Reputation and Adverse Selection: Theory and Evidence from eBay*

Maryam Saeedi

Carnegie Mellon University

June 13, 2018

Abstract

How can actors in a marketplace introduce mechanisms to overcome possible inefficiencies caused by adverse selection? In this paper, I use a unique dataset that follows sellers on eBay over time to show that reputation is a major determinant of variations in price. I develop a model of sellers' dynamic behavior where sellers have heterogeneous qualities unobservable by consumers. Using reputation as a signal of quality, I structurally estimate the model to uncover buyers' utility and sellers' costs and underlying qualities. The results show that removing the reputation mechanism increases low-quality sellers' market share, lowers prices, and consequently reduces sellers' profit by 66% and consumer surplus by 35%.

JEL Classification: D8, L2

*I am indebted to Patrick Bajari and Thomas Holmes for their valuable advice. I would also like to thank V.V. Chari, Alessandro Dovis, Konstantin Golyaev, Hugo Hopenhayn, Larry Jones, Patrick Kehoe, Kiyoo-il Kim, Tina Marsh, Ellen Mcgrattan, Minjung Park, Amil Petrin, Erick Sager, Jack Shen, Ali Shourideh, Neel Sundaresan, Steve Tadelis, Robert Town, Joel Waldfogel, Thomas Youle, and Ariel Zetlin-Jones, as well as the participants of the Applied Micro Workshop and Applied Micro Seminar of the University of Minnesota, IIOC, the Ohio State University, Carnegie Melon University, Arizona State University, State University of New York at Stony Brook, UCLA, Charles River Associates, Revolution Analytics Inc., eBay Research Labs' Reading Group, Universidad Catolica, Chile, Said Business School, Oxford University. Any remaining errors are my own. Correspondence: msaeedi@andrew.cmu.edu

1 Introduction

Asymmetric information is known to lead to inefficiencies in markets. In particular, in online markets where trading is decentralized and with little repeat interactions, asymmetric information problems can be magnified.[1] In these markets, the sellers are better informed about the characteristics of the items they sell or their level of expertise than are their costumers. The sellers on eBay, for example, can misrepresent the objects they sell or mishandle the shipping of the items sold. Reputation mechanisms are often used to mitigate these asymmetric information problems. However, while the positive role of reputation in overcoming asymmetric information problems is known, its quantitative effect on marketplaces is still not well understood.[2] To quantify the effect of reputation mechanisms, in this paper, I develop and estimate a dynamic model for the eBay reputation mechanism. My quantitative analysis highlights the main channel through which reputation affects eBay. In particular, I show that the equilibrium effects of removing the reputation mechanism will greatly outweigh the static effect of reputation in terms of the price premium reputable sellers receive.

The eBay marketplace is plagued by information asymmetries, and the problem is partially alleviated by the eBay reputation system.[3] Moreover, as Bar-Isaac and Tadelis [2008] mention, eBay provides a very good environment for economists to study the effects of reputation on sellers' actions and profits. First, economists can observe all of the sellers' characteristics observable by buyers. Second, sellers and buyers have little to no interactions with each other outside the eBay website; thus, buyers do not have additional information, unobservable to a researcher, about

[1] Examples of these markets include eBay, Amazon Marketplace, Alibaba, and Taobao in retail; Airbnb and VRBO in room and house sharing; Uber and Lyft in transportation; Care.com in child care; Rover in pet care; and Upwork (former oDesk) in the freelance and labor market, as studied in Cai et al. [2013], Fan et al. [2016], Zervas et al. [2015] and Filippas et al. [2017]. However, asymmetric information problems are prevalent in offline marketplaces as well. They have been shown to exist in insurance markets, Finkelstein and McGarry [2006], credit markets, Crawford et al. [2015], and financial markets, Ivashina [2009], among many others.

[2] For example, see Holmstrom [1999], Mailath and Samuelson [2001], Board and Meyer-ter Vehn [2010], and Board and Meyer-ter Vehn [2011].

[3] See, for example, Resnick et al. [2006], Brown and Morgan [2006], Lucking-Reiley et al. [2007], Kollock [1999], and Yamagishi and Matsuda [2002]. Bajari and Hortaçsu [2004] surveys the literature on eBay's feedback system.

sellers. Third, economists can track sellers over time, which gives them extra information about sellers that is unobservable to buyers. This information can then be used to estimate the underlying model parameters.

To quantify the value of reputation, I examine sellers on eBay and use a unique dataset that follows them over time. First, I analyze the determinants of price variation in a set of homogeneous goods (iPods). Second, I develop and estimate a model of sellers' behavior over time, in which sellers have heterogeneous unobserved qualities and build up their reputation by selling objects and acquiring the eBay registered-store status (eBay store from now on), Powerseller status, or both. Finally, using the estimated model, I perform counterfactuals to analyze the effect of reputation on profits and market outcomes, as well as possible alternative methods to overcome adverse selection. The counterfactuals highlight the dynamic role of reputation in reducing the market share of high-quality sellers. Specifically, absent any reputation mechanism, consumer surplus declines by 35%, most of which is due to a change in market structure as market share of high-quality sellers decreases and market share of low-quality sellers increases.

In order to empirically analyze the role of reputation, I examine the data on sellers of iPods between 2008 and 2009. The dataset follows sellers on eBay and includes the number of items sold, the final price of items sold, and items' and sellers' characteristics. Consistent with other studies on eBay, considerable variation exists in the prices of iPods sold. In this context, there are two main variables of interest related to reputation: *Powerseller status* and *eBay store status*. Using these two variables as proxies for reputation, I show that reputation has a significant role in explaining price variations. In particular, prices of new iPods are positively correlated with reputation. Among sellers of new iPods, being a Powerseller, *ceteris paribus*, increases prices by approximately 3%, while being an eBay store, *ceteris paribus*, increases prices by approximately 4%.

While this empirical evidence is suggestive of the value of reputation, it ignores its effect on sellers' incentives to build their reputation over time (i.e., achieve Powerseller or eBay store status). In particular, a high-quality seller without Powerseller or eBay store status might sell a higher quantity of goods than a low-quality seller in order to become a Powerseller or eBay store which results in a change in market share of sellers of different quality. To quantify this effect, I develop a dynamic structural model of sellers' behavior where sellers are privately informed about their quality and acquire reputation over time.

Formally, the model consists of two sets of agents: buyers and sellers. Buyers are short-lived, derive utility from purchased goods, and do not observe the quality of the objects bought by previous buyers. Sellers are long-lived, can sell different quantities of goods over time, and are heterogeneous in the *quality* of the goods they are selling. The higher the quality of a good, the higher the buyer's utility from purchasing one unit of that good. Sellers quality has a persistent component and transitory independent and identically distributed (i.i.d.) component for each time period. To capture adverse selection, I assume that qualities are privately known to sellers; therefore, buyers do not observe the quality of an object.

In this environment, I introduce eBay's reputation system: eBay store and Powerseller statuses. High-quality sellers can choose to pay a monthly fee to become an eBay store, and they must fulfill two requirements to become a Powerseller: sell more than the quantity threshold set by eBay and have a quality higher than the quality threshold. These statuses are valued by sellers, because buyers use them as signals to distinguish high-quality sellers from low-quality ones. Since buyers prefer high-quality sellers to low-quality sellers and are willing to pay a higher price for items sold by high-quality sellers, sellers are willing to pay the monthly fee or sell more than the static optimum in order to achieve these status.

This model economy constitutes a dynamic game between many sellers. Specifically, my dataset contains 769 sellers whose decisions depend on their quantity choices in the past. Due to this high number of sellers and the large state space at the individual level, the standard methods of estimation of dynamics games, namely the estimation of Markov perfect equilibrium by Pakes and McGuire [2001] is impossible – due to the "curse of dimensionality" as discussed by Pakes and McGuire [2001]. To overcome this problem and achieve tractability, I use the *oblivious equilibrium* concept introduced by Weintraub et al. [2008]. Under this equilibrium concept, the state of the game is given by market aggregates. Roughly speaking, the assumption is that sellers are small and, as a result, their actions do not affect the state of the industry. This means that any given seller does not need to take into account other sellers' response to her decision during this period or the next, and the seller does not need to keep track of the history of her opponents' actions. This greatly reduces the dimensionality of the state space that one needs to keep track of.[4] Given that the dataset has a large number of sellers and the largest seller accounts for only 5% of the market, this equilibrium concept is reasonable and closely approximates Markov perfect equilibrium, as discussed by Weintraub et al. [2006].

I estimate the model in two steps. First, I identify the stochastic process for qualities by using an important theoretical insight from the model. In particular, I show that in the model higher-quality sellers (as measured by their persistent level of quality) must always sell higher quantities. This relationship allows the parametric estimation of qualities using the quantity choices of sellers over time. In the next step, using the approach of Bajari et al. [2007], I identify sellers' cost parameters. This procedure is based on the assumption that the observed data are the outcome of sellers' maximization problem based on their actual cost function; therefore, revealed profit conditions can be used to back out cost parameters.

Using the above model, I perform two counterfactuals to estimate the added value of having a reputation system and the effect of substituting it with a warranty mechanism. In the first counterfactual, I remove eBay's reputation system altogether and solve the equilibrium in which sellers have no means of signaling their quality to buyers. Absent any reputation mechanism,

[4] To some extent, oblivious equilibrium is similar to monopolistic competition when firms are infinitesimal and their decisions change their price and profits but do not affect the market as a whole.

the model becomes a static model with adverse selection. As a result, higher-quality sellers face lower prices and therefore lower market shares. In contrast, the market share of low-quality sellers increases. This decline in average quality in the market leads to further decrease in prices and consequently a shrinkage of the market and its unraveling. Specifically, removing the reputation system decreases buyers' surplus by 35%, total sellers' profit by 66%, and eBay's profit by 38%. As demonstrated by this discussion, these changes in total welfare are well beyond the observed premium for high-quality sellers and are driven by the fact that higher-quality sellers cannot build reputation by increasing their sales. In presence of the reputation mechanism, even among non-Powersellers, the higher-quality sellers have incentives to sell more items given that they are more likely to become a Powerseller in the future. These large effects emphasize the importance of a reputation mechanism particularly for online marketplaces, where buyers and sellers usually do not meet before making transactions.

Finally, I estimate the effect of substituting the reputation system with a warranty mechanism. I assume sellers can provide a warranty, promising that the quality of the sold item is higher than a threshold set by the marketplace. In this case, high-quality sellers' profit increases, and so does total sellers' profit. Depending on the level of quality threshold, buyers' welfare and eBay's profit can increase or decrease compared to the no-signaling case. Additionally, at certain quality thresholds, the surplus values near those in the marketplace with a reputation mechanism.[5]

Related Literature. This paper contributes to two lines of literature: theoretical papers on reputation and empirical work on reputation systems, both of which are summarized in Bar-Isaac and Tadelis [2008]. Although many papers have explored these topics, to the best of my knowledge, this paper is the first to empirically estimate the role of reputation based on a dynamic model of seller behavior.

[5]In 2010, eBay added a site-wide warranty mechanism called eBay Buyer Protection, which mandates that sellers refund buyers if the sold items are not as described or if the items are not received. The effect of adding this policy and its interaction with reputation is studied in Hui et al. [2015] using regression discontinuity methods. My coauthors and I estimate that adding warranty increases the total welfare by 2.9%.

In this paper, reputation can help mitigate adverse selection problems, as discussed in the literature on modeling reputation as beliefs about behavioral types (Milgrom and Roberts [1982], Kreps and Wilson [1982], Holmstrom [1999], and Mailath and Samuelson [2001] to name a few).[6] The closest paper is perhaps Mailath and Samuelson [2001], where a seller has private information about the product she sells and can exert effort to increase its quality. Market participants can learn her type from observing signals of quality. However, my paper departs from this line of research in that I abstract from learning. Reputation in the model developed here is the mechanism introduced by the marketplace (in this case eBay) that can help signal sellers' private types.

Another literature line has empirically examined the role of the reputation mechanism, including on eBay (see Bajari and Hortaçsu [2004] and Dellarocas [2005] for a summary of this line of literature). Examples of major empirical work in this area are Resnick and Zeckhauser [2002], Melnik and Alm [2002], Houser et al. [2006], Resnick et al. [2006], Masclet and Pénard [2008], and Reiley et al. [2007]. These papers have studied the role of the feedback system on eBay, finding a positive correlation between the price of an item and the feedback that a seller has received. Cabral and Hortacsu [2010] empirically studied the feedback system in a dynamic setup, and they found that the first negative feedback increases sellers' exit rate, but consecutive negative feedback ratings do not have large effects on sellers' performance and exit rates. I build on these studies by providing evidence of the effect of Powerseller status, eBay store status, or both on sellers' revenues and profits and by structurally estimating the value of reputation using a dynamic model of reputation.

Reputation and adverse selection play key roles in my empirical and theoretical analyses. A few studies have pointed out the significance of the adverse selection problem on eBay. Using a novel approach, Yin [2003] showed that the final price of an object is negatively correlated with the dispersion in the perceived value of that object. This observation implies that the higher the

[6] Many researchers have applied the techniques introduced in this literature to address applied questions. See Chari et al. [2014], Board and Meyer-ter Vehn [2010], and Board and Meyer-ter Vehn [2011], for example.

dispersion in the perceived value, the higher the discount at which buyers are willing to purchase the good. This points to the existence of information asymmetries and their negative effects on the final price of an item. Lewis [2011], however, showed that by selectively revealing information, sellers decrease the dispersion of the perceived value of an object, thereby increasing its final price. He considered the number of photos and the amount of text a seller provides for an object sold to be the main source of revealing information, and found that the final price increases with the number of photos and the amount of text uploaded on the auction page.

The rest of the paper is organized as follows: The dataset analyzed and an overview of the market structure on eBay are presented in Section 2. Then, Section 3 presents the dynamic model of sellers' behavior and their interactions with buyers through eBay; Section 4 describes the identification procedure for the deep parameters of the model; and Section 5 describes the estimation of the model. Finally, Section 6 presents two counterfactual exercises to estimate the value of reputation, and Section 7 concludes the paper.

2 Data

The dataset consists of all eBay transactions involving iPods over eight months in 2008 and 2009. I have chosen iPods because they are homogeneous goods, and I can precisely control for their characteristics. In addition, there were few or no promotions for them outside eBay at the time. I collected the data using a spider program.[7] For each transaction, I have data on items' characteristics, sellers' characteristics, and listing format. iPods come in different models and generations (Figure 1). The newest model at the time of study was iPod Touch, introduced in 2007, and the oldest was iPod Classic, introduced in 2001.

eBay is one of the biggest e-commerce websites, and the largest auction website, where users can sell or buy a wide variety of items. In early 2008, eBay counted hundreds of millions of registered

[7] This program, written in Python, searched for completed iPod listings and saved the information contained on the eBay website into a file. The program ran frequently to collect new data points for a period of 8 months.

users, more than 15,000 employees, and revenues of almost $7.7 billion.[8] On eBay sellers can sell an item either through an auction or by setting a fixed price for their item, an option called "Buy it Now". The auction mechanism is similar to a second-price auction: A seller sets the starting bid of an auction, and bidders can bid for the item. Each bidder can observe all previous bids, except for the current highest bid. A bidder needs to bid an amount higher than the current second highest bid plus a minimum increment.[9] If this value is higher than the current highest bid, the bidder becomes the new highest bidder; otherwise, the bidder becomes the second highest bidder. The winner has to pay the second highest bid plus the increment or his or her own bid, whichever is smaller. Auctions last for 3 to 10 days and have a predetermined ending time that cannot be changed once the auction is active. While the eBay auction mechanism is of significant importance and has been the subject of an extensive literature, I abstract from modeling it explicitly. Instead, I assume that sellers choose the number of items to sell and prices are determined by a demand function derived from a standard discrete choice model.

Figures 2 and 3 show a snapshot of a finished auction page and bid history for the same item, respectively. At the top of the page, there is information about the object and its bid history, and on the top right side of the page, information about the seller. The rest of the page contains detailed information about the object sold in the auction. Bidders have access to the bid history page, which shows previous bidders' short-form IDs, their bids, and the time they submitted their bid.[10]

2.1 The eBay Reputation Mechanism

The reputation mechanism on eBay consists of various signals: *seller feedback number*, *seller feedback percentage*, *detailed seller ratings*, *eBay store*, and *Powerseller*. After each transaction, buyers

[8] source: http://en.wikipedia.org/wiki/eBay

[9] The increment is between 2% and 5 % of the current price and is set by eBay.

[10] eBay stopped showing the complete ID of bidders in 2007, mentioning the following reasons: to keep the eBay community safe, enhance bidder privacy, and protect eBay's members from fraudulent emails.

can leave a positive, negative, or neutral feedback rating accompanied by a textual comment. A summary of feedback history is available on the auction page: *seller feedback number* and *seller feedback percentage*. The feedback number is the difference between the number of positive and negative feedback ratings received by a seller. The feedback percentage is the percentage of positive feedback ratings from all positive or negative feedback ratings received by a seller. Since 2007, buyers can additionally rate sellers according to four criteria: item as described, communication, shipping time, and shipping and handling charges. This extra information, known as *detailed seller ratings*, is not shown on the auction page, but is accessible through sellers' web page by clicking on the sellers' ID on the listing page.

There are two additional reputation signals available on the listing page: *eBay store* and *Powerseller*. In order to be eligible for becoming an *eBay store*, sellers must follow eBay's policies and have a high seller standard rating.[11] An eBay store benefits from discounted listing fees, but it must pay a fixed monthly fee to eBay. Sellers acquire the *Powerseller* badge if they meet various quality and quantity thresholds.[12] Enrollment into this program is free of charge for qualified sellers.

To make a simplifying assumption, I control only for Powerseller and eBay registered store status as signals of sellers' quality. I do not include information that is not presented on the listing page, detailed seller ratings and textual feedback ratings, because, as reported by Nosko and Tadelis [2014], who used the click-through data of eBay users, less than 0.1% of buyers visit sellers' history page. I also do not include the feedback number and feedback percentage ratings, because in my dataset, the standard deviation of the feedback percentage is very low. Most sellers have a feedback percentage higher than 99%, and the median of the feedback percentage for sellers is

[11] The seller standard rating includes various requirements, such as few open disputes, few low detailed seller ratings, and no outstanding balance.

[12] The requirements for becoming a Powerseller are as follows.
Three-month requirement: a minimum of $1,000 in sales or 100 items per month, for three consecutive months.
Annual requirement: a minimum of $12,000 or 1,200 items for the prior twelve months.
Achieve an overall feedback rating of 100, of which 98% or more is positive.
Account in good financial standing.
Following eBay rules.

100%.[13] Moreover, the effect of these two variables on price is very small, as shown in more detail in the next section. In addition, Table 9 in the Appendix shows that there is a very small effect for the feedback percentage and a negative sign for the number of feedback ratings. This finding might be due to the quality and quantity requirements for obtaining Powerseller status. Therefore, by controlling for this status, I do not find a big effect from the other two signals.

The variables used in this paper are summarized in Table 1, which shows that 36% of the listings in my dataset are sold by eBay registered stores and 48% are sold by Powersellers. I further focus on sellers with at least 10 sales since sellers with very few transactions may not be optimizing. The total number of sellers is thus reduced from 46552 to 769, which makes the problem more tractable; in addition, the total number of transactions is reduced to 84301, about half of the original dataset.

Most of the items (92%) were sold using an auction method. Therefore, in the model section, I assume that sellers are setting the quantity, and the price is determined in the market. In an eBay auction setting, sellers can set a secret reserve value; if the final bid is lower than this value, the trade will not occur. However, only 4% of listings have used this option; thus, I do not model it as an option for sellers.

I have also collected a set of items' characteristics, such as the condition of the item (new, refurbished, or used), the level of internal memory, and the make of iPod. Most iPods sold on eBay are used items, with 25% of listings being new and 19% being refurbished. One would expect a higher effect of seller reputation for used items, given the higher sources of asymmetric information (the battery may not be working, the screen may be scratched, or the touchpad screens may not work properly, and so on), and in Appendix B, I show that this is in fact true.

2.2 Reputation and Price

eBay store and Powerseller statuses signal sellers' quality, showing that sellers are following eBay's rules closely and have a good track record. As the data in Table 2 indicate, the final prices of

[13] This is also reported by Nosko and Tadelis [2014].

items sold are higher when sellers are Powersellers or eBay stores. The first column of Table 2 shows the average price of all iPods. Having eBay store status or Powerseller status increases the average final price for iPods sold by these sellers. This increase in price may be due to a selection bias—Powersellers or eBay stores may sell higher-value items compared to other sellers—which can be partly accounted for by controlling for the item characteristics. Specifically, I use the regression specification (which I later use to estimate the buyers' demand) to show the fitted values for an iPod with the average characteristics observed in the dataset. The average fitted prices are shown in the third and fourth columns of Table 2; the average price goes up for Powersellers or eBay stores by 15% and 5%, respectively.

The last two rows of Table 2 illustrate the effect that moving from the lowest 25 percentile to the highest 25 percentile of feedback ratings has on price; when controlling for item characteristics, the effect on price is positive but very small, and sellers with higher positive feedback receive a slightly higher price. This finding is consistent with findings in Cabral and Hortacsu [2010].[14]

Another reason for price differences can be unobservable heterogeneity between sellers or items sold. However, the number of sellers with status change is not large enough to exploit the price variation as sellers change their status.[15] Table 8 shows the results I obtain when I control for more observable differences between sellers such as the amount of text or the number of pictures provided in a listing. Adding these controls does not reduce the effect of reputation signals. Further, the regression for "New iPod Nano" still results in positive effects of the reputational badges.

[14] Cabral and Hortacsu [2010] further studied the dynamic effect of negative feedback on sellers' actions. They found that the most important effect of negative feedback is preventing the lowest-quality sellers from participating, as the first instance of negative feedback substantially increases the probability of exit for sellers. However, the authors did not find a big effect of negative feedback on prices.

[15] In a more recent paper, Hui et al. [2015], my co-authors and I investigated the effect of Top-rated seller status by studying sellers whose status changed. The eBay Top-rated status replaced Powerseller status in 2009. In this paper, we used various studies to show that the increase in price as a result of obtaining the reputation badge is due to signaling rather than learning or other omitted variables.

3 Model

To capture the dynamic effect of reputation, I develop a simplified dynamic model of reputation based on Mailath and Samuelson [2001]. There are two major players in this market: buyers and sellers, and eBay is the informed market designer that sets the rules for signaling mechanisms.[16] Sellers have heterogeneous qualities, which are unobservable to buyers but observable to eBay, while buyers care about the sellers' quality but observe only the signals provided by eBay. Hence, the signaling mechanism helps buyers distinguish high-quality sellers from low-quality sellers and helps higher-quality sellers to obtain a higher profit.

3.1 eBay

eBay sets different mechanisms for sellers to signal their quality. I assume that eBay observes the quality of sellers based on the history of sellers in the market, which is not explicitly modeled. The two mechanisms in this paper are similar to those of Powerseller status and store status used on eBay during the data collection period. Sellers who sell more than Q^p, a threshold set by eBay, for three consecutive periods and have a quality level r_{jt} higher than μ^p, another threshold set by eBay, are signaled as Powersellers. A seller does not pay any fixed or monthly fee to be included in this program. Sellers have to maintain these quality and quantity thresholds every month to keep their Powerseller badge; otherwise, they lose their badge. Similarly, sellers who have quality r_{jt} higher than the threshold set by eBay, μ^s, can register their account as an eBay store, and they have to pay a monthly fee to eBay, c^s, to participate in this program.

3.2 Buyers

There is a unit measure of buyers. Buyers are short-lived and cannot track sellers over time. Each period, a buyer decides to buy one of the items offered by one of the sellers or to buy the outside

[16] In this paper, I am not solving for the optimal reputation system and eBay is not a player. I am modeling the signaling to be similar to the rules in effect on eBay.com during the period data were collected.

good. In this case, the outside good is assumed to be an MP3 player of another brand available on eBay.[17] To simplify the choice of buyers, the auction framework of eBay is not explicitly modeled, and it is assumed that buyers choose the option that gives them the highest utility level. Buyers care about the quality of sellers but do not observe it directly.[18] Instead, they observe two signals that are correlated with sellers' quality: Powerseller status and store status.

Buyer i gets random utility u_{ijt} from purchasing an iPod from seller j at time period t:

$$u_{ijt} = -\alpha p_{jt} + \beta_r r_{jt} + \xi_t + \xi_{jt} + \epsilon_{ijt},$$

where price p_{jt} is the price of the item sold by seller j at time period t; quality r_{jt} is the quality of seller j at time period t, which is unobservable to buyers; seasonal shock to utility is denoted by ξ_t; and the seller-time unobservable quality is denoted by ξ_{jt}. The random variable ϵ_{ijt} is the unobservable utility random variable with the logit distribution. As it is common in settings with the logit preferences and as stated in Berry [1994], the demand function faced by each seller depends on their quantity, status, total quantity in the market, and total quantity of the outside good. For now, I refer to this demand function as $p\left(q; \phi_p, \phi_s, \Omega\right)$, where Ω is the vector of market characteristics that affect demand; that is, total sales and the total quantity of the outside good. In Section 4.1, I will discuss the determination of this demand function in detail.

3.3 Sellers

There are N sellers. Sellers are born with a persistent level of quality, η_j. At the beginning of each period, they get a shock to their quality, γ_{jt}, which is i.i.d. distributed with distribution G.[19] This

[17] This assumption helps me to control for seasonality. Another sensible outside good is buying an iPod directly from Apple; however, I could not get access to monthly sales of iPods through Apple.

[18] Sellers' quality represents information about them that is important for buyers, such as sellers' accuracy of the item description, shipping speed, or their communication skills.

[19] In this paper, I abstract from modeling moral hazard; I assume that sellers' quality is determined exogenously and sellers cannot select their quality directly by making an effort or investing in quality; this assumption has been formulated in prior theory papers (e.g., Holmstrom [1999], Mailath and Samuelson [2001], Board and Meyer-ter Vehn [2010], and Board and Meyer-ter Vehn [2011]). The main reason for this simplifying assumption is lack of data. I do not directly observe sellers' intrinsic quality nor their level of effort to increase their quality.

will result in sellers' quality at period t to be

$$r_{jt} = \eta_j + \gamma_{jt}.$$

The timing of sellers' decision is shown in Figure 4. After knowing their level of quality, sellers can deduce their Powerseller status, ϕ^p, based on their past quantity and their quality level using the following rule:

$$\phi^p_{jt} = 1 \Leftrightarrow \begin{cases} q_{jt-1}, q_{jt-2}, q_{jt-3} > Q^p \\ r_{jt} > \mu^p. \end{cases} \tag{1}$$

Next, they choose the number of items they intend to sell, $\tilde{q}_{jt}$. However, the actual number of iPods they can obtain may be different from this number, and it is drawn from a negative binomial distribution with mean $\tilde{q}_{jt}$ and the dispersion parameter κ, denoted by $dF(q_{jt}|\tilde{q}_{jt})$. This randomness captures the variation in quantities sold even after controlling for various observables. It can be interpreted as an inventory shock as well as other sources of uncertainty in sellers' decision. For example, sellers might not have a secure supply line of iPods but can put in effort in order to find iPods to sell. After the realization of quantity, they need to decide about their store status, ϕ^j_t; they can choose to become a registered store and pay the monthly fee of c^s only if $r_{jt} > \mu^s$; otherwise, their store status will be set to zero.

Sellers' profit function at time t is

$$\pi(q_{jt}, \phi^p_{jt}, \phi^s_{jt}, \Omega_t) = p(q_{jt}, \phi^p_{jt}, \phi^s_{jt}, \Omega_t)q_{jt} - cq_{jt} - c^s\phi^s_t,$$

where c is the marginal cost of acquiring an iPod for sellers. Given that most transactions at this time were done through the auction mechanism and not the fixed price mechanism, Buy it Now, I assume that the sellers choose the number of items to sell and the price is the outcome of the buyers' problem, as discussed in detail in Section 4.1.

Sellers' problem can be formulated as follows:

Given $\mathbf{q}_{-} = \{q_{jt-1}, q_{jt-2}, q_{jt-3}\}$, their persistent level of quality, the shock to their quality, and market characteristics, sellers choose the intended number of items to sell and store status to maximize

$$V(\eta_j, \gamma, \mathbf{q}_{-}, \Omega) = \max_{\tilde{q}_j, \phi_j^s} \int \left(\pi(q_j, \phi_j^p, \phi_j^s, \Omega) + \beta \int V(\eta_j, \gamma', \mathbf{q}'_{-}, \Omega') g(\gamma) d\gamma \right) dF(q_j|\tilde{q}_j) \tag{2}$$

subject to

$$\begin{aligned} \mathbf{q}'_{-} &= (q_j, q_{j,-1}, q_{j,-2}) \\ \phi_j^s &= 0 \quad \text{if} \quad \eta_j + \gamma < \mu^s \\ \phi_j^p &= 1 \quad \text{if} \quad \begin{cases} q_{jt-1}, q_{jt-2}, q_{jt-3} > Q^p \\ \eta_j + \gamma > \mu^p. \end{cases} \end{aligned} \tag{3}$$

Let $q^*(\eta, \gamma, \mathbf{q}'_{-}, \Omega)$ be a non-negative real number, sellers' status choice $\phi^{s*}(\eta, \gamma, \mathbf{q}'_{-}, \Omega)$ be a binary function that solves the above problem, and β be sellers' discount factor. Here, I assume that the market characteristics follow a deterministic path over time. I will explain the needed assumptions for this to hold in the following section when I describe in detail the oblivious equilibrium concept used in this paper.

It is important to emphasize a major simplifying assumption regarding the type of iPods sold. I assume that a seller selects the number of items to sell but not the actual type of iPod the seller will end up selling. As mentioned in the data section, there are various models and generations of iPods, and they have different marginal costs to obtain. There can be strategic choices of type of iPod to sell, as the sellers may be able to charge different markups for different types of iPods, but I abstract from modeling them in this paper. Instead, to simplify, I assume that an item with the average characteristics of iPods in my dataset is traded, and sellers do not choose the types of iPods to sell.[20] The data on prices will give me the appropriate value of each characteristic, and I will find the marginal cost for this averaged product, simply called iPod from now on.[21]

Finally, there is no entry into this economy after period 0, and there is no permanent exit from the market, either. Sellers can decide to sell no items in one period, which can be interpreted as exiting the market; however, they can return to the market without paying a fee in the next period.

3.4 Oblivious Equilibrium

The above environment describes a dynamic game wherein sellers act dynamically and strategically. In other words, a seller's past choices of quantity affect that seller's current quantity choice through their effect on the seller's current Powerseller or store status. In addition, other sellers' current and past quantity choices affect a given seller's quantity choice. The literature pioneered by Pakes and McGuire [2001] often uses Markov perfect equilibrium (MPE) to analyze players' behavior in such dynamic games. However, in my setting with a large number of sellers (769), the state variable for the game becomes items sold by each seller as well as their reputation status in the past two periods—a vector of dimension 3076! This is the curse of dimensionality, which is discussed in depth by Pakes and McGuire [2001].

In order to curb this problem, I use the oblivious equilibrium as developed by Weintraub et al. [2008]. Under this equilibrium concept, it is assumed that aggregate sales in the market and the size of the outside market are the state variables; that is, all histories of the game that lead to the same level of aggregate sales and outside market size must lead to the same strategies. This is in contrast with MPE in which variations in the past individual states affect equilibrium strategies beyond their

[20] This iPod would be 20% iPod Nano, 30% iPod Classic, 20% new, 10% refurbished, and so on.

[21] This assumption may not be so restrictive as long as sellers do not have access to a very strong upstream source of inventory where they can actually select the type of items to sell on eBay. Given that I observe most sellers do sell a variety of items, different models, different conditions, and so on, I expect this to be the case.

effect on average industry variables. This equilibrium concept is particularly applicable to markets that are comprised of many small sellers. Note that with many smaller sellers, an approximate law of large numbers implies that the distributions of individual states are somewhat invariant, which would then mean that industry average becomes the relevant state variable in the game. In my sample, there are 769 sellers, the biggest seller in the market has around 5% of the market in one month, and most sellers' market share is less than 1%; therefore, oblivious equilibrium can be a good approximation.

The application of the oblivious equilibrium has two key implications: first, the only relevant endogenous state variable is market size (in addition to the size of the outside market); second, each seller is small and, as a result, a seller's choice does not affect the state of the industry. In other words, sellers are *oblivious* to their knowledge of other sellers' state.[22] Finally, as Weintraub et al. [2010] have shown, this equilibrium concept can be applied to a non-stationary industry, which is required in my setting due to the seasonality in demand.

Given this discussion, I can proceed with the formal definition of the equilibrium. As mentioned, the relevant state variables are the total quantity of items sold by sellers and size of outside good. Given the non-stationarity of the data, mainly arising from seasonality, I assume that all players have perfect foresight and thus have a common perceived path for the aggregate state, $\{\Omega_t\}_{t\geq 1}$, where $\Omega_t = \{Q_t, q_{0t}\}$. Given perfect foresight, this perceived path coincides with its actual realizations. As a slight abuse of notation, I use $\Omega = \{\Omega_t\}_{t\geq 1}$ for the perceived and actual path of aggregate states.

A non-stationary oblivious equilibrium is defined as a set of policy functions, $q^*(q_-, \eta, \gamma, \Omega)$ and $\phi_s^*(q, q_-, \eta, \gamma, \Omega)$, buyers' beliefs, market characteristics, Ω, and pricing functions, $p(q; \Omega, \phi_s, \phi_p)$, such that

- given policy functions, characteristics of sellers, and buyers' beliefs, $p(q; \Omega, \phi_s, \phi_p)$ is the out-

[22]Additionally, Weintraub et al. [2006] have shown that as the number of sellers grows, oblivious equilibrium converges to Markov perfect equilibrium using the framework of Ericson and Pakes [1995].

come of buyers' utility optimization;

- $q^*(q_-, \eta, \gamma, \{\Omega_t\})$ and $\phi^*_s(q, q_-, \eta, \gamma, \{\Omega_t\})$ are maximizing sellers' value function given $\{q_-, \eta, \gamma\}$ and $\{\Omega_t\}$;
- buyers' beliefs are consistent with sellers' behavior; and
- total market quantities, Q_t, are consistent with individual sellers' choices — that is, the sum of expected quantities across sellers is equal to Q_t.

3.5 Analysis of Quantity Choice

Before proceeding with the estimation of the model, I present a theoretical result that will help in the identification procedure. Specifically, I show that in any equilibrium, the quantity choice of a seller is increasing in that seller's quality. This result is used to identify the sellers' persistent level of quality, η by associating it to sellers' fixed effect in a quantity regression described in section 5.2. The following proposition summarizes this result:

Proposition 1. *Consider any equilibrium with an associated path of market aggregates $\mathbf{\Omega}$. Suppose that $p(q; \phi^s, \phi^p, \Omega)$ is increasing in ϕ^s and ϕ^p and is supermodular in (q, ϕ^s, ϕ^p). Then, the choice of optimal quantity associated with the Bellman equation (2) is increasing in the persistent level of quality, η. That is, if $\eta' > \eta$ and q' and q are optimal choices of quantity given individual states $(\eta', \gamma, \mathbf{q}_-)$ and $(\eta, \gamma, \mathbf{q}_-)$, respectively, then $q' > q$.*

In this section, I only sketch the proof, leaving the details for Appendix A. Recall Equation 2 in Section 3.3. To prove the proposition, I use a method similar to Hopenhayn and Prescott [1992], adopted from Topkis [1998], and I show that the objective function has increasing differences. First note that the optimal choice of ϕ^s does not affect future values. Moreover, the only way past quantities affect current profits is through their effects on Powerseller status. Thus, I can define the following auxiliary period profit function:

$$\hat{\pi}\left(q,\eta,\gamma,\phi^p\right)=\max_{\phi^s\in\{0,1\}} p\left(q;\phi^s,\phi^p\right)q-cq-c^s\phi^s$$

subject to

$$\phi^s \quad = \quad 0 \quad \text{if} \quad \eta+\gamma<\mu^s.$$

I prove the proposition in three steps:

Step 1. Period profit function $\int\hat{\pi}\left(\eta,\gamma,q,\phi^p\right)dF\left(q|\tilde{q}\right)$ is supermodular in $(\eta,\tilde{q})$, in (η,q_{-i}) for $i=1,2,3$, and $(\tilde{q},q_{-i})$ for $i=1,2,3$.

Step 2. The solution to Equation 2 is supermodular in (η,q_{-i}) for $i=1,2,3$.

Step 3. The policy function is increasing in quality η.

The intuition for this result is the existence of a dynamic complementarity between sellers' persistent level of quality and their quantity choices. That is, a seller with a higher persistent level of quality will have a higher probability to meet the quality eligibility for the Powerseller status in the future, which results in a price premium given the assumption on the demand function; I will describe this in detail in section 4.1. Thus, a seller with a high persistent level of quality has more incentive to sell more items in order to meet the quantity eligibility for the Powerseller status. Proposition 1 also makes it clear that the only determinant of seller size dynamics is reputation, and sellers are willing to increase their size in anticipation of future Powerseller and store status. In the absence of these mechanisms, though, sellers have no incentive to change their size over time.

The assumption required for this result is that the demand function faced by each seller increases with Powerseller and store status. Moreover, this functional form must be supermodular in (q,ϕ^s,ϕ^p). As I show in section 4.1, both of these assumptions are satisfied by the buyers' utility specification.

Note that in this proof, I rely heavily on the concept of oblivious equilibrium; thus, in any equilibrium, the only statistic relevant for every given seller is the path of aggregate quantity produced in the market. Since sellers only care about the path of market aggregates, the above proposition establishes for any sequence of market aggregates that sellers' quantity choice is monotone in its quality. This is regardless of what these aggregate quantities are arising from and whether they are consistent with equilibrium behavior. In other words, the proposition can be interpreted as stating that the best response of any given seller to a path of market aggregates is monotone in that seller's persistent level of quality.

Another implication of the model for the quantity choice of sellers is that optimal quantity choice can be represented as a function of sellers' persistent level of quality, their Powerseller and store statuses, and their quantity in the last two periods. Controlling for Powerseller and store statuses allows me to drop sellers' transitory shock to quality, γ_{jt}, as well as their quantity three periods before, q_{-3}.

Lemma 1. *The policy function* $q^*(\eta, \gamma, \mathbf{q}_-, \Omega)$ *can also be represented as* $q^*(\eta, \phi^s, \phi^p, q_{-1}, q_{-2}, \Omega)$.

Proof. Sellers choose the quantity of items to sell after their Powerseller status has been determined and they have selected store status. The profit function of sellers, $\pi(q_j, \phi_j^p, \phi_j^s, \Omega)$, and their expectation of continuation value function, $\int V(\eta_j, \gamma', \mathbf{q}'_-, \Omega')g(\gamma)d\gamma$, are not directly a function of γ or q_{-3}. Therefore, sellers' choice of quantity should not depend on γ or q_{-3} after controlling for ϕ_j^p and ϕ_j^s. □

The above lemma will help in modeling sellers' choice of quantity in Section 5. Note that Proposition 1 can also be extended to the policy function with this new representation, and this policy function is also weakly increasing in the persistent level of quality.

4 Identification Procedure

In this section, I describe the method to identify the main parameters of the model: sellers' unobservable quality, their cost parameters, and buyers' utility function. These are the deep parameters of the model that affect buyers' and sellers' decisions, and they are used in my counterfactual analysis. An important step in this identification is using the monotonicity result from Section 3.5.

4.1 Identification of Demand

Before discussing the identification of sellers' unobservable quality and their cost parameters, I describe in more detail the formulation of demand and its identification. As mentioned earlier, buyers do not directly observe sellers' quality, but it affects their utility. If buyers do not observe any signal from sellers, their expected utility from buying an item is

$$E(u_{ijt}) = -\alpha p_{jt} + \beta_r E(r_{jt}) + \xi_t + \xi_{jt} + \epsilon_{ijt}.$$

Assuming that a seller sells only one type of good each period, then the market share of seller j at time t, given that the distribution of error terms is coming from the logit distribution, is

$$s_{jt} = \frac{exp(-\alpha p_{jt} + \beta_r E(r_{jt}) + \xi_t + \xi_{jt})}{1 + \sum exp(-\alpha p_{j't} + \beta_r E(r_{j't}) + \xi_t + \xi_{j't})}.$$

Following Berry [1994], I assume the utility of the outside good is normalized to 0. Then I can decompose the formula for the market share, using the formula of outside good share, s_{0t}:

$$log(s_{jt}) - log(s_{0t}) = -\alpha p_{jt} + \beta_r E(r_{jt}) + \xi_t + \xi_{jt}.$$

Therefore,

$$p_{jt} = (-log(s_{jt}) + log(s_{0t}) + \beta_r E(r_{jt}) + \xi_t + \xi_{jt})/\alpha.$$

The demand function can be generalized in the case when buyers observe signals of quality, such as Powerseller status, ϕ^p_{jt}, and store status, ϕ^s_{jt}. In this case, buyers' expected utility function is

$$E(u_{ijt}|\phi^p_{jt},\phi^s_{jt}) = -\alpha p_{jt} + \beta_r E(r_{jt}|\phi^p_{jt},\phi^s_{jt}) + \xi_t + \xi_{jt} + \epsilon_{ijt}.$$

The same set of analyses as above leads to the following pricing function:

$$p_{jt} = (-log(s_{jt}) + log(s_{0t}) + \beta_r E(r_{jt}|\phi^p_{jt},\phi^s_{jt}) + \xi_t + \xi_{jt})/\alpha,$$

where $E(r_{jt}|\phi^p_{jt},\phi^s_{jt})$ is the expectation of a seller's quality based on the seller's two reputational signals. This expectation is endogenously determined by equilibrium decisions of sellers in the market and is subject to change based on the market setup.

Note that ϕ^p_{jt} and ϕ^s_{jt} are binary variables and can only be 0 or 1. Let $\bar{r}_{mn} = E(r_{jt}|\phi^p_{jt} = m, \phi^s_{jt} = n)$. Then, $E(r_{jt}|\phi^p_{jt},\phi^s_{jt})$ can be written as

$$E(r_{jt}|\phi^p_{jt},\phi^s_{jt}) = \bar{r}_{00} + (\bar{r}_{10} - \bar{r}_{00})\phi^p_{jt} + (\bar{r}_{01} - \bar{r}_{00})\phi^s_{jt} + (\bar{r}_{00} - \bar{r}_{10} - \bar{r}_{01} + \bar{r}_{11})\phi^p_{jt}\phi^s_{jt}.$$

Substituting the above expression into the demand function formula leads to the following formulation:

$$\begin{aligned} p_{jt} &= [-log(s_{jt}) + log(s_{0t}) + \xi_t + \xi_{jt}]/\alpha \\ &\quad + \beta_r/\alpha[\bar{r}_{00} + (\bar{r}_{10} - \bar{r}_{00})\phi^p_{jt} + (\bar{r}_{01} - \bar{r}_{00})\phi^s_{jt} + (\bar{r}_{00} - \bar{r}_{10} - \bar{r}_{01} + \bar{r}_{11})\phi^p_{jt}\phi^s_{jt}] \\ &= [-log(s_{jt}) + log(s_{0t})]/\alpha + \bar{r}_{00}\beta_r/\alpha + \beta_p\phi^p_{jt} + \beta_s\phi^s_{jt} + \beta_{ps}\phi^p_{jt}\phi^s_{jt} + [\xi_t + \xi_{jt}]/\alpha. \end{aligned} \quad (4)$$

This formula can be used to estimate the parameters of the demand function, which gives me a mean to estimate the deep parameters of buyers' utility function. The estimation method for the above formula will be discussed in Section 5. Note that since the demand function is additively separable between q_{jt} and $\left(\phi^s_{jt}, \phi^p_{jt}\right)$, it is obviously supermodular in (q, ϕ^s) and (q, ϕ^p), i.e., its cross partials are 0. Moreover, as I show in section 5, the interaction coefficient β_{ps} is positive, which implies that the demand function is supermodular in (q, ϕ^s, ϕ^p).

4.2 Identification Procedure

Three main sets of parameters need to be identified: buyers' utility function, sellers' unobservable quality, and sellers' cost parameters. The first two are estimated using a three-step procedure, while the third is identified using a two-step method similar to Bajari et al. [2007]. [23]

The three-step procedure is used to estimate the unobservable sellers' quality and buyers' utility function. The following is an overview of the procedure:

1 Estimating the structural demand function. This gives me an estimate of α, β_p, β_s, and β_{ps}.

2 Estimating the realized policy functions. Given Proposition 1 and Lemma 1, the quantity choice of sellers can be used to identify sellers' quality. If I control for Powerseller and store statuses and sellers' quantity choice in the last two periods, I find that sellers' fixed effects are an index of sellers' persistent level of quality.

3 Using the simulated method of moments to estimate the remaining five parameters of the model. The five moments are number of Powersellers, number of registered stores, number of sellers who have both statuses, and two moments from the demand function (Equation 4):

$$\begin{aligned} &(\bar{r}_{10} - \bar{r}_{00})/\beta_p - (\bar{r}_{01} - \bar{r}_{00})/\beta_s = 0 \\ &(\bar{r}_{10} - \bar{r}_{00})/\beta_p - (\bar{r}_{00} - \bar{r}_{10} - \bar{r}_{01} + \bar{r}_{11})/\beta_{ps} = 0. \end{aligned} \tag{5}$$

The five parameters to be estimated are the quality thresholds for Powerseller and store statuses, μ_p and μ_s; the coefficient of quality in the utility function of buyers, β_r/α; the

[23] Some of the steps in identifying unobservable quality and cost parameters do overlap.

parametric variable that converts the index of persistent quality found in the previous step to that value; and the variance of random shocks to quality. The estimation procedure will be discussed in detail in Section 5.2.3.

The next step is estimating the cost parameters of sellers using a two-step estimator method similar to the method in Bajari et al. [2007]. The method uses the basics of revealed profit to estimate the deep parameters of the model and, in this case, to estimate the cost parameters: the average monthly cost sellers should pay to become a store on eBay, c^s, and the average cost of obtaining an iPod for sellers, c.

In the first step of this method, I estimate the structural demand function of buyers and policy functions of sellers. Then, assuming the estimated policy functions are the optimal choices of sellers, any perturbation of these functions should yield a value function lower than the original value function when I use the original policy function. The cost parameters are those that satisfy the above condition. The two-step estimation procedure is as follows:

1A Estimating the structural demand function.

1B Estimating the realized policy functions.

2A Perturbing the policy functions.

2B Simulating the model using the original policy functions and the perturbed policy functions.

2C Defining the loss function as a function of the model parameters,

$$\sum_{sellers,perturbations} (V_{perturbed}(\theta^c) - V_{original}(\theta^c))^2 \mathbf{1}[(V_{perturbed}(\theta^c) - V_{original}(\theta^c)) > 0],$$

where C is the vector of cost parameters; $V_{perturbed}(\theta^c)$ is the value function using perturbed policy functions; $V_{original}(\theta^c)$ is the value function using the original policy functions; and $\mathbf{1}[V_{perturbed}(\theta^c) - V_{original}(\theta^c) > 0]$ is an indicator function that is equal to 1 if

$V_{perturbed}(\theta^c) - V_{original}(\theta^c) > 0$, and 0 otherwise.[24] If this expression is positive, it means that the seller's value function is higher for perturbed policy functions, which cannot be the case if θ_c is the true cost parameter. The summation is taken over all sellers and perturbations.

2D Estimating the cost parameters after minimizing the loss function as defined above. Under the true cost parameters of the model, the estimated policy functions are optimal. Therefore, the cost parameters that minimize the above loss function are consistent. Note that given the inequality conditions, there is a possibility of set-identification if the above inequalities hold for a range of cost parameters.

5 Estimation

In this section, I estimate the deep parameters of the model using the identification procedure explained in Section 4.2.

5.1 Estimating Structural Demand

To estimate the structural demand function, I use the demand Equation 4 derived in Section 4.1. This formula translates into a simple OLS regression of price over the logarithm of sellers' share minus the outside good's share, Powerseller status, store status, and characteristics of the item and listing. Various variables for the characteristics of listings and items are considered; in Section B I review some of the different setups for demand. Unlike for the demand function in the model, I include item characteristics in this regression; however, in simulations, I use the average characteristics of iPods in my dataset when picking the iPod.

Table 3 shows the results of the regression. The effect of changes in $log(s_0) - log(s_j)$ is captured by $1/\alpha$ and is positive. This means that when sellers sell more items, they sell at a lower price

[24]The results are robust to other definitions of loss function, i.e., sum of the absolute values of deviations.

per unit. Therefore, the demand function is elastic. Moreover, the coefficient of Powerseller status is positive, which shows that the expectation of quality is higher for Powersellers. Finally, the coefficient of store status is positive, which shows that the expectation of quality is higher for sellers who are a store than for others. Both of these observations are consistent with the results in Section 3: high-quality sellers become Powersellers and registered stores. It should also be noted that since the coefficient of the interaction term between Powerseller and store status is positive, the demand function is supermodular in (q, ϕ^s, ϕ^p) and therefore satisfies the assumption of Proposition 1.

The above regression also determines how the characteristics of the iPods sold affect their price. Compared to used iPods, new iPods have a premium of \$39.45 on average, and refurbished iPods have a premium of \$14.20 on average. Each extra gigabyte of internal memory on an iPod results in an extra \$1.49 in price. I also include a fixed effect for the type of iPods: Nano, Touch, Classic, Mini, Video, and Shuffle; their coefficients are as expected, with the highest for Touch and the lowest for Shuffle. Month fixed effects are also included to treat the seasonal fixed effects; each month is considered to be a period. In Appendix B I carry out additional robustness checks on this demand formulation by adding more characteristics of sellers and focusing on a subset of data.

It should be noted that there is some concern about the possible endogeneity of sellers' market share, as sellers may adjust their supply in response to predicting a better-than-average month on eBay. However, the usual instruments cannot be applied to this setting, because of a lack of information about the sellers beyond their choices of quantity over time. To control for seasonality to the best of my abilities, I include monthly time fixed effects.

5.2 Estimating Policy Functions and Sellers' Quality

In this section, I estimate sellers' policy functions and their persistent level of quality. Sellers have two policy functions in this model: number of items to sell and the decision to become an eBay store.. In what follows, I describe, in detail, the estimation of these policy functions as well as that of sellers' persistent level of quality.

5.2.1 Quantity Choice

One of the decisions that sellers make each period is the expected number of items they list on eBay. Sellers' optimal quantity choice depends on their persistent level of quality, Powerseller status, store status, and choice of quantity in the past two periods, as discussed in Section 3.5. I can control for all of these variables except for the unobservable persistent level of quality, η_j. In order to estimate the persistent level of quality, I use the theoretical result in Proposition 1, which states that sellers' quantity choice is an increasing function of their persistent level of quality. This implies that after controlling for all other variables, a seller's fixed effect can be interpreted as being a function of quality. In Section 5.2.3, I use the estimated fixed effects in this section to further estimate this function of quality.

As mentioned in Section 3.3, sellers choose their expected number of items to sell, but the actual number of items they end up selling may not be exactly the same. This assumption allows me to match the variations in quantity observed in the data even after controlling for observables such as Powerseller or store status as well as past quantities.[25] To capture this randomness, I use a negative binomial distribution. Specifically, I assume that q_{jt} is a random variable that takes values in $\mathbb{N} \cup \{0\}$ and its probability distribution satisfies the following recursion

$$\Pr\left(q_{jt}=0\right)=\left(\frac{r}{r+\rho_{jt}}\right)^{r}, \Pr(q_{jt}=k)=\Pr\left(q_{jt}=k-1\right)\frac{r+k-1}{k}\left(\frac{\rho_{jt}}{\rho_{jt}+r}\right).$$

In the above, $r=1/\kappa$ is the dispersion parameter to be estimated, and ρ_{jt},which is the average

[25] Without considering the inventory shock, the model is rejected, as there can be a seller who has the same observable characteristics in two periods and different quantity choices.

value of quantities sold by seller j at time t, has the following specification

$$\rho_{jt} = \exp\left(\left[\phi^s_{jt}, \phi^p_{jt}, q_{j,t-1}, q_{j,t-2}\right] \cdot \boldsymbol{\beta} + \nu_j + \delta_{t,j}\right).$$

In the above specification, $\boldsymbol{\beta}$ captures how observables affect average quantity, and ν_j is a seller fixed effect to be estimated. These estimated coefficients together with the dispersion parameter κ are shown in Table 4. As mentioned above, sellers' choice of quantity is positively correlated with their persistent level of quality. This means that sellers with a higher persistent level of quality should have a higher fixed effect in this formulation as well. In section 5.2.3, I will use this relationship to parametrically identify the persistent level of quality from this fixed effect.

Figure 5 shows the proportion of times that sellers in the market have sold n number of items in the data versus the distribution implied by the negative binomial distribution. It also shows the results for an alternative estimation using a Poisson distribution. As it can be seen, the negative binomial distribution does well in fitting the data.

5.2.2 Registered-Store Status

Sellers who meet certain quality requirements can register as eBay stores, for which they pay a monthly fee, and can display the eBay store badge on their listing pages. I assume that sellers decide on their store status each period after knowing the shock to their quality and, therefore, their Powerseller status. In this section, I estimate this policy function using the observed data on store choices of sellers.

According to the model, this decision is a function of sellers' state variables, the shock to their quality, and Powerseller status. However, when sellers make the decision about store status, the value of current shock will not affect their profit that period and the following periods after controlling for their Powerseller status; therefore, it will not enter their decision-making process. Thus, sellers' choice of store status, after qualifying for it, is a function of their Powerseller status, persistent level of quality, and their quantity in the past two periods: $\phi^{s*}(\eta, \phi^p, q_{-1}, q_{-2}, \Omega)$. I use the index for quality, ν_j, estimated in the previous section to control for sellers' persistent level of quality. This decision is a binary choice for sellers, and I model it using the logit distribution. Table 4 shows the results of the regression.

5.2.3 Estimating Unobservable Quality

In this section, I estimate sellers' unobservable persistent level of quality as well as some remaining parameters from the demand and thresholds for Powerseller and store status. To do so, I use the simulated method of moments where I minimize the difference between five moments from the simulated model and the data.

To estimate sellers' unobservable persistent level of quality, η_j, I need to relate it to the estimated sellers' fixed effect, ν_j, from Section 5.2.1. As discussed before, the estimated sellers' fixed effect is a function of η_j and–through Proposition 1–is a non-decreasing function of this value. I also assume the following parametric formulation for η_j:[26]

$$\eta_j = \text{sign}(\nu_j)|\nu_j|^{\lambda}.$$

The next step is constructing the two moments found from the demand formulation, shown in Equation 5, where they relate parameters from the demand estimation to the expected values of quality of sellers with and without Powerseller and registered store status. The other three moments are percentage of Powersellers, percentage of registered stores, and percentage of Powersellers and registered stores.

Finally, by minimizing the joint differences between the simulated moment conditions using the model and the calculated moments using the data, I estimate the value of λ, μ^p, μ^s, variance of random shocks to quality, γ_{jt}, and β_r/α, the coefficient of r_{jt} in the demand function. Table 5

[26]Given that ν takes negative numbers, the functional form is designed to be valid for negative numbers and at the same time be increasing to be consistent with the main proposition of the paper.

shows the estimated values for λ and β_r/α. Note that β_r/α is positive; therefore, buyers enjoy buying an item from a seller with a higher level of quality.

5.3 Perturbations and Cost Estimation

In the second step, I perturb the policy functions, simulate the actions of sellers over time, and estimate the value functions of sellers for each perturbation. This helps in determining out-of-equilibrium revenue values for sellers. To get the perturbations, one should perturb only one seller at a time to be consistent with the equilibrium definition. Moreover, the perturbations should be big enough to change sellers' policy functions and therefore the actions they take. To estimate the cost parameters, I perturb the policy function associated with the number of sales and store status. I make sure that perturbations both increase and decrease the number of items sellers choose to sell, in both large and small magnitudes.

Having the perturbed actions of sellers and also their original simulated actions over time, I can estimate the expected value function for sellers, given a set of initial conditions for cost parameters. Optimality of choice by sellers means that true cost parameters should result in higher expected value functions driven by original policy functions compared to those driven by perturbed policy functions. To estimate the cost parameters, that is, the parameters that minimize this function, I construct a loss function, as explained in Section 4.2.

Table 6 shows the estimated cost parameters for two different specifications. In the first specification, I force the monthly cost of becoming a registered store to be 0, and I estimate the marginal cost of acquiring an iPod for sellers that rationalize sellers' choices. Note that this is the cost associated with the average iPod in the data, as I do not explicitly model the choice of iPod. In the second specification, I jointly estimate the marginal cost of acquiring an iPod for sellers and the monthly fee for becoming a registered store. The monthly fee charged by eBay for store status is between \$15 and \$300 for different types of stores, which I abstract from modeling, and my estimate for the monthly fee is \$44.40 per month, which is within the range of actual monthly costs for store status. Standard deviations are estimated using the bootstrapping method. To do that, I resample data with substitution and estimate the mean and standard errors of the mean to estimate the standard error of the cost parameters.

6 Counterfactuals: Value of Reputation

Given the estimated values for sellers' cost and quality, and buyers' utility function, I can run various counterfactuals to estimate the overall effect of reputation on the market. In the first counterfactual, I assume sellers have no means of signaling their quality, and I recalculate the equilibrium demand and surplus. In the second counterfactual, I assume sellers whose quality level is higher than a set threshold can provide warranty to buyers. Warranty promises that a seller's quality is higher than this threshold; otherwise, the seller must refund the buyer in full. Warranty increases high-quality sellers' profit and decreases low-quality sellers' profit, by giving buyers the ability to distinguish between the two groups. Therefore, in this case warranty can be used as a substitute for reputation.

6.1 No Reputation Mechanism

As mentioned before, Powerseller status and store status are tools used by eBay to signal sellers' quality. This helps a high-quality seller to sell more products on eBay, and it helps buyers recognize high-quality sellers and have an overall better experience in the marketplace. To find the overall effect of reputation, I run a counterfactual in which no Powerseller, store status, or any other reputation signal exists. Note that even though I do not find big effects for other reputation signals in the presence of Powerseller and store status, they may become relatively more important if they are the only signal to rely on. Without any quality signals, sellers are all pooled together; therefore, high-quality sellers cannot receive price or quantity premiums.

In the absence of reputational signals, buyers' demand function will change, as well as the problems that sellers face. On the one hand, buyers will no longer observe any reputational signals for quality, so buyers cannot infer sellers' quality based on these signals. On the other hand, as sellers cannot signal their quality levels to buyers, those with different quality levels will face the same problem.

6.1.1 Sellers' Problem

In the absence of reputation in the market, sellers are pooled together and cannot signal their quality. I assume they are competing in a Cournot equilibrium, given the new demand function in the market. I further assume there is a symmetric equilibrium in the market, and that sellers face the same inventory shock as in the original problem.[27] Sellers maximize their expected profit, given their marginal cost and the number of items sold by other sellers in the market. Their period t profit function is

$$\pi(q_{jt}) = p(q_{jt}, \Omega)q_{jt} - cq_{jt}.$$

The expectation is taken over the inventory shock that sellers face. The dispersion is assumed to be unaffected by the policy change and to remain the same for all sellers. Sellers' choice is not dynamic in this setup, as it is assumed that buyers do not observe sellers' past actions. The goal is to find the Nash equilibrium; that is, sellers' choice of average number of items to sell should be optimum, given other sellers' choice. Sellers face the same final price and the same sellers' problem; therefore, I assume they face a symmetric equilibrium, $q_{jt} = q_t$.

[27] In Section 5.2.1, I assumed that sellers faced an inventory shock; and even though the average number of items they sell is equal to their optimal number, they may end up with a larger or smaller number of items sold. The dispersion was assumed to be the same for all sellers in the market. When performing counterfactuals, I assume that this dispersion parameter remains the same. The number of items sold will follow a negative binomial distribution with a fixed dispersion rate for all sellers.

6.1.2 Demand Function

The demand function in the counterfactual study is assumed to follow the same utility function estimated in the original problem. However, in the new setup, buyers do not observe sellers' quality, nor do they observe any signals related to that quality; hence, only the expected value of sellers' quality affects buyers' expected utility function.[28] The expectation is taken over all the listings and sellers in the market. Note that since sellers cannot signal their quality, there is no observable heterogeneity among sellers. Given sellers' quality comes from distribution function L estimated from the original problem, buyers' expected utility function is

$$E(u_{ijt}) = -\alpha p_{jt} + \beta_r E(r) + \xi_t + \xi_{jt} + \epsilon_{ijt}$$

where $E(r)$ is the unconditional average value of quality. Given the above utility function and assuming that ϵ_{ijt} follows an extreme value distribution, the demand function as explained in Section 4.1 will be as follows:

$$p_{jt} = \frac{1}{\alpha}(-log(s_{jt}) + log(s_{0t})) + \frac{\beta_r}{\alpha}E(r) + \frac{\xi_t}{\alpha} + \frac{\xi_{jt}}{\alpha}, \tag{6}$$

where α has the same parametric values as estimated parameters in Table 3 in the previous section, implying they are invariant to the changing policies of eBay. I use the results in Section 5.2.3 to estimate $\beta_r/\alpha E(r)$, which gives me an estimate of β_r/α and also an estimate of $r_{jt} = \eta_j + \gamma_{jt}$. Note that γ is distributed i.i.d. with a mean of 0.

6.1.3 Estimation and Results

To estimate the optimal level for q_t, I start with an initial guess, q. I then estimate the best response of a seller when the other sellers choose q using simulation, and call this optimal level $B(q)$. Next,

[28] In contrast, in the original problem, buyers were able to observe the two signals of quality, which they used to get conditional expected values of sellers' quality, as discussed in Section 4.1.

I change the initial guess to a number between q and $B(q)$. I repeat these two steps until q and $B(q)$ are close enough, thereby achieving the optimal level of q_t.

After solving for sellers' new policy functions, I simulate the model to obtain sellers' expected value function, eBay's profit, and buyers' surplus.[29] The results are shown in Table 7. As a result of the policy change, consumer surplus has decreased by 35%, eBay's profit by 38%, and the total sellers' profit by 66%.

One of the reasons I get large effects as a result of removing the reputation mechanism is that in my original setting, even among sellers who are not Powersellers, those with higher quality levels sell more. Recall that to become a Powerseller, a seller must reach a quantity threshold set by eBay for three months. Therefore, high-quality non-Powersellers have incentive to sell more items than their static optimal values because they are more likely to become Powersellers in the future. This results in a high level for the average quality of items sold even by non-Powersellers, as shown in Figure 6. Figure 7 shows the number of items sold by both Powersellers and non-Powersellers. Sellers with higher quality values sell more, and Powersellers have extra incentive to sell more to retain their Powerseller status. In contrast, in the case of no signaling, sellers with higher levels of quality do not have incentives to sell more items and the market share of all seller types becomes the same. Thus, in this case, the average price of items drops even below the price of items sold by non-Powersellers. This will result in a substantially lower number of items sold in the marketplace and, subsequently, lower profit for eBay and even lower consumer surplus despite lower prices. These results emphasize the importance of dynamics in the value of reputation. The effects found here from removing the reputation mechanism are much higher than the static values I found for reputation or those discussed in the literature. Furthermore, these large effects emphasize the importance of a reputation mechanism for different marketplaces, especially newly introduced e-commerce marketplaces, and the importance of their design.

[29] Even though the sellers' problem is static, I simulate it over time to both account for seasonal demand shock and compare the counterfactual problem with the original setup more easily.

An important issue to note here is that I abstract from modeling moral hazard given the data limitations. However, the effects of removing the reputation mechanism will be even larger in the presence of moral hazard, as sellers will have little incentive to exert effort and their quality will drop in the absence of reputation. In this study, the result comes mainly from a change in market structure by shifting the market share from high-quality sellers to low-quality sellers, but the additional reduction in the quality of sellers can exaggerate the result even further.

6.2 Warranty

Warranty can be a substitute for a reputation mechanism. In this section, I consider a simplified setup for a warranty mechanism. Implementing a warranty mechanism involves many screening costs, such as confirming buyers' claim and having a mean of getting a refund from seller to buyer. However, I abstract from such costs in this simplified setup, and I assume sellers can voluntarily provide warranty for the items they sell on the market. They guarantee that their quality is above a set threshold and if proven wrong, the buyer can return the item for full refund. For simplicity, I assume that the threshold is fixed for all sellers and announced by eBay. There is no track record of sellers and therefore the problem is static. In this case, sellers will be divided into two groups: sellers with high quality who provide warranty and sellers with low quality who do not provide warranty.[30]

Buyers' expected utility function depends on the warranty option and can be represented as

$$E(u_{ijt}|w) = -\alpha p_{jt} + \beta_r E(r_{jt}|w) + \xi_t + \xi_{jt} + \epsilon_{ijt},$$

where w is equal to 1 if the seller provides warranty, and 0 otherwise. The demand function from the above formula is similar to the ones in the previous section, the only difference being the warranty option. I assume that in equilibrium, all sellers within the high- and low-quality groups choose

[30] Note that by assumption sellers with low quality have no incentive to pretend to be of high quality, and sellers with high quality benefit from being grouped with high-quality sellers; therefore, they choose to provide warranty.

to sell the same number of items, q_H and q_L, respectively. However, sellers still face inventory shocks. To find q_L and q_H, I start from an initial value and first find the fixed point of the best response function for low-quality sellers, given q_H. Then, given q_L^* from the last step, I find the fixed point of best response for high-quality sellers. I continue these steps until convergence. The surplus numbers are shown in Table 7. Different levels of surplus are shown for quality thresholds corresponding to top 5%, 10%, and 50% of sellers providing warranty. Note that when the threshold goes to $-\inf$ or $+\inf$, the surplus values converge to the case of no signaling.

As shown in Table 7, warranty can be a good substitute for reputation. When sellers provide warranty, total sellers' profit goes up. This higher total profit is a result of high-quality sellers getting separated from low-quality sellers and obtaining substantially higher profits. However, low-quality sellers suffer as a result of receiving even a lower price compared to that in the case with no reputation mechanism. Consumer surplus and eBay's surplus can be lower or higher than in the case with no reputation mechanism: When only a few sellers are signaled, the price for the products sold by these sellers increases considerably, and the quantity of products sold by other sellers decreases even more, which in turn leads to lower consumer surplus and eBay's profit. When the threshold is at 50%, all surplus values are higher than those in the case with a reputation mechanism. Note that this increase does not necessarily imply that warranty is a better choice than reputation, for two main reasons. First, the reputation mechanism set by eBay is not the optimal mechanism, and it can be improved. Second, the warranty considered in this case is simplified, and it is difficult to implement the thresholds.

As mentioned above, implementing a warranty mechanism is costly, as it involves screening sellers and buyers for various claims, similar to policing the marketplace. On the other hand, a reputation mechanism can achieve similar surplus levels by using just the track record of transactions. For a marketplace such as eBay, substituting the reputation mechanism with warranty would be a total loss as long as the cost of screening is non-negligible. However, the marketplace can add warranty as well as having a reputation mechanism, and in fact, eBay has done so in 2010. Unfortunately, using this setup, I cannot predict the effect of adding warranty, because the problem becomes intractable. In a follow-up paper, Hui et al. [2015], my coauthors and I study the effect of adding warranty to the eBay marketplace, and we estimate that it leads to a 2.9% increase in total welfare.

7 Conclusion

In this paper, I quantify the value of eBay's reputation mechanism. To do so, I develop a dynamic model of sellers' behavior where they have heterogeneous qualities, unobservable by consumers, and reputation is used as a signal to buyers in order to improve allocations. By structurally estimating this model, I uncover deep parameters of buyers' utility and sellers' costs, as well as sellers' unobservable qualities. The estimated model suggests that reputation has a positive effect on the expected profits of high-quality sellers, as well as their market share. To establish the value of reputation, I perform a counterfactual. The findings show that removing the reputation mechanisms put in place by eBay drops consumer surplus by 35%, eBay's profit by 38%, and total sellers' profit by 66%. This is a result of significant change in the market share of high-quality sellers and consequently the unraveling that causes a significant drop in total market size. In addition, the effect of adding warranty as a substitution for the reputation mechanism is studied. The findings indicate that adding warranty overcomes part of the inefficiencies due to adverse selection.

The results of this paper are of significant importance for the analysis and design of reputation systems. Specifically, the quantitative model can be used to shed light on optimal design of reputation systems, a task that is crucial for online marketplaces. In subsequent work, Hopenhayn and Saeedi [2018], we use variants of this model to study some aspects of designing reputation mechanisms. Further quantitative studies are needed and are left for future work.

References

Bajari, P., C. Benkard, and J. Levin (2007): "Estimating dynamic models of imperfect competition," *Econometrica*, 75, 1331–1370. 1, 4.2, 4.2

Bajari, P. and A. Hortaçsu (2004): "Economic Insights from Internet Auctions," *Journal of Economic Literature*, 42, 457–486. 3, 1

Bar-Isaac, H. and S. Tadelis (2008): "Seller Reputation," *Foundations and Trends® in Microeconomics*, 4, 273–351. 1

Berry, S. (1994): "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, 242–262. 3.2, 4.1

Board, S. and M. Meyer-ter Vehn (2010): "A Reputational Theory of Firm Dynamics," . 2, 6, 19

——— (2011): "Reputation for quality," *UCLA Working Paper*. 2, 6, 19

Brown, J. and J. Morgan (2006): "Reputation in online auctions: The market for trust," *California Management Review*, 49, 61–82. 3

Cabral, L. and A. Hortacsu (2010): "The dynamics of seller reputation: Evidence from ebay*," *The Journal of Industrial Economics*, 58, 54–78. 1, 2.2, 14

Cai, H., G. Z. Jin, C. Liu, and L.-A. Zhou (2013): "More trusting, less trust? an investigation of early e-commerce in china," Tech. rep., National Bureau of Economic Research. 1

Chari, V., A. Shourideh, and A. Zetlin-Jones (2014): "Adverse Selection, Reputation and Sudden Collapses in Secondary Loan Markets," *American Economic Review*. 6

Crawford, G., N. Pavanini, and F. Schivardi (2015): "Asymmetric information and imperfect competition in lending markets," . 1

DELLAROCAS, C. (2005): "Reputation mechanisms," *Handbook on Economics and Information Systems*. 1

ERICSON, R. AND A. PAKES (1995): "Markov-perfect industry dynamics: A framework for empirical work," *The Review of Economic Studies*, 62, 53. 22

FAN, Y., J. JU, AND M. XIAO (2016): "Reputation premium and reputation management: Evidence from the largest e-commerce platform in China," *International Journal of Industrial Organization*, 46, 63–76. 1

FILIPPAS, A., J. HORTON, AND J. M. GOLDEN (2017): "Reputation in the Long-Run," Tech. rep., CESifo Working Paper. 1

FINKELSTEIN, A. AND K. MCGARRY (2006): "Multiple dimensions of private information: evidence from the long-term care insurance market," *American Economic Review*, 96, 938–958. 1

HOLMSTROM, B. (1999): "Managerial incentive problems: A dynamic perspective," *The Review of Economic Studies*, 66, 169. 2, 1, 19

HOPENHAYN, H. AND E. PRESCOTT (1992): "Stochastic monotonicity and stationary distributions for dynamic economies," *Econometrica: Journal of the Econometric Society*, 1387–1406. 3.5, A, A

HOPENHAYN, H. AND M. SAEEDI (2018): "Reputation Signals and Market Outcomes," *mimeo*. 7

HOUSER, D., J. WOODERS, AND M. HALL (2006): "Reputation in auctions: Theory, and evidence from eBay," *Journal of Economics & Management Strategy*, 15, 353–369. 1

HUI, X., M. SAEEDI, Z. SHEN, AND N. SUNDARESAN (2015): "Reputation and Regulations, Evidence from eBay," *Forthcoming, Management Science*. 5, 15, 6.2

Ivashina, V. (2009): "Asymmetric information effects on loan spreads," *Journal of financial Economics*, 92, 300–319. 1

Kollock, P. (1999): "The production of trust in online markets," *Advances in group processes*, 16, 99–123. 3

Kreps, D. and R. Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–279. 1

Lewis, G. (2011): "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, forthcoming. 1

Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves (2007): "PENNIES FROM EBAY: THE DETERMINANTS OF PRICE IN ONLINE AUCTIONS," *The Journal of Industrial Economics*, 55, 223–233. 3

Mailath, G. and L. Samuelson (2001): "Who Wants a Good Reputation?" *Review of Economic Studies*, 68, 415–441. 2, 1, 3, 19

Masclet, D. and T. Pénard (2008): "Is the eBay feedback system really efficient? An experimental study," *Economics Working Paper Archive (University of Rennes 1 & University of Caen)*. 1

Melnik, M. and J. Alm (2002): "Does a seller's ecommerce reputation matter? Evidence from eBay auctions," *Journal of Industrial Economics*, 337–349. 1

Milgrom, P. and J. Roberts (1982): "Predation, Reputation, and Entry Deterrence," *Journal of economic theory*, 27, 280–312. 1

Nosko, C. and S. Tadelis (2014): "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment," . 2.1, 13

Pakes, A. and P. McGuire (2001): "Stochastic algorithms, symmetric Markov perfect equilibrium, and the'curse'of dimensionality," *Econometrica*, 1261–1281. 1, 3.4

Reiley, D., D. Bryan, N. Prasad, D. Reeves, M. Hall, K. Engines, T. Inc, and C. San Mateo (2007): "Pennies from eBay: The determinants of price in online auctions," *Journal of Industrial Economics*, 55, 223–233. 1

Resnick, P. and R. Zeckhauser (2002): "Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system," *Advances in Applied Microeconomics: A Research Annual*, 11, 127–157. 1

Resnick, P., R. Zeckhauser, J. Swanson, and K. Lockwood (2006): "The value of reputation on eBay: A controlled experiment," *Experimental Economics*, 9, 79–101. 3, 1

Topkis, D. (1998): *Supermodularity and complementarity*, Princeton Univ Pr. 3.5, A, A, A

Weintraub, G., C. Benkard, and B. Van Roy (2006): "Oblivious equilibrium: A mean field approximation for large-scale dynamic games," *Advances in Neural Information Processing Systems*, 18, 1489. 1, 22

——— (2008): "Markov perfect industry dynamics with many firms," *Econometrica*, 76, 1375–1411. 1, 3.4

Weintraub, G. Y., C. L. Benkard, P. Jeziorski, and B. Van Roy (2010): "Nonstationary oblivious equilibrium," Tech. rep., Working papers//the Johns Hopkins University, Department of Economics. 3.4

Yamagishi, T. and M. Matsuda (2002): "Improving the lemons market with a reputation system: An experimental study of internet auctioning," *University of Hokkaido, Tech. Rep., May*. 3

Yin, P. (2003): "Information Dispersion and Auction Prices," . 1

Zervas, G., D. Proserpio, and J. Byers (2015): "A first look at online reputation on Airbnb, where every stay is above average," . 1

Appendix A Proof of Proposition 1

Before proceeding with the proof of Proposition 1, I will show a preliminary lemma that will be used frequently. Other versions of this can be found in Topkis [1998]; see Theorem 2.7.6 for a special case:

Lemma 2. *Suppose that the function* $g(x,y,z)$ *is supermodular in* (x,y,z) *and the correspondence* $\Gamma(x,y)$ *is supermodular, i.e., if* $z \in \Gamma(x,y), z' \in \Gamma(x',y')$*, then* $z \wedge z' \in \Gamma(x \wedge x', y \wedge y')$ *and* $z \vee z' \in \Gamma(x \vee x', y \vee y')$ *where* $x \wedge x' = \min\{x, x'\}$ *and* $x \vee x' = \max\{x, x'\}$*. Then the function* $f(x,y) = \max_{z \in \Gamma(x,y)} g(x,y,z)$ *is supermodular in* (x,y)*.*

Proof. Let $z \in \Gamma(x,y)$ and $z' \in \Gamma(x',y')$. Then, by supermodularity of g, we must have that

$$g(x,y,z) + g\left(x',y',z'\right) \leq g\left(x \vee x', y \vee y', z \vee z'\right) + g\left(x \wedge x', y \wedge y', z \wedge z'\right).$$

Since $\Gamma(\cdot)$ is supermodular, we have that $z \vee z' \in \Gamma(x \vee x', y \vee y')$ and $z \wedge z' \in \Gamma(x \wedge x', y \wedge y')$. Therefore, by definition of $f(x,y)$ we must have that

$$g\left(x \vee x', y \vee y', z \vee z'\right) \leq f\left(x \vee x', y \vee y'\right)$$
$$g\left(x \wedge x', y \wedge y', z \wedge z'\right) \leq f\left(x \wedge x', y \wedge y'\right).$$

Thus, we have that

$$\forall z \in \Gamma(x,y), z' \in \Gamma\left(x',y'\right), g(x,y,z) + g\left(x',y',z'\right) \leq f\left(x \vee x', y \vee y'\right) + f\left(x \wedge x', y \wedge y'\right).$$

Since the above holds for all $z \in \Gamma(x,y)$, it must also hold for the solution of $\max_{z \in \Gamma(x,y)} g(x,y,z)$, and the same thing can be said about x', y', z'. Hence,

$$f(x,y) + f\left(x',y'\right) \leq f\left(x \vee x', y \vee y'\right) + f\left(x \wedge x', y \wedge y'\right).$$

This concludes the proof. □

I will use this lemma to establish the main claim in Proposition 1.

Proof. Recall Equation 2 in Section 3.3. To prove the proposition, I use a method similar to Hopenhayn and Prescott [1992], adopted from Topkis [1998], and I show that the objective function has increasing differences. First note that the optimal choice of ϕ^s does not affect future values. Moreover, the only way past quantities affect current profits is through their effects on Powerseller status. Thus, I can define the following auxiliary period profit function:

$$\hat{\pi}\left(q, \eta, \gamma, \phi^{p}\right)=\max_{\phi^{s} \in\{0,1\}} p\left(q ; \phi^{s}, \phi^{p}\right) q-c q-c^{s} \phi^{s}=\max_{\phi^{s} \in\{0,1\}} \pi\left(q ; \phi^{s}, \phi^{p}\right) \tag{7}$$

subject to

$$\phi^{s} \quad = \quad 0 \quad \text{if} \quad \eta+\gamma<\mu^{s}.$$

To ease notation, I have suppressed the dependence of functions on Ω; the proof of the general case is identical but notationally more cumbersome. Note that given the determinants of Powerseller status, ϕ^p is given by the following function:

$$\tilde{\phi}^{p}\left(\eta, \gamma, \mathbf{q}_{-}\right)=1 \quad \text{iff} \quad \begin{cases} q_{-1}, q_{-2}, q_{-3}>Q^{p}. \\ \eta+\gamma>\mu^{p} \end{cases}$$

I prove the proposition in three steps:

Step 1. Period profit function $\int \hat{\pi}\left(\eta, \gamma, q, \tilde{\phi}^{p}\left(\eta, \gamma, \mathbf{q}_{-}\right)\right) d F(q \mid \tilde{q})$ is supermodular in $(\eta, \tilde{q})$, in (η, q_{-i}) for $i=1,2,3$, and $(\tilde{q}, q_{-i})$ for $i=1,2,3$.

Step 2. The solution to Equation 2 is supermodular in (η, q_{-i}) for $i=1,2,3$.

Step 3. The policy function is increasing in quality η.

As a general principle, most of the proof consists of repeated applications of Topkis [1998]'s result on supermodular functions.

Step 1. To establish the first step I proceed as follows:

Supermodularity of $\int \tilde{\pi} dF$ with respect to $(\eta, \tilde{q})$

To show this, I first establish that $\hat{\pi}(\cdot)$ is supermodular in (q, η) and in (ϕ^p, η). Since $\tilde{\phi}^p$ is increasing in η, the period profit function $\tilde{\pi}(\eta, \gamma, q, \mathbf{q}_-) = \hat{\pi}\left(\eta, \gamma, q, \tilde{\phi}^p(\eta, \gamma, \mathbf{q}_-)\right)$ is supermodular in (q, η). The supermodularity of $\int \tilde{\pi} dF$ follows from this property and the fact that $F(q|\tilde{q})$ is increasing in the sense of first order stochastic dominance.

Now, to show the claim, note that the function $\pi(q; \phi^s, \phi^p)$ is supermodular in (q, ϕ^s, ϕ^p). This is because if $q < q'$, then

$$p(q; \phi^s, \phi^p) + p\left(q'; \hat{\phi}^s, \hat{\phi}^p\right) \leq p\left(q'; \phi^s \vee \hat{\phi}^s, \phi^p \vee \hat{\phi}^p\right) + p\left(q; \phi^s \wedge \hat{\phi}^s, \phi^p \wedge \hat{\phi}^p\right)$$

due to the supermodularity of $p(\cdot; \cdot, \cdot)$ and, therefore,

$$p(q; \phi^s, \phi^p) - p\left(q; \phi^s \wedge \hat{\phi}^s, \phi^p \wedge \hat{\phi}^p\right) \leq p\left(q'; \phi^s \vee \hat{\phi}^s, \phi^p \vee \hat{\phi}^p\right) - p\left(q'; \hat{\phi}^s, \hat{\phi}^p\right).$$

Since the function $p(\cdot; \phi^p, \phi^s)$ is increasing in ϕ^p and ϕ^s, then the two sides of the above inequality are positive and, therefore,

$$q\left[p(q; \phi^s, \phi^p) - p\left(q; \phi^s \wedge \hat{\phi}^s, \phi^p \wedge \hat{\phi}^p\right)\right] \leq q'\left[p\left(q'; \phi^s \vee \hat{\phi}^s, \phi^p \vee \hat{\phi}^p\right) - p\left(q'; \hat{\phi}^s, \hat{\phi}^p\right)\right],$$

and, as a result,

$$
\begin{aligned}
qp\left(q;\phi^s,\phi^p\right)+q'p\left(q';\hat{\phi}^s,\hat{\phi}^p\right) \leq & q'p\left(q';\phi^s\vee\hat{\phi}^s,\phi^p\vee\hat{\phi}^p\right)+qp\left(q;\phi^s\wedge\hat{\phi}^s,\phi^p\wedge\hat{\phi}^p\right) \\
qp\left(q;\phi^s,\phi^p\right)+q'p\left(q';\hat{\phi}^s,\hat{\phi}^p\right)-c^s\left(\phi^s+\hat{\phi}^s\right) \leq & q'p\left(q';\phi^s\vee\hat{\phi}^s,\phi^p\vee\hat{\phi}^p\right)+qp\left(q;\phi^s\wedge\hat{\phi}^s,\phi^p\wedge\hat{\phi}^p\right) \\
& -c^s\left(\phi^s+\hat{\phi}^s\right) \\
\pi\left(q;\phi^s,\phi^p\right)+\pi\left(q';\hat{\phi}^s,\hat{\phi}^p\right) \leq & \pi\left(q;\phi^s\vee\hat{\phi}^s,\phi^p\vee\hat{\phi}^p\right)+\pi\left(q';\phi^s\wedge\hat{\phi}^s,\phi^p\wedge\hat{\phi}^p\right),
\end{aligned}
$$

which establishes that $\pi\left(q;\phi^s,\phi^p\right)$ is supermodular in q,ϕ^s,ϕ^p. Moreover, the correspondence defined by

$$
\Gamma\left(\eta,\gamma\right)=\begin{cases}\{0\} & \eta+\gamma<\mu^s \\ \{0,1\} & \eta+\gamma\geq\mu^s\end{cases}
$$

is obviously monotone in η and, as a result, supermodular in $\left(\eta,q,\phi^p\right)$, and this is the case since Γ is independent of $\left(q,\phi^p\right)$. Since $\hat{\pi}\left(q,\eta,\gamma,\phi^p\right)=\max_{\phi^s\in\Gamma(\eta,\gamma)}\pi\left(q;\phi^s,\phi^p\right)$, three applications of Lemma 2 implies that $\hat{\pi}\left(q,\eta,\gamma,\phi^p\right)$ is supermodular in $\left(q,\eta\right)$, $\left(q,\phi^p\right)$ and $\left(\eta,\phi^p\right)$.

Given the supermodularity of $\hat{\pi}$ and the fact that by its definition, $\tilde{\phi}^p\left(\eta,\gamma,\mathbf{q}_-\right)$ is monotone in η, for a pair $q'>q,\eta'>\eta$, we have

$$
\begin{aligned}
\hat{\pi}\left(q',\eta,\gamma,\tilde{\phi}^p\left(\eta,\gamma,\mathbf{q}_-\right)\right)-\hat{\pi}\left(q,\eta,\gamma,\tilde{\phi}^p\left(\eta,\gamma,\mathbf{q}_-\right)\right) & \leq\hat{\pi}\left(q',\eta,\gamma,\tilde{\phi}^p\left(\eta',\gamma,\mathbf{q}_-\right)\right)-\hat{\pi}\left(q,\eta,\gamma,\tilde{\phi}^p\left(\eta',\gamma,\mathbf{q}_-\right)\right) \\
& \leq\hat{\pi}\left(q',\eta',\gamma,\tilde{\phi}^p\left(\eta',\gamma,\mathbf{q}_-\right)\right)-\hat{\pi}\left(q,\eta',\gamma,\tilde{\phi}^p\left(\eta',\gamma,\mathbf{q}_-\right)\right).
\end{aligned}
$$

In other words, the function $\hat{\pi}\left(q,\eta,\gamma,\tilde{\phi}^p\left(\eta,\gamma,\mathbf{q}_-\right)\right)=\tilde{\pi}\left(\eta,\gamma,q,\mathbf{q}_-\right)$ is supermodular in $\left(q,\eta\right)$. This, together with the fact that $F\left(\cdot|\tilde{q}\right)$ is increasing in $\tilde{q}$ in the sense of first order stochastic dominance, in turn implies that $\int\tilde{\pi}\left(\eta,\gamma,q,\mathbf{q}_-\right)dF\left(q|\tilde{q}\right)$ is supermodular in $\left(\tilde{q},\eta\right)$. To see this note that if $\eta' > \eta$ and $\tilde{q}' > \tilde{q}$, then

$$\int \tilde{\pi}\left(\eta', \gamma, q, \mathbf{q}_{-}\right) dF\left(q|\tilde{q}\right) - \int \tilde{\pi}\left(\eta, \gamma, q, \mathbf{q}_{-}\right) dF\left(q|\tilde{q}\right) = \int \left[\tilde{\pi}\left(\eta', \gamma, q, \mathbf{q}_{-}\right) - \tilde{\pi}\left(\eta, \gamma, q, \mathbf{q}_{-}\right)\right] dF\left(q|\tilde{q}\right).$$

By supermodularity of $\tilde{\pi}$ in $(\eta, \tilde{q})$, $\tilde{\pi}\left(\eta', \gamma, q, \mathbf{q}_{-}\right) - \tilde{\pi}\left(\eta, \gamma, q, \mathbf{q}_{-}\right)$ must be increasing in q and, therefore, by applying the definition of first order stochastic dominance:

$$\int \left[\tilde{\pi}\left(\eta', \gamma, q, \mathbf{q}_{-}\right) - \tilde{\pi}\left(\eta, \gamma, q, \mathbf{q}_{-}\right)\right] dF\left(q|\tilde{q}'\right) \geq \int \left[\tilde{\pi}\left(\eta', \gamma, q, \mathbf{q}_{-}\right) - \tilde{\pi}\left(\eta, \gamma, q, \mathbf{q}_{-}\right)\right] dF\left(q|\tilde{q}\right),$$

which implies the supermodularity of $\int \hat{\pi} dF$ in $(\eta, \tilde{q})$.

Supermodularity of $\int \hat{\pi} dF$ with respect to (η, q_{-i}): Since $\int \hat{\pi} dF$ is symmetric with respect to q_{-i} for all $i = 1, 2, 3$, it is sufficient to focus on only one past quantity, say q_{-1}. By application of Lemma 2, we know that $\hat{\pi}$ is supermodular in (q, ϕ^p). Since $\tilde{\phi}^p$ is increasing in q_{-1}, this implies that $\tilde{\pi}$ is supermodular in (η, q_{-1}).

Supermodularity of $\int \hat{\pi} dF$ with respect to $(q_{-i}, \tilde{q})$: This property is straightforward, as it follows a similar argument as above. In particular, an increase in q_{-i} could lead to Powerseller status for the seller and thus a higher price. The proof follows because $\hat{\pi}$ is supermodular in (q, ϕ^p).

Step 2. Here I show that the solution to the functional equation above is supermodular. To do so, since the set of continuous supermodular functions is closed, it is sufficient to show that the transformation associated with the Bellman equation preserves supermodularity. That is, for any function $v\left(\eta, \gamma, \mathbf{q}_{-}\right)$ that is supermodular in (η, q_{-i}), the following function is also supermodular in (η, q_{-i}):

$$\hat{v}\left(\eta, \gamma, \mathbf{q}_{-}\right) = \max_{\tilde{q}} \int \left[\tilde{\pi}\left(\eta, \gamma, q, \mathbf{q}_{-}\right) + \beta \int v\left(\eta, \gamma', \left(q, q_{-1}, q_{-2}\right)\right) g\left(\gamma\right) d\gamma\right] dF\left(q|\tilde{q}\right).$$

To show this, note that the function

$$\tilde{v}(\eta, \gamma, q, \mathbf{q}_-) = \tilde{\pi}(\eta, \gamma, q, \mathbf{q}_-) + \beta \int v\left(\eta, \gamma', (q, q_{-1}, q_{-2})\right) g(\gamma)\, d\gamma$$

is supermodular in $(\tilde{q}, q_{-i}), (\tilde{q}, \eta), (\eta, q_{-i})$ for all $i = 1, 2, 3$. Therefore, by Lemma 1 in Hopenhayn and Prescott [1992], the function $\hat{v}(\eta, \gamma, \mathbf{q}_-)$ is also supermodular. This concludes Step 2.

Step 3. Given the above steps, the objective function in the Bellman equation is supermodular in (η, q) and (η, q_{-i}). Now suppose to the contrary of the proposition that there exists $\eta' > \eta$ such that the optimal solution under $(\eta', \gamma, \mathbf{q}_-)$, q', is lower than the optimal solution under $(\eta, \gamma, \mathbf{q}_-)$, q. Given $\gamma, \mathbf{q}_-$, define the following function:

$$f(\eta, q) = \hat{\pi}(\eta, \gamma, q, \mathbf{q}_-) + \beta \int v\left(\eta, \gamma', (q, q_{-1}, q_{-2})\right) g(\gamma)\, d\gamma,$$

which is supermodular in (η, q). Hence,

$$f(\eta, q) - f\left(\eta, q'\right) \leq f\left(\eta', q\right) - f\left(\eta', q'\right).$$

By optimality of q under η and uniqueness of the policy function, the LHS of the above inequality is positive; hence, the RHS is also positive. This contradicts the fact that q' is optimal under η'. Hence, the policy function $q^*(\eta, \gamma, \mathbf{q}_-)$ must be increasing in η. Similarly, I can show that it is increasing in q_{-i}. □

Appendix B Demand Function Estimation Robustness

As mentioned in Section 2, I estimate a structural demand function based on buyers' utility function. In this section, I run a simple OLS regression of price over additional characteristics of sellers, and characteristics of items sold by them, to show the robustness of the results when it comes to the effect of Powerseller and store status. The results in this section show that when I control for sellers with a high level of sales, I still see the positive effect of Powerseller and store status. Moreover, when I control for the condition of the items sold (i.e., if they are new or used), I see that Powerseller and store statuses still have a positive effect, with a higher effect when looking only at used items.

Table 8 reports the OLS results. The first column includes only the seller characteristics. In addition to Powerseller and store statuses, I also include other sellers' characteristics: number of days a seller has been active in the market, which I will call "age", the amount of information entered by sellers on the listing page, whether sellers have provided their phone number in their listing page, the existence of an "About-me" page,[31] and whether the listing was in a fixed price format (i.e., Buy it Now).

Table 8 also shows that being a Powerseller or a store on eBay has a positive effect on price. The coefficient of the age variable shows that being on eBay for one additional year gives a seller about a 3-dollar boost in the final price. Additionally, having more text on the listing page has a positive effect on price.[32] The About-me coefficient has a negative effect on price, because the option of having an About-me webpage was more popular during the starting days of eBay. However, iPods are a newer subcategory on eBay, and most of the big sellers in this category are newer sellers. Therefore, the coefficient of the About-me variable picks up the effect of older versus newer sellers.

Column II represents the coefficients when I consider only the characteristics of the items sold on eBay. As expected, if the condition of the iPod is new or refurbished, it results in a price premium. High internal memory of iPods also results in higher prices. I also add dummy variables for different brands of iPods, which also have the expected coefficients.

Column III of Table 8 includes both seller and item characteristics. The effect of Powerseller and store statuses is lower compared to the results in Column I. This shows that Powersellers and registered stores tend to sell better-quality products; when I control for item characteristics, the

[31] Sellers can enter a webpage called About-me and explain their business on this page for buyers to see.

[32] Note that the two variables, text and description size, represent different measures of information entered on the webpage. They are highly correlated, and having only one of them in the regression results in a positive coefficient.

effect of Powerseller and store statuses diminishes. However, the effect of these reputation-related variables is still very high; the premium on Powerseller status is \$29, which is about 15% of the average price of iPods sold. The premium on store status is about \$8.6, which is about 5% of the average price of iPods sold.

Column IV represents only sellers with more than 25 sales in my sample. The effect of store and Powerseller status declines when I focus only on this sample of data, because this is a pool of sellers with a higher volume of sales, and thus more experienced. Therefore, the signal for these sellers is less important than for smaller sellers with a lower volume of sales.

Buyers take the reputation of sellers more into account when they are buying an item with a less predetermined value (i.e., used goods versus new goods). Table 9 shows the regression results for used versus new items. Powerseller and store statuses have remarkably higher effects for a used item versus a new item. The market value of a new iPod is predetermined. In this case, buyers may be more confident to buy from a more trustworthy seller because they expect a better shipping experience and better communications, or, in the extreme case, because they are afraid of receiving a used iPod instead of a new iPod from a less reputable seller. On the other hand, when a seller is buying a used iPod, there are many aspects of the item quality that can be misrepresented by a fraudulent seller, leading to a very high value of reputation for sellers.

In the last column of Table 9, I include feedback scores and feedback percentages in the regressors in the third column. After the end of a transaction, seller and buyer can leave feedback for each other, and this can be positive, negative, or neutral. Feedback percentage is the percentage of positive feedback among all feedback ratings that a seller has received. Feedback score is the number of positive feedback ratings received minus the number of negative feedback ratings received by a seller. Many of the papers on the effects of eBay's reputation mechanism focus only on sellers' feedback scores and feedback percentages. This regression shows that, controlling for Powerseller and store status, these two variables do not have a high effect on the final price. Feedback percentage is a number between 0 and 100, with an average of 99% for the active sellers on the market. When I compare a seller with a perfect feedback percentage (100%) and a seller in the 25% percentile (98% feedback percentage), the effect of feedback percentage on price is $0.75. The coefficient on feedback score is negative when I control for sellers' size, showing that this coefficient does not carry additional information for buyers in terms of reputation.

List of Figures

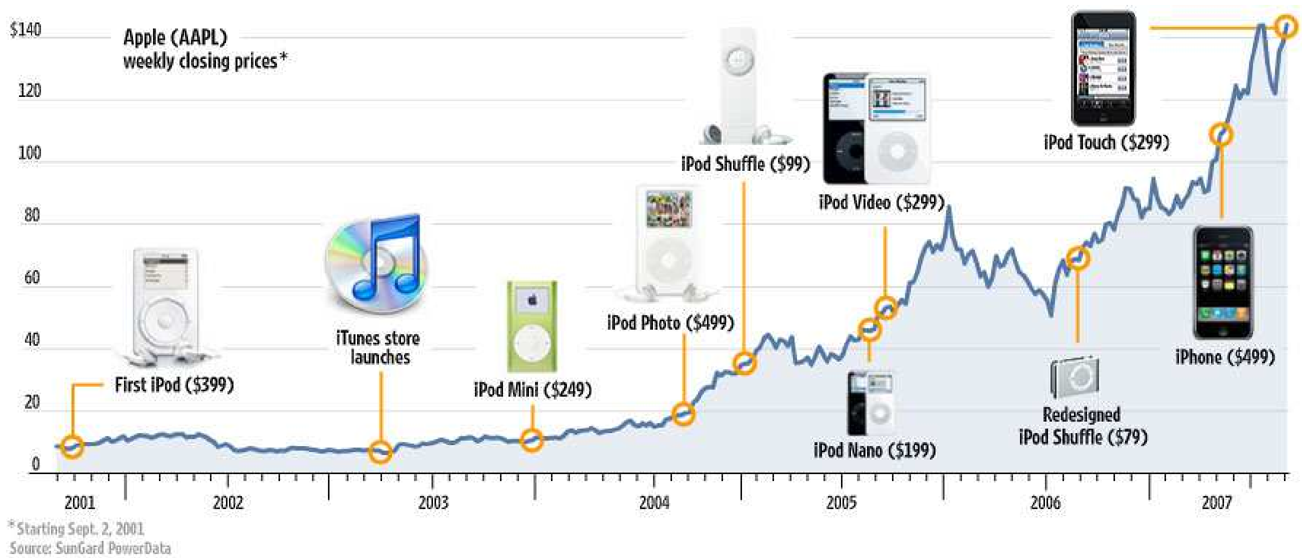

Figure 1: Different models of iPods and their prices over time.

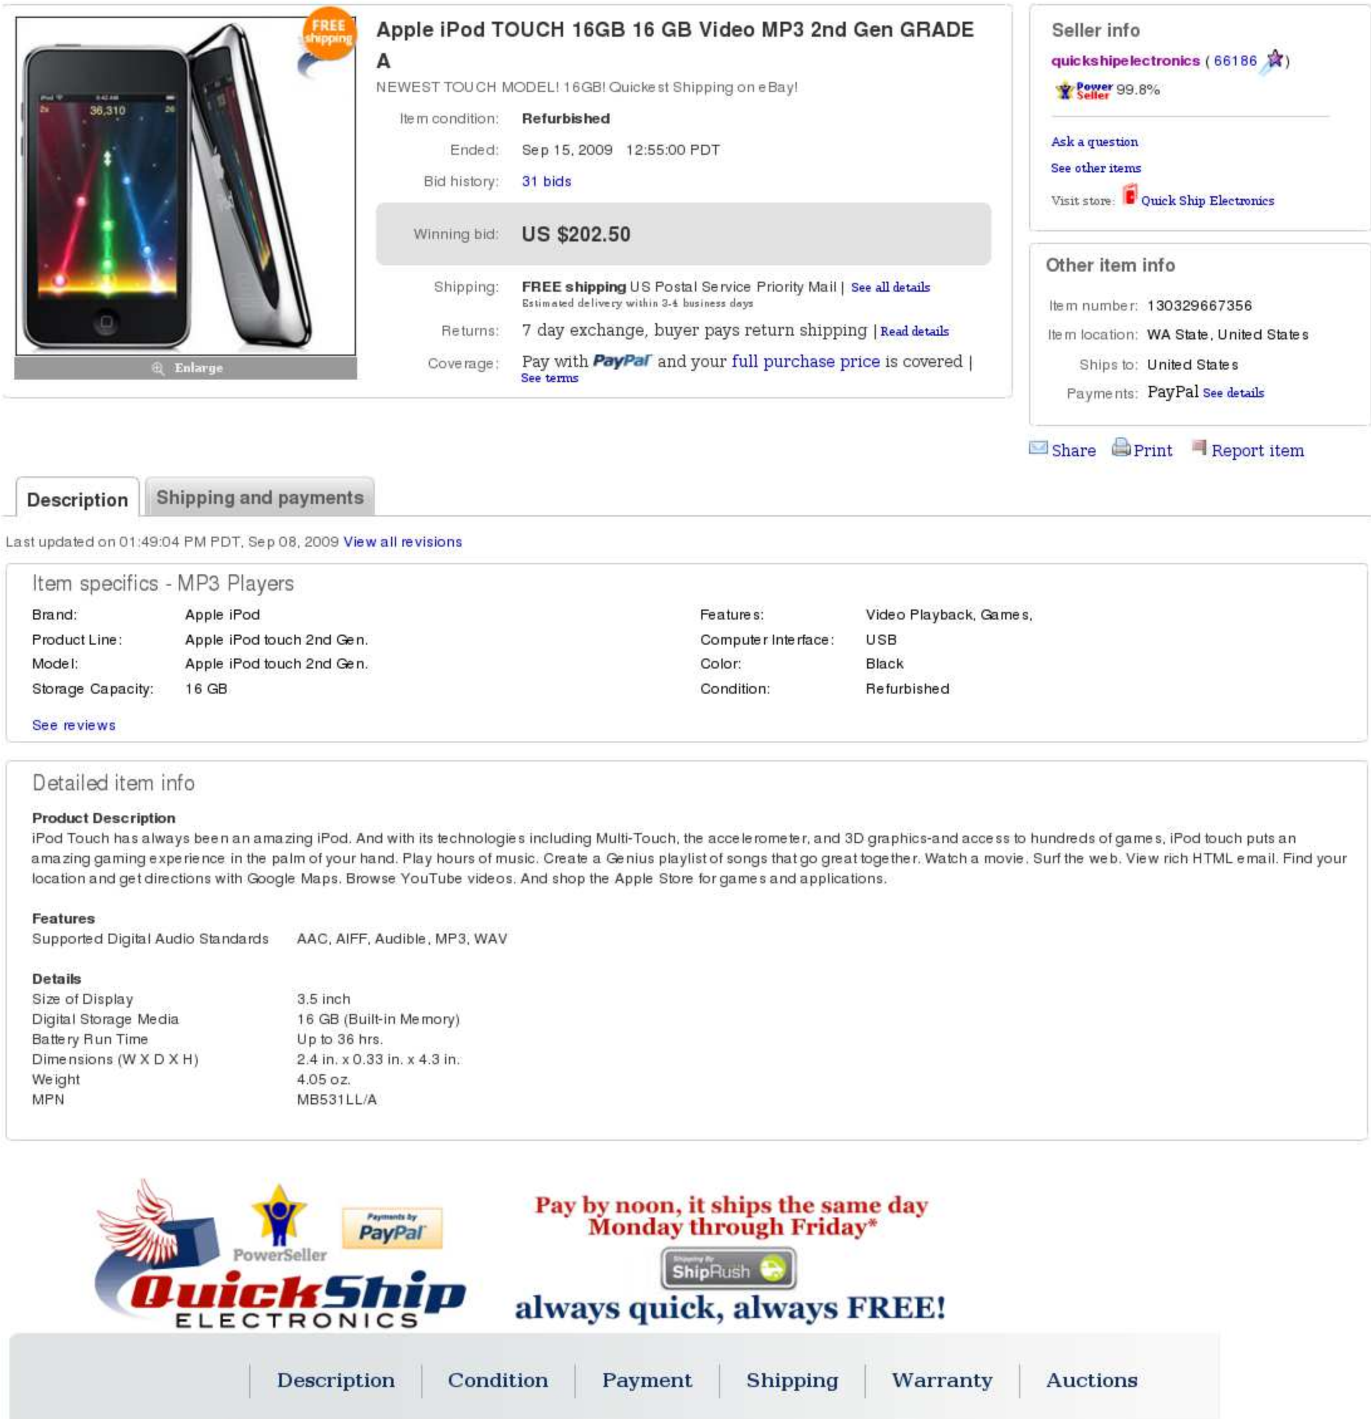

Figure 2: Snapshot of an iPod Auction

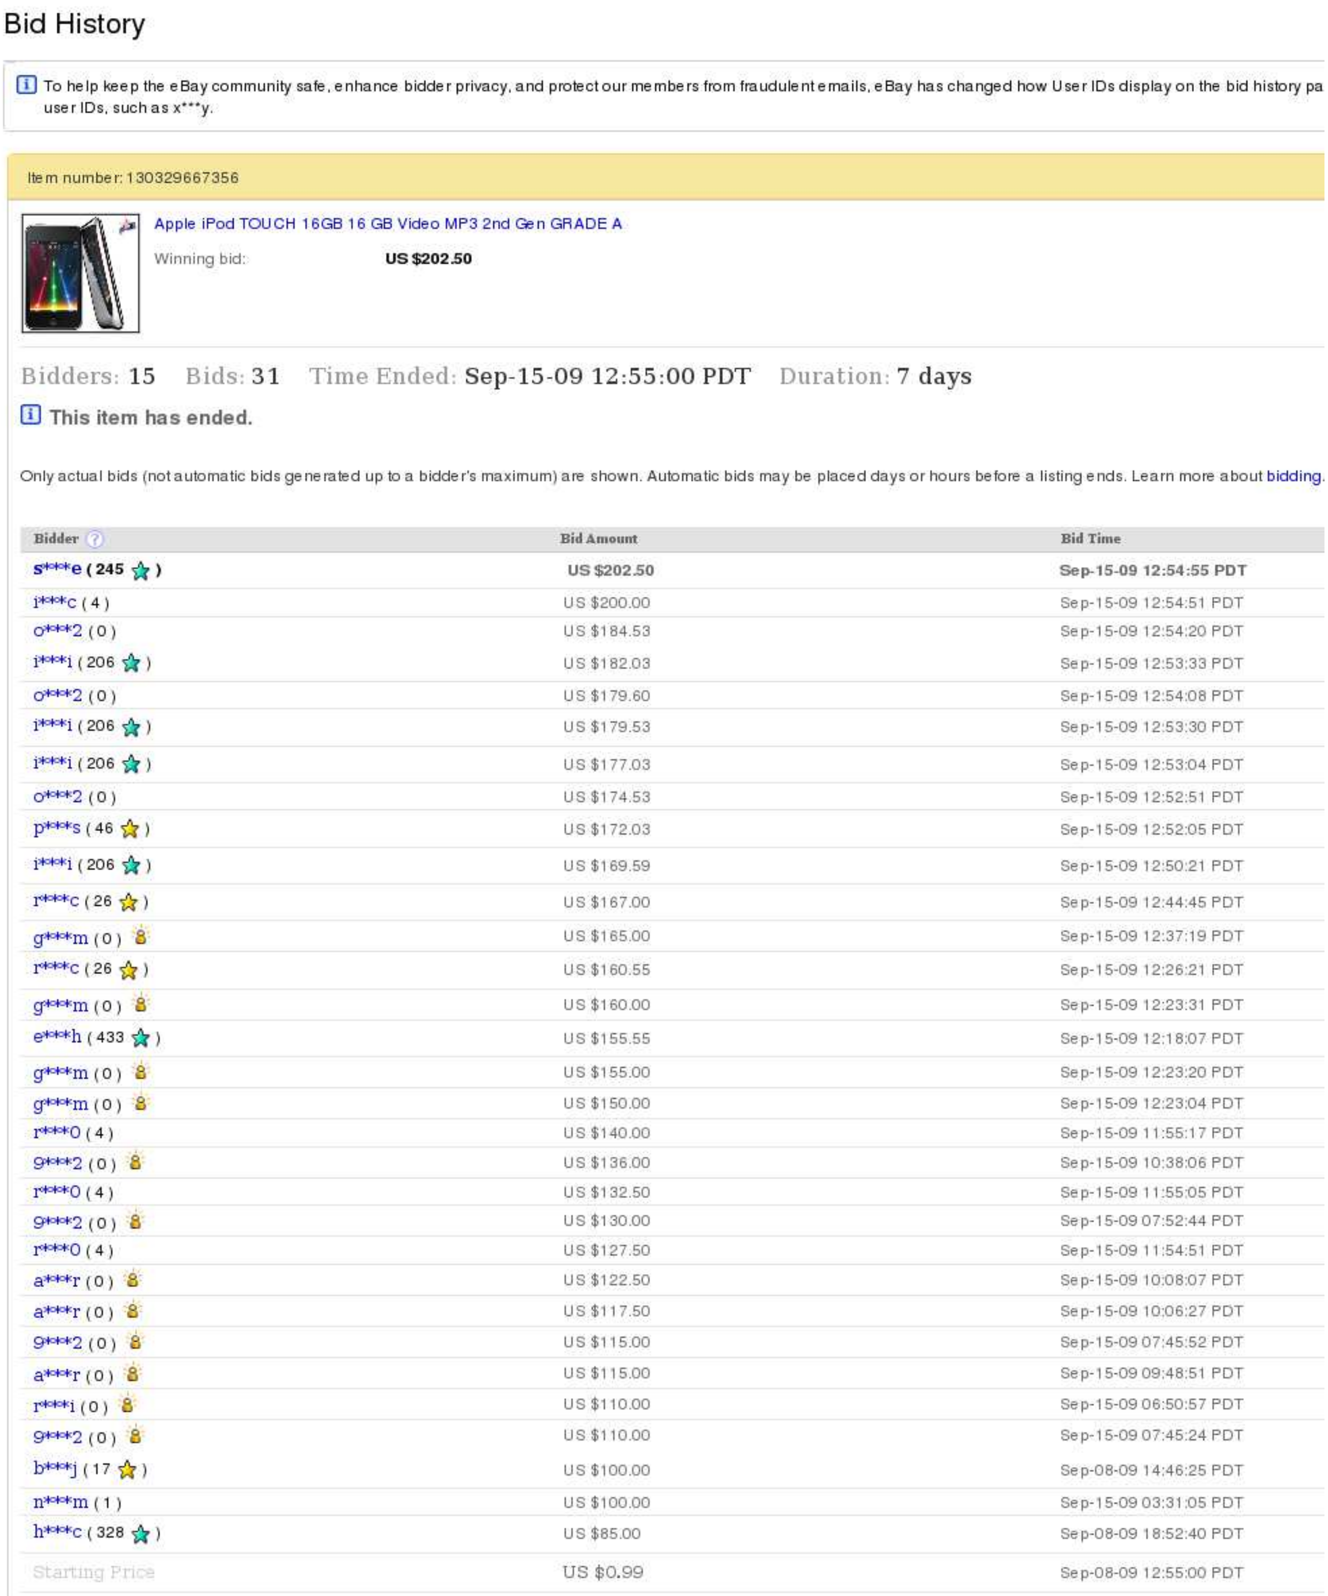
Bid History

To help keep the eBay community safe, enhance bidder privacy, and protect our members from fraudulent emails, eBay has changed how User IDs display on the bid history pa
user IDs, such as x***y.

Item number: 130329667356

Apple iPod TOUCH 16GB 16 GB Video MP3 2nd Gen GRADE A

Winning bid: **US $202.50**

Bidders: **15** Bids: **31** Time Ended: **Sep-15-09 12:55:00 PDT** Duration: **7 days**

This item has ended.

Only actual bids (not automatic bids generated up to a bidder's maximum) are shown. Automatic bids may be placed days or hours before a listing ends. Learn more about bidding.

Bidder	Bid Amount	Bid Time
s*e (245)**	**US $202.50**	**Sep-15-09 12:54:55 PDT**
i***c (4)	US $200.00	Sep-15-09 12:54:51 PDT
o***2 (0)	US $184.53	Sep-15-09 12:54:20 PDT
i***i (206)	US $182.03	Sep-15-09 12:53:33 PDT
o***2 (0)	US $179.60	Sep-15-09 12:54:08 PDT
i***i (206)	US $179.53	Sep-15-09 12:53:30 PDT
i***i (206)	US $177.03	Sep-15-09 12:53:04 PDT
o***2 (0)	US $174.53	Sep-15-09 12:52:51 PDT
p***s (46)	US $172.03	Sep-15-09 12:52:05 PDT
i***i (206)	US $169.59	Sep-15-09 12:50:21 PDT
r***c (26)	US $167.00	Sep-15-09 12:44:45 PDT
g***m (0)	US $165.00	Sep-15-09 12:37:19 PDT
r***c (26)	US $160.55	Sep-15-09 12:26:21 PDT
g***m (0)	US $160.00	Sep-15-09 12:23:31 PDT
e***h (433)	US $155.55	Sep-15-09 12:18:07 PDT
g***m (0)	US $155.00	Sep-15-09 12:23:20 PDT
g***m (0)	US $150.00	Sep-15-09 12:23:04 PDT
r***0 (4)	US $140.00	Sep-15-09 11:55:17 PDT
9***2 (0)	US $136.00	Sep-15-09 10:38:06 PDT
r***0 (4)	US $132.50	Sep-15-09 11:55:05 PDT
9***2 (0)	US $130.00	Sep-15-09 07:52:44 PDT
r***0 (4)	US $127.50	Sep-15-09 11:54:51 PDT
a***r (0)	US $122.50	Sep-15-09 10:08:07 PDT
a***r (0)	US $117.50	Sep-15-09 10:06:27 PDT
9***2 (0)	US $115.00	Sep-15-09 07:45:52 PDT
a***r (0)	US $115.00	Sep-15-09 09:48:51 PDT
r***i (0)	US $110.00	Sep-15-09 06:50:57 PDT
9***2 (0)	US $110.00	Sep-15-09 07:45:24 PDT
b***j (17)	US $100.00	Sep-08-09 14:46:25 PDT
n***m (1)	US $100.00	Sep-15-09 03:31:05 PDT
h***c (328)	US $85.00	Sep-08-09 18:52:40 PDT
Starting Price	US $0.99	Sep-08-09 12:55:00 PDT

Figure 3: Snapshot of Bid History Page

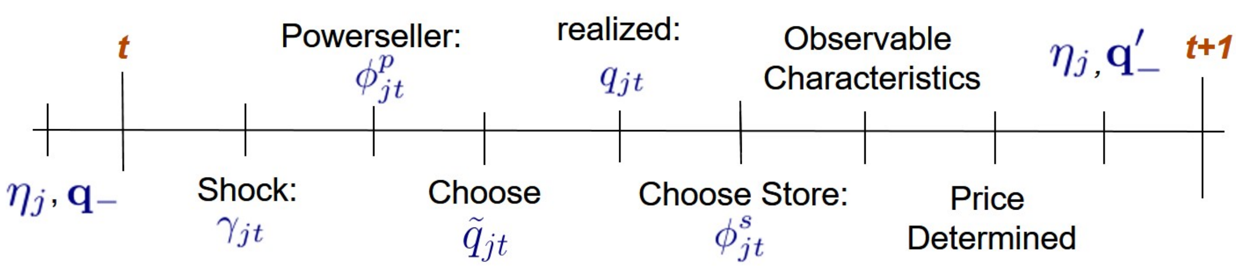

Figure 4: Timing of Sellers' Choices and Shocks

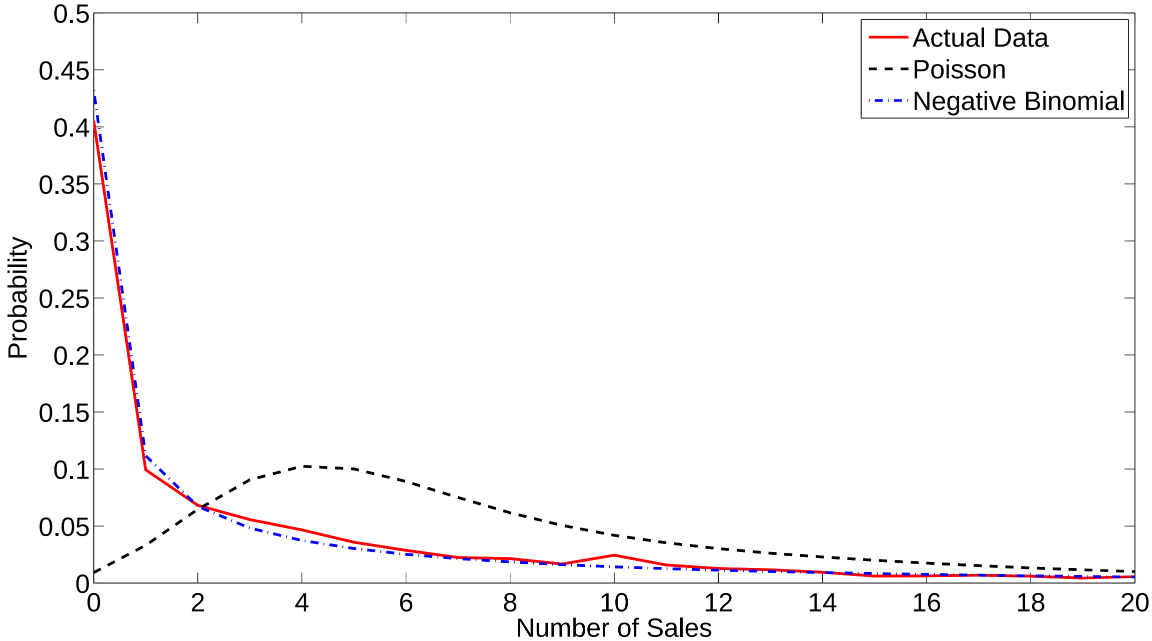

Figure 5: Probability Distribution of the Number of Sales, Original vs. Fitted Negative-Binomial

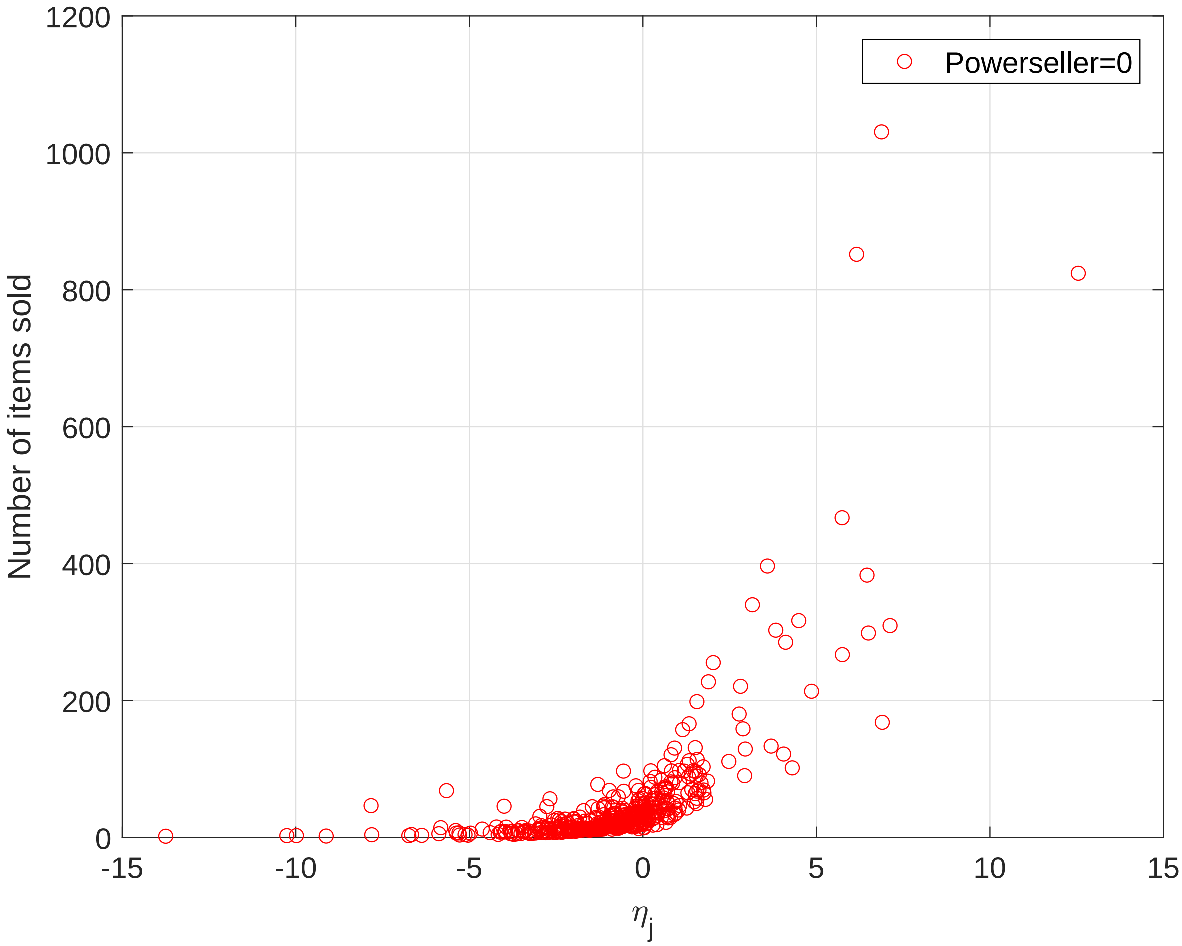

Figure 6: Total Quantity Sold by Non-Powersellers of Different Quality

Notes: The x-axis shows the persistent level of quality of sellers.

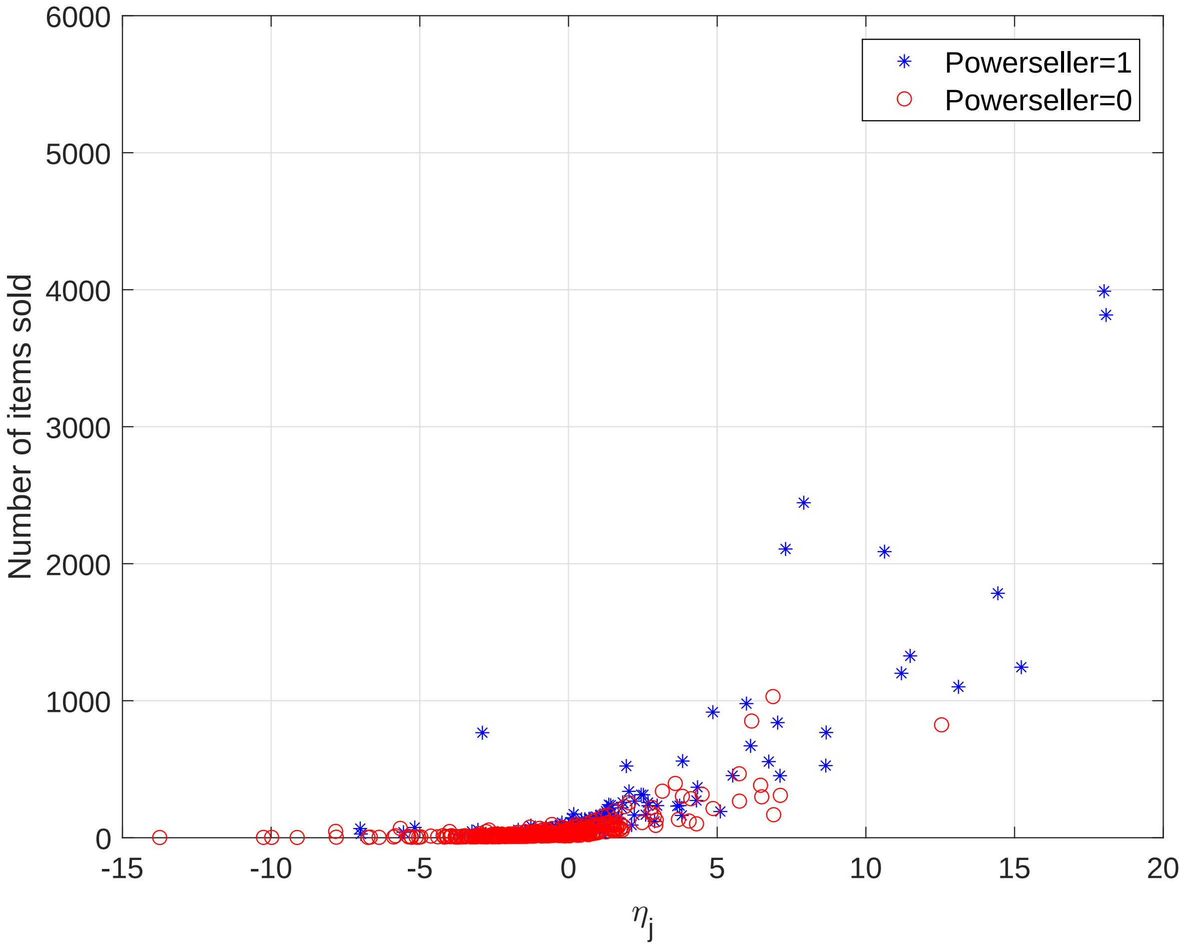

Figure 7: Total Quantity Sold by Both Powersellers and Non-Powersellers

Notes: The x-axis shows the persistent level of quality of sellers.

List of Tables

Table 1: Data Summary

Characteristics of Listings and iPods Sold

Variable	Obs.	Mean	Std. Dev.	Min	Max
eBay registered store	174,280	0.36	0.48	0	1
Powerseller	174,280	0.48	0.50	0	1
Feedback Number	174,154	14,120.3	48,971.8	-3	1,026,575
Feedback Percentage	22,366	99.22	1.88	33.3	100
Sold with Buy it Now	174,273	0.08	0.27	0	1
Buy it Now option	174,280	0.29	0.45	0	1
Secret Reserve	174,280	0.04	0.27	0	2
Number of Bidders	146,597	7.29	4.82	0	30
Items Sold	167,199	1.00	1.84	0	180
New Item	174,280	0.25	0.43	0	1
Refurbished Item	174,280	0.19	0.40	0	1
Internal Memory	159,234	19.68	27.51	1	240

Table 2: Reputation and Price

	Average Prices		Fitted Values	
	All iPods	New iPod Nano	Average Item	New, Nano, 8GB
All Sellers	$131.81	$132.95	$136.51	$135.34
Non-Powersellers and Non-Store	$130.70	$130.15	$122.18	$131.19
Registered Stores	$135.96	$134.09	$128.80	$139.96
Powersellers	$134.95	$137.44	$137.79	$140.90
Powersellers and Stores	$139.90	$135.29	$145.35	$142.09
Lowest 25% Feedback	$134.60		$135.34	
Highest 25% Feedback	$136.68		$135.51	

Note: The first two columns represent the average price of items sold by each group of sellers. The third and fourth columns show the fitted values for the price of items sold by each group. The third column shows the average characteristics of an iPod sold during the time period, and the fourth column shows the price of a new iPod Nano sold by each group.

Table 3: First-Stage Estimation, Demand

Price		
	Coef.	Std. Dev.
$log(s_0) - log(s_j)$	2.14	0.09
Powerseller	4.77	0.41
Registered Store	1.44	0.60
Powerseller* Store	6.66	0.68
New	39.45	0.34
Refurbished	14.20	0.34
Internal Memory	1.49	0.01
iPod Model Fixed Effect	✓	
Monthly Fixed Effect	✓	
R^2	0.95	

Table 4: First-Stage Estimation, Policy Functions

	Coef.	Std. Dev.	
Quantity Choice	Registered Store	0.65	0.34
	Powerseller	0.33	0.15
	q_{-1}	0.003	0.0007
	q_{-2}	-0.001	0.0004
	Dispersion	0.90	0.03
Registered Store	Powerseller	1.54	0.10
	q_{-1}	0.013	0.002
	q_{-2}	0.008	0.001
	Fixed Effect	-0.37	0.04
	Constant	-2.33	0.10

Note: The first part of the table is the coefficients for negative binomial regression. This regression also includes fixed effects for sellers. The second part is the estimation of a binary logit model for store status choice for sellers.

Table 5: Parametric Estimation Unobserved Quality

Effect of Quality on Price

	Parameter	Std. Dev.
λ	0.38	0.04
β_r/α	4.72	0.19

Table 6: Cost Estimations

	Specifications			
	I		II	
	Coef.	Std. Dev.	Coef.	Std. Dev.
iPod's Cost	132.72	0.51	134.08	0.48
Registered Store	–	–	44.40	10.59

Table 7: Change in Per Period Consumer Surplus, Sellers' Profit, and eBay's Profit

	Original	No Reputation	Warranty 5%	10%	50%
Total Consumers' Surplus	116,070	75,593	43,570	97,168	182,740
Total Sellers' Profit	60,037	20,633	48,631	114,10	285,090
eBay's Profit	107,440	67,082	32,940	89,490	211,470

Notes: The numbers in this table are for the surpluses in the last period where all sellers are assumed to be active.

Table 8: Regression Result for iPod

	Price			
	I	II	III	IV
Powerseller	80.04		29.26	9.29
	(0.75)		(0.81)	(0.31)
Registered Store	40.67		8.62	4.31
	(0.65)		(0.42)	(0.36)
Age	0.01		0.008	0.005
	(0.00)		(0.0002)	(0.0001)
Phone	21.19		0.68	-5.39
	(0.72)		(0.50)	(0.40)
Text	-0.003		-0.001	-0.0004
	(8.0E-05)		(4.3E-05)	(4E-05)
Description	0.001		0.0004	0.0002
	(2.4E-05)		(1.4E-05)	(1.2E-05)
About Me	-14.89		-15.07	-5.69
	(0.91)		(0.53)	(0.37)
Buy it Now	26.20		36.62	5.38
	(3.26)		(2.09)	(0.54)
New		31.02	29.43	48.27
		(0.52)	(0.55)	(0.34)
Refurbished		11.04	3.32	12.42
		(0.39)	(0.45)	(0.32)
Internal Memory		1.43	1.40	1.41
		(0.02)	(0.02)	(0.008)
Nano		87.72	46.16	64.89
		(0.34)	(1.05)	(0.30)
Mini		52.02	3.62	34.02
		(0.60)	(1.25)	(0.46)
Classic		44.33	2.50	24.94
		(1.80)	(1.98)	(0.70)
Shuffle		27.82	-14.37	7.07
		(0.31)	(1.05)	(0.34)
Touch		195.66	152.11	179.61
		(0.52)	(1.17)	(0.41)
Video		58.99	19.69	43.63
		(1.16)	(1.50)	(0.58)
R^2	0.72	0.93	0.94	0.92

I: Only Sellers' Characteristics
II: Only Item Characteristics,
III: Both Sellers' and Item Characteristics,
IV: Both Sellers' and Item Characteristics, Sellers > 25 Sales
Standard errors in parentheses
* p<0.05 ** p<0.01 *** p<0.001

Table 9: Regression Result for iPod, New vs. Used Items

	Price			
	Original	New Items	Used Items	Feedback
Powerseller	29.27***	6.37***	35.95***	17.41***
	(0.82)	(1.51)	(0.91)	(0.80)
Registered Store	8.62***	0.36	11.53***	15.49***
	(0.42)	(1.09)	(0.45)	(0.42)
Age	0.008***	0.01***	0.006***	0.008***
	(0.0002)	(0.0007)	(0.0002)	(0.0002)
Phone	0.68	-7.84***	5.58***	-3.95***
	(0.49)	(1.28)	(0.58)	(0.46)
Description Size	0.0004***	-0.0001*	0.0006***	0.0005***
	(0.00001)	(0.00004)	(0.00002)	(0.00001)
Text	-0.001***	0.001***	-0.002***	-0.001***
	(0.00004)	(0.0001)	(0.00004)	(0.00004)
About me	-15.07***	-1.16	-13.75***	-15.77***
	(0.53)	(1.59)	(0.55)	(0.48)
Buy it Now	36.62***	-31.29***	66.24***	24.95***
	(2.09)	(3.24)	(2.33)	(2.07)
New	29.43***			36.96***
	(0.55)			(0.54)
Refurbished	3.31***		0.51	15.13***
	(0.44)		(0.47)	(0.41)
Internal Memory	1.40***	1.55***	1.36***	1.48***
	(0.017)	(0.07)	(0.08)	(0.02)
Nano	46.16***	101.40***	41.17***	38.79***
	(1.051)	(2.67)	(1.14)	(0.99)
Mini	3.62**		-4.41**	-1.76
	(1.25)		(1.35)	(1.28)
Classic	2.50	45.35***	-0.23	-12.87***
	(1.98)	(8.16)	(2.05)	(1.98)
Shuffle	-14.37***	19.41***	-14.00***	-15.40***
	(1.06)	(2.37)	(1.15)	(0.98)
Touch	152.1***	209.0***	147.6***	147.1***
	(1.17)	(3.28)	(1.26)	(1.09)
Video	19.69***	106.8***	16.17***	15.63***
	(1.49)	(4.56)	(1.54)	(1.43)
Feedback Percentage				0.37***
				(0.006)
Feedback Score				-0.00006***
				(0.000002)
R^2	0.94	0.96	0.94	0.95

Standard errors in parentheses
* p<0.05 ** p<0.01 *** p<0.001