the data, an important conclusion of our paper is that eBay's reputation system gives way to noticeable strategic responses from both buyers and sellers. That is, the mechanism has "bite". Obviously, this does not imply that the current structure of the

increase over time until the seller's type becomes bad. This is consistent with eBay data. However, it depends crucially on the assumption of word-of-mouth effects, which we believe are not that important on eBay.

[49]See also Friedman (1971), Telser (1980) for earlier work.

[50]This is not entirely true. Sellers can set minimum bids and "buy now" options. In the limit, a very low "buy now" option would essentially amount to a low posted price.

[51]This is similar to the idea of "building trust." See Datta (1997), Watson (1999), Watson (2002).

system is optimal. In fact, we believe an exciting area for future research is precisely the design of an efficient reputation mechanism.

To conclude, we should also mention that 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 3). A natural next step is thus to study the strategic motives underlying various agents' feedback behavior. This we plan to do in a new empirical project (Cabral, Hortaçsu and Yin, 2003).

Appendix: Derivation of $\Delta(P_0)$

From the text,

$$\Delta(P_0) \equiv \left(V^H(P_0,0) - V^H(0,0)\right) - \left(V^L(P_0,0) - V^L(0,0)\right).$$

In order to compute $\Delta(P_0)$, we must first compute the value functions. For a good seller, the value function may determined recursively. Let $v(P,N)$ be the buyers' willingness to pay given history (P,N).[52] Then

$$V^H(P,0) = v(P,0) + \delta V^H(P+1,0),$$

leading to

$$\begin{aligned} V^H(P,0) &= \sum_{k=0}^{\infty} \delta^k v(P+k,0) \\ &= \sum_{k=0}^{\infty} \delta^k \frac{\mu_0 + (1-\mu_0)\alpha}{\mu_0 + (1-\mu_0)\alpha^{P+k}}, \end{aligned}$$

where we use the fact that, by Bayesian updating, $v(P,0) = \frac{\mu_0+(1-\mu_0)\alpha}{\mu_0+(1-\mu_0)\alpha^P}$.

For a bad seller, we know that, if $N > 0$, then $V^L(P,N) = \beta/(1-\delta)$. When $N = 0$, we have

$$\begin{aligned} V^L(P,0) &= v(P,0) + (1-\alpha)\delta V^L(P,1) + \alpha\delta V^L(P+1,0) \\ &= v(P,0) + (1-\alpha)\delta V^L(P,1) + \\ &\qquad +\alpha\delta\left(v(P+1,0) + (1-\alpha)\delta V^L(P,1) + \alpha\delta V^L(P+2,0)\right) \\ &= \ldots \\ &= \sum_{k=0}^{\infty} (\alpha\delta)^k v(P+k) + \frac{1-\alpha}{1-\alpha\delta} \cdot \frac{\beta}{1-\delta} \\ &= \sum_{k=0}^{\infty} (\alpha\delta)^k \frac{\mu_0 + (1-\mu_0)\alpha}{\mu_0 + (1-\mu_0)\alpha^{P+k}} + \frac{(1-\alpha)\beta}{(1-\alpha\delta)(1-\beta)}, \end{aligned}$$

Finally,

$$\begin{aligned} \Delta(P_0) &= \sum_{k=0}^{\infty} \left((\alpha\delta)^k - \delta^k - (\alpha\delta)^{P_0+k} + \delta^{P_0+k}\right) \frac{\mu_0 + (1-\mu_0)\alpha}{\mu_0 + (1-\mu_0)\alpha^{P_0+k}} - \\ &\qquad - \sum_{k=0}^{P_0-1} \left((\alpha\delta)^k - \delta^k\right) \frac{\mu_0 + (1-\mu_0)\alpha}{\mu_0 + (1-\mu_0)\alpha^k}. \end{aligned}$$

[52] Earlier, we introduced willingness to pay, $v(\mu,\rho)$, as a function of buyers' believes. For simplicity, if with some abuse of notation, we now write $v(P,N)$ as the willingness to pay induced by the beliefs following history (P,N).

Appendix: Data set descriptive statistics

We start with the description of the transactions data for the two types of coins, given in the first two columns of Table 12. 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. Finally, the modal length of the coin auctions was 7 days, ranging from 1 day to 10 days.

The third column of Table 12 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 cards, and 80% 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

Table 12: Summary statistics of transactions data.

	Am'an Eagle	Mint Set	Think-pad	Beanie Baby
Sale Price	50.8	77.8	578.6	11.1
	(10.7)	(21.6)	(413.6)	(4.3)
Highest Bid	50.0	75.8	529.5	10.7
	(14.0)	(25.7)	(429.9)	(3.8)
newvar Percent. items sold	.90	.84	.85	.52
	(.30)	(.37)	(.35)	(.49)
Minimum bid set by seller	20.0	38.3	104.7	9.8
	(23.1)	(38.2)	(260.7)	(5.0)
Number of Bidders	6.8	7.5	21.5	1.7
	(4.6)	(6.9)	(16.5)	(2.9)
Seller's eBay Rating	1596	1475	12442	2634
	(1639)	(2250)	(11628)	(4371)
Winning Bidder's eBay Rating	147	118	68	154
	(304)	(181)	(244)	(296)
Percent. sellers using PayPal	.54	.61	.78	.76
	(.49)	(.49)	(.41)	(.43)
Percentage sellers accepting credit cards	.79	.66	.83	.39
	(.41)	(.47)	(.38)	(.49)
Percentage listings with object photo	.41	.33	.81	.38
	(.49)	(.47)	(.39)	(.49)
Number of words in description	271	241	683	301
	(116)	(140)	(192)	(164)
Auction Length (days)	5.9	5.5	4.6	5.3
	(2.2)	(2.7)	(1.8)	(2.3)
Number observations	216	298	264	555
Number unique sellers	72	157	62	238
Market HHI	342	112	2756	195

Notes: 1. Standard errors in parentheses.

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.[53]

The last column of Table 12 provides a description of the Holiday Teddy Beanie Baby 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.[54] 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.

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 may 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.[55]

---

[53] 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.

[54] 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.

[55] 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.

References

Abbring, Japp H., Pierre-André Chiappori, and Jean Pinquet (2003), "Moral Hazard and Dynamic Insurance Data," forthcoming in the *Journal of the European Economic Association*.

Asch, Solomon (1946), "Forming impressions of personality," *Journal of Social and Abnormal Psychology* **41**, 258–290.

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

Bajari, Pat, and Ali Hortaçsu (2003), "Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay," *Rand Journal of Economics* **34**, 329–356.

Bajari, Pat, and Ali Hortaçsu (2004), "Economic Insight from Internet Auctions," forthcoming in *Journal of Economic Literature*.

Bar-Isaac, Heski (2002), "Reputation and Survival: Learning in a Dynamic Signalling Model," *Review of Economic Studies* **70**, 1–21.

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

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

Cohen, Adam (2002), *The Perfect Store: Inside eBay*, Little Brown and Company.

Datta, Saikat (1997), "Building Trust," London School of Economics, April.

Dellarocas, Chrysanthos (2002), "The Digitization of Word-of-Mouth: Promise and Challenges of Online Reputation Mechanisms," forthcoming in *Management Science*.

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

Dewan, S. and V. Hsu (2001), "Trust in Electronic Markets: Price Discovery in Generalist Versus Specialty Online Auctions," mimeo.

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

Eaton, D. H. (2002), "Valuing Information: Evidence from Guitar Auction on eBay," mimeo.

Ederington, L. H. and M. Dewally (2003), "A Comparison of Reputation, Certification, Warranties, and Information Disclosure as Remedies for Information Asymmetries: Lessons from the On-Line Comic Book Market," mimeo.

Ericson, Richard, and Ariel Pakes (1995), "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economics Studies* **62**, 53–82.

Friedman, Eric J., and Paul Resnick (2001), "The Social Cost of Cheap Pseudonyms," *Journal of Economics and Management Strategy* **10**, 173–199.

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

Gale, Douglas, and Robert W Rosenthal (1994), "Price and Quality Cycles for Experience Goods," *Rand Journal of Economics* **25**, 590–607.

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

Hoch, S., and Y. Ha (1986), "Consumer learning : advertising and the ambiguity of product experience," *Journal of Consumer Research* **13**, 221–233.

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

Hopenhayn, Hugo (1992), "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica* **60**, 1127–1150.

Houser, D. and J. Wooders (2003), "Reputation in Auctions: Theory and Evidence from eBay," University of Arizona Working Paper.

Jin, Gin Z., and A. Kato (2004), "Blind Trust Online: Experimental Evidence from Baseball Cards," University of Maryland Working Paper.

Jovanovic, Boyan (1982), "Selection and Evolution of Industry," *Econometrica* **50**, 649–670.

Kalyanam, K. and S. McIntyre (2001), "Returns to Reputation in Online Auction Markets," mimeo.

Klein, B, and K Leffler (1981), "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy* **89**, 615–641.

Kreps, David M., Paul Milgrom, John Roberts and Robert Wilson (1982), "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory* **27**, 245–252.

Kreps, David M., and Robert Wilson (1982), "Reputation and Imperfect Information," *Journal of Economic Theory* **27**, 253–279.

Levin, Dan, and James Smith (1994), "Equilibrium in Auctions with Entry," *American Economic Review* **84**, 585–599.

Livingston, J. (2002), "How Valuable is a Good Reputation? A Sample Selection Model of Internet Auctions," mimeo.

Lucking-Reiley, D., D. Bryan, N. Prasad, N., and D. Reeves (2000), "Pennies from eBay: the Determinants of Price in Online Auctions," University of Arizona Working Paper.

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

Mailath, George J, and Larry Samuelson (1998), "Your Reputation Is Who You're Not, Not Who You'd Like To Be," University of Pennsylvania and University of Wisconsin.

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

Melnik, M. I. and J. Alm (2002), "Does A Seller's eCommerce Reputation Matter? Evidence from eBay Auctions," *Journal of Industrial Economics* **50**, 337-350.

Milgrom, Paul, and John Roberts (1982), “Predation, Reputation, and Entry Deterrence,” *Journal of Economic Theory* **27**, 280–312.

Ockenfels, Axel, and Alvin E. Roth (2003), “Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction,” forthcoming in *Games and Economic Behavior.*

Phelan, Christopher (2001), “Public Trust and Government Betrayal,” Federal Reserve Bank of Minneapolis Staff Report No. 283, January.

Porter, Robert H (1983), “Optimal Cartel Trigger Price Strategies,” *Journal of Economic Theory* **29**, 313–338.

Resnick, Paul, and Richard Zeckhauser (2001), “Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System,” Harvard Kennedy School Working Paper.

Resnick, Paul, Richard Zeckhauser, J. Swanson, and K. Lockwood. (2003), “The Value of Reputation on eBay: A Controlled Experiment,” Harvard Kennedy School Working Paper.

Roth, Alvin E., and Axel Ockenfels (2002), “Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon Auctions on the Internet,” *American Economic Review* **92**, 1093–1103.

Shapiro, Carl (1983), “Premiums for High Quality Products as Rents to Reputation,” *Quarterly Journal of Economics* **98**, 659–680.

Snyder and Canto (1979), “Testing Hypotheses about other people: The use of historical knowledge,” *in Journal of Experimental Social Psychology* **15**, 330–342.

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

Tadelis, S. (2002), “The Market for Reputations as an Incentive Mechanism,” *Journal of Political Economy* **92**, 854–882.

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

Watson, Joel (1999), “Starting Small and Renegotiation,” *Journal of Economic Theory* **85**, 52–90.

Watson, Joel (2002), “Starting Small and Commitment,” *Games and Economic Behavior* **38**, 176–199.

Figure 2: Sample eBay Auction Listing

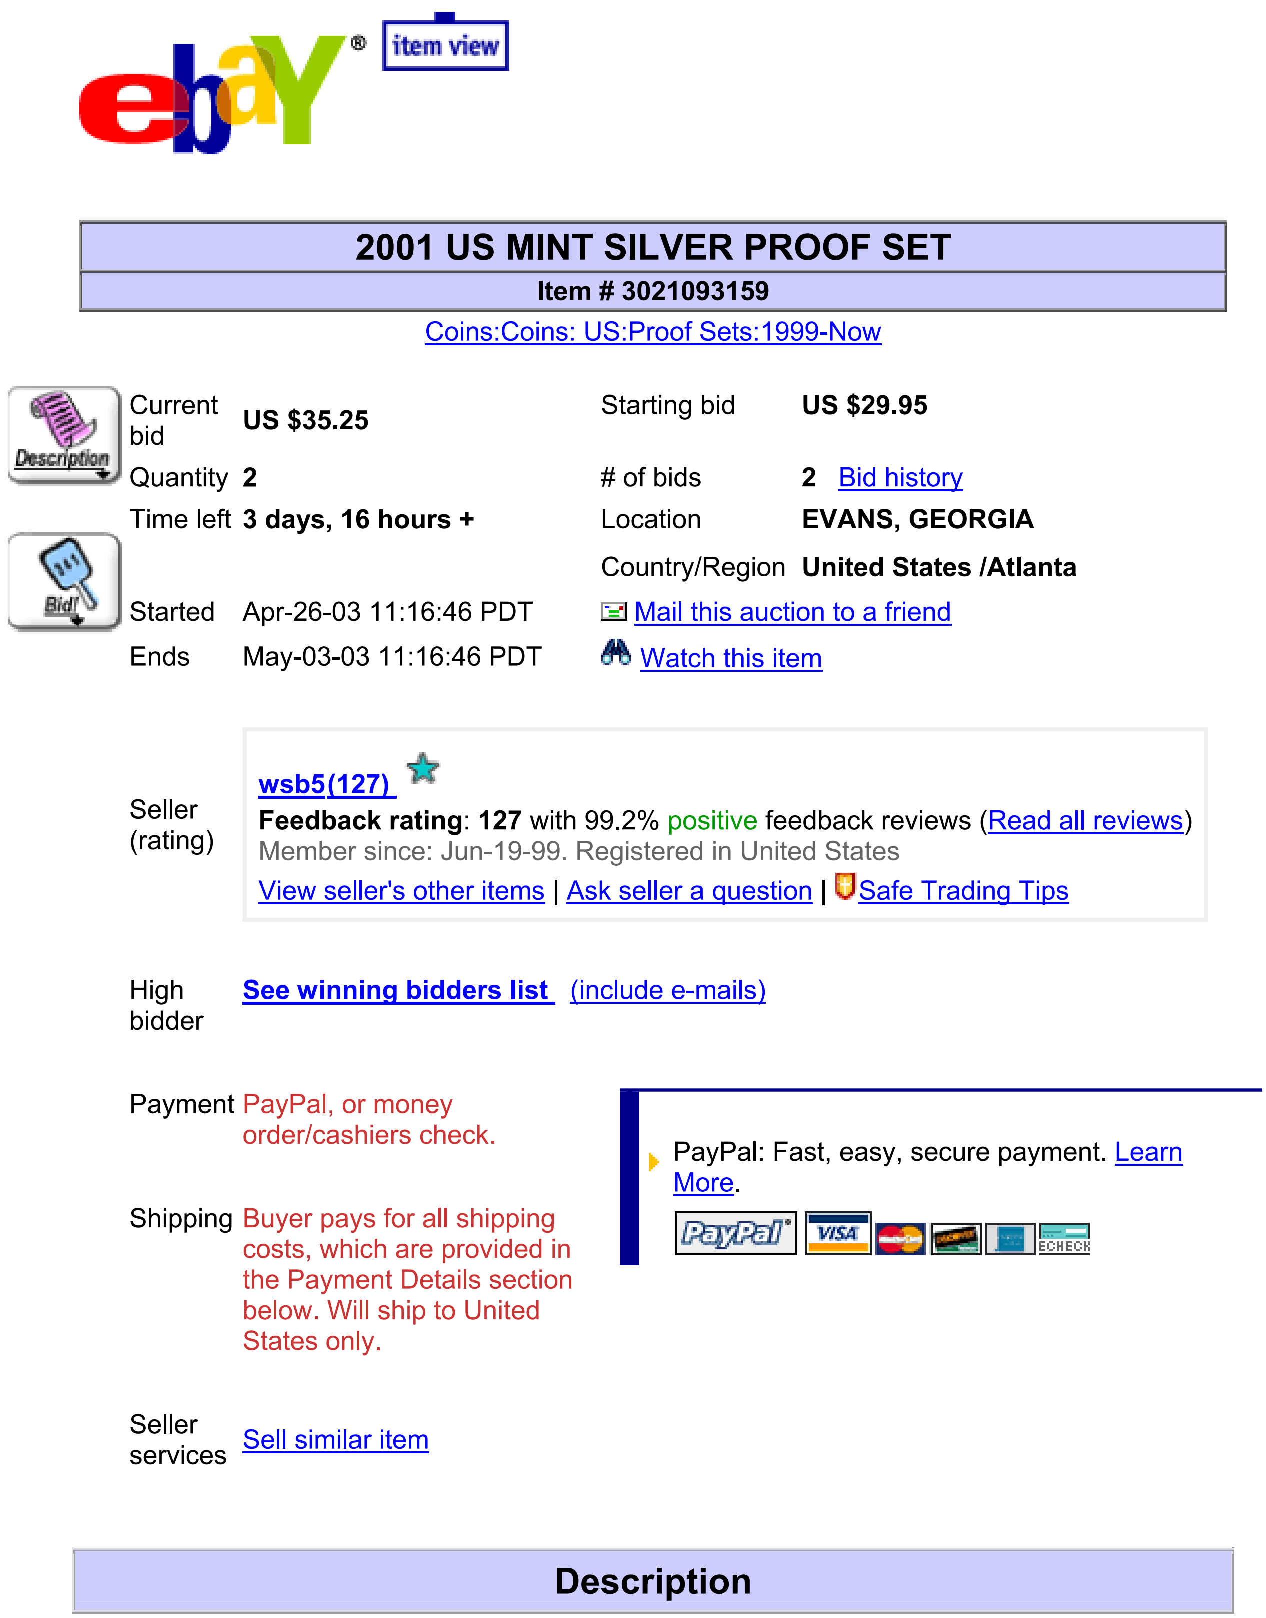

Set has 10 coins. Five state quarters and the penny,nickle,dime,half dollar and golden dollar.

Payment Details	Payment Instructions
United States Shipping and handling US $4.00 Additional shipping per item US $2.00 Shipping insurance per item (optional)US $1.30	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.

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 3: 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

Figure 4:How "bearsylvania" became a seller