

NBER WORKING PAPER SERIES

SALES MECHANISMS IN ONLINE MARKETS:
WHAT HAPPENED TO INTERNET AUCTIONS?

Liran Einav
Chiara Farronato
Jonathan D. Levin
Neel Sundaresan

Working Paper 19021
<http://www.nber.org/papers/w19021>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2013

We are grateful to Pat Bajari, Preston McAfee and many seminar participants for helpful comments. We appreciate support from the National Science Foundation, the Stanford Institute for Economic Policy Research, the Toulouse Network on Information Technology, and the Alfred P. Sloan Foundation. The data for this study were obtained through a contract between the Stanford authors (Einav, Farronato, Levin) and eBay Research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w19021.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Liran Einav, Chiara Farronato, Jonathan D. Levin, and Neel Sundaresan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Sales Mechanisms in Online Markets: What Happened to Internet Auctions?

Liran Einav, Chiara Farronato, Jonathan D. Levin, and Neel Sundaesan

NBER Working Paper No. 19021

May 2013

JEL No. D02,D43,D44,D47,D82,L13,L86

ABSTRACT

Consumer auctions were very popular in the early days of internet commerce, but today online sellers mostly use posted prices. Data from eBay shows that compositional shifts in the items being sold, or the sellers offering these items, cannot account for this evolution. Instead, the returns to sellers using auctions have diminished. We develop a model to distinguish two hypotheses: a shift in buyer demand away from auctions, and general narrowing of seller margins that favors posted prices. Our estimates suggest that the former is more important. We also provide evidence on where auctions still are used, and on why some sellers may continue to use both auctions and posted prices.

Liran Einav
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
and NBER
leinav@stanford.edu

Jonathan D. Levin
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
and NBER
jdlevin@stanford.edu

Chiara Farronato
Stanford University
Department of Economics
579 Serra Mall
Stanford, CA 94305-6072
chiarafarronato@gmail.com

Neel Sundaesan
eBay Research Labs
2065 Hamilton Avenue
San Jose, CA 95125
nsundaesan@ebay.com

An online appendix is available at:
<http://www.nber.org/data-appendix/w19021>

1 Introduction

In the early days of the internet, many observers speculated that technology would shift retail markets in the direction of more dynamic and flexible pricing mechanisms. The *Economist* (2000) wrote that the internet had introduced “the possibility of a permanent worldwide bazaar in which no prices are ever fixed for long, all information is instantly available, and buyers and sellers spend their lives haggling to try to get the best deals” (see also Hall, 2002). The best example is eBay, which by 2001 had become a dominant platform for consumer auctions, and was the third-ranked web site in terms of time spent by consumers.¹ Its growth was enabled by the development of proxy bidding, which allowed buyers to submit a maximum bid and have the computer respond to opponent bids up to this maximum, thereby lowering transaction costs because bidders did not need to be physically present or even paying attention in real time in order to participate in a dynamic auction.

Since this time, online commerce has grown enormously but internet auctions have not. Today most online commerce takes place at posted retail prices, as has been the case offline for at least a hundred years (Surowiecki, 2011). Figure 1 shows this evolution on eBay: the share of listings and transaction volume attributable to auction sales has fallen well below fifty percent.² Figure 2 shows a similar pattern, this time for Google searches involving the terms “online auctions” and “online prices.” In this paper, we explore this evolution, which we view as interesting not just because of the fast growth of online commerce, but because the choice of sales mechanism, and competition between sales mechanisms, is one of the classic problems in microeconomic theory.

We combine a simple model with rich data from eBay to tease apart competing explanations. One immediate hypothesis is that there has been a shift in the composition of internet sellers, or of items being sold online. We show that, at least on eBay, this does not appear to explain much of the move to posted prices. Instead, the shift has occurred *within* natural groupings of sellers and products. A second possibility is a change in consumer preferences.

¹Source: <http://blog.compete.com/2007/10/01/top-ranked-web-sites-popularity-2001/>.

²The sharp drop in the fall of 2008 partly reflects a change in eBay policy that made posted price listings cheaper and allowed posted price listings to remain active for longer periods. The figure omits a variety of other sales mechanisms such as “hybrid” auctions that allow a preemptive posted price purchase, or posted prices with an option to make a lower offer. These represent a smaller share of listings and transactions.

Ten years ago, internet auctions were a form of online entertainment. Today, YouTube, Facebook, and other online diversions may have increased the demand for convenient shopping. Finally, the price discovery benefits of auctions have declined over time: either because online search has made it easier to find comparable prices, or (relatedly) because increased competition has narrowed seller margins.

We start in Section 2 by documenting some empirical facts about the decline in auctions, and about where and when they are currently used. Many of the patterns are expected: auctions are favored for used and idiosyncratic items, and by less experienced sellers. Others are more surprising. Compositional changes do not explain the decline in auctions. And there are only small differences in the selling behavior of different cohorts of sellers. The latter two facts suggest that changes in seller incentives, rather than seller composition, have been important. These patterns together also provide a starting point for a theoretical model.

In Section 3, we develop a model that captures the key trade-off between the price discovery that results from running an auction, and the convenience of using a posted price. The model is parsimonious. It has just a few parameters, which capture how reduced uncertainty about an item’s value, greater retail competition, and greater demand for convenience all favor posted prices. Nevertheless, it matches the basic patterns in the data, including a robust pattern that auction listings tend to have higher sale probabilities but lower prices.

In Section 4, we analyze the returns to auction and posted price selling at different points in time. The main difficulty is accounting for differences in the types of items sold by auction and posted price, and the types of sellers doing the selling. Items listed by auction can be quite different than items listed by posted price, even after controlling for their measurable characteristics. Potential confounders include the quality of the free-text item description, the attractiveness of posted pictures, and the identity of the seller. Moreover, except for a limited set of products with catalogue identifiers, we lack information on product characteristics apart from the seller-chosen listing title, and a somewhat coarse category classification.

To make progress, we adopt a strategy used by Elfenbein et al. (2012, 2013) and Einav et al. (2013). We take advantage of a ubiquitous feature of eBay’s platform, namely that

sellers frequently list the same item, either simultaneously or over time, while varying their sale format or other pricing parameters. We construct a large dataset of matched listings spanning from 2003 (when auctions were the dominant sales mechanism) to 2009 (by which time posted prices had overtaken auctions). We focus on cases where a seller listed the same item by auction and by posted price, enabling an “apples-to-apples” comparison.

The comparison reveals the growth of a large “auction discount”. In 2003, auction prices were on average within five percent of matched posted price sales. By 2009, the discount had grown to over sixteen percent. The drop cannot be explained by sellers lowering their auction reserve prices. In fact, reserve prices generally increased relative to matched posted prices. In addition, we show that success rates for both formats fell considerably, but not differentially, over this time period.

In Section 5, we connect these findings to seller incentives. To do this, we need to account for the flexibility that sellers have to adjust their posted prices and/or reserve prices. We use variation in these prices, within sets of matched listings, to estimate the combinations of sale probability and sale price that can be achieved with different mechanisms. The estimates can be interpreted as “listing-level” demand curves. Comparing the estimates from 2003 and 2009 shows that listing-level demand has fallen for both formats, but much more for auctions. The qualitative features of the estimates also align with our theoretical model.

We therefore decompose the time-series change into a general fall in listing-level demand, and a specific fall in auction demand, and use the model to back out the effects on seller profitability under auctions and price posting. We find that most of the impetus toward price posting has arisen from the relative fall in auction demand, rather than the overall increase in listing competition. This has been particularly true in categories such as Collectibles, Jewelry and Clothing that feature relatively differentiated products. The increase in listing competition seems to have played a bigger role in categories such as Electronics and Computers, which have seen a similar shift toward posted pricing.

In the final sections of the paper, we provide some additional evidence on the changing nature of buyer demand. We first show that compositional changes in the set of buyers do not appear to have driven the increased demand for posted prices. However, at any point in time, there is significant heterogeneity in the pool of buyers. Buyers who do a lot of

shopping on eBay tend to focus more on auctions, and get better bargains when they are bidding. This suggests that customer segmentation could be one reason why sellers have continued to use auctions for items that are not particularly idiosyncratic or unusual. From this perspective, the rationale for auctions becomes less about price discovery and closer to the use of couponing and other promotional strategies seen in more traditional retail markets. We also discuss a few other factors that might have some explanatory power over the mix of sales formats, including platform fees and the general speed of market adaptations.

Our paper relates to a theoretical literature on the choice between auctions and posted prices, and to work on online auctions and online commerce. Prior theoretical work includes Wang (1993), Lu and McAfee (1996), and Kultti (1999). The literature on online auctions is surveyed by Bajari and Hortacsu (2004). Hammond (2010) focuses specifically on the choice between auctions and price posting in selling compact disks online. He finds that sellers with a higher opportunity cost tend to choose posted prices and that seller heterogeneity can account for the coexistence of the two mechanisms. Bauner (2011) studies the use of auctions and price posting using eBay sales of Major League Baseball tickets. He finds that not only opportunity costs, but heterogeneity across buyers, broadly divided into “fixed price lovers” and “neutrals,” help explain the coexistence of auction and fixed price sales. In a separate study of Major League Baseball tickets that focuses on dynamic price adjustment, Sweeting (2012) observes that auction prices are typically below posted price transaction levels. Finally, Ariely and Simonson (2003) and Malmendier and Lee (2011) analyze auction prices relative to posted prices in an earlier time period. They find that auction prices can be anomalously high, and argue that bidding behavior can contradict the predictions of standard economic models.

2 Setting, Data, and Preliminary Evidence

Our analysis focuses on eBay, which started in 1995 as an online auction site. Many auction sites were active at the time, but eBay became the most successful (Lucking-Reiley, 2000). Its marketplace transaction volume reached 9.3 billion dollars in 2001 (it has subsequently risen to over \$67 billion as of 2012). Initially, eBay was known for its used and idiosyncratic

items – the first item sold was a broken laser pointer – but it has evolved over time into a market that pools consumer sellers with small to large retailers. Nowadays, firms such as Buy.com and Toys“R”Us offer their inventory on eBay. In the summer of 2002, eBay began allowing sellers to list items at posted prices, using the “Buy it Now” (BIN) format.³ Transactions have shifted steadily toward this format, as was shown in Figure 1.

While this evolution has mirrored a broader trend in e-commerce, there are two features that make eBay a particularly interesting environment to study pricing mechanisms. The first is the coexistence of auctions and posted prices, which allows us to look at selection — who uses auctions and for what types of goods — and to compare the performance of auction and posted price mechanisms in a controlled fashion. The second is the platform’s gradual transition, which allows us to ask what underlying features of the economic environment might be responsible for the change. In this section, we start with a description of the stylized facts — first in the cross-section, and then in the time-series.

2.1 Who Sells What by Auction?

We begin by using recent data to document how auctions and posted prices are used for different types of items and by different types of sellers. Because of the time and effort they require from buyers, one expects auctions to be used in cases where it is difficult for the seller to forecast the appropriate price: either because the item is idiosyncratic, or consumer demand is uncertain, or because the seller has relatively little experience or has little interest in keeping the item or trying a second time to sell it. These types of predictions can be assessed in a relatively controlled setting because eBay has a robust market both for widely available goods and idiosyncratic or used goods, as well as a mix of professional and occasional sellers.

The data clearly show a greater prevalence of auctions for items that are more idiosyncratic. There are various ways to categorize items based on their “idiosyncrasy”. As one example, consider the distinction between used and new goods. Of the items listed in 2009

³A hybrid format has been introduced earlier, in which sellers list the item for auction but also set a price at which a buyer can preempt the auction by purchasing the item at that price before the first bid is submitted. We return to this hybrid format, which is still used, in Section 5.

that were labeled as “used” (i.e. the word “used” appears in the item’s 10-word listing title), 79% were offered by auction and only 21% by posted price.⁴ In contrast, of the items labeled “new” (i.e. the word “new” appears in the listing title) 44% were offered by auction and 56% by posted price. To the extent that it is harder to find comparable prices for used goods, or there is more uncertainty about buyer valuations, this is what one would expect.

Another way to identify more idiosyncratic items starts from the observation that many listings (about 73% in 2009) have “duplicates” — near-identical matches where the same seller has posted the same item with exactly the same title (but potentially with different pricing or sales format). Listings with duplicates generally are products that the seller is sourcing or has multiple units of inventory, whereas unique or truly idiosyncratic items (e.g. a lock of teen heartthrob Justin Bieber’s hair) would have no duplicate. Of the individual listings that appeared in 2009, about 75% are auctions, compared to 57% of the duplicates.⁵

Distinguishing between more and less idiosyncratic items helps explain why auctions are more or less prevalent across broad product categories. For instance, auctions are much more common in selling collectibles or clothing, whereas posted prices are much more common for selling electronics or computers. Figure 3 shows that product categories that have a higher fraction of “new” or duplicate listings (and hence less “idiosyncrasy”) also have a higher frequency of posted price listings. A similar pattern shows up for earlier years, although with a lower overall prevalence of posted prices.

We can identify similar patterns for sellers. Here, we rely on eBay’s internal classification of sellers based on past transaction volume on eBay. We compare experienced (“business”) sellers and inexperienced (“occasional”) sellers. There are also two groups of sellers with intermediate experience. Business sellers are much more likely to use posted prices. Of the items listed by business sellers, 65% have a posted price, compared to only 13% of the items listed by occasional sellers. It is also the case that sellers gravitate toward posted prices as they gain experience. If we look at sellers who entered the market between 2004 and

⁴In this calculation we limit attention to only “pure” auction listings and “pure” posted price listings, omitting other (less common) types, such as the hybrid auctions mentioned earlier. Unless specifically noted otherwise, we use this convention through the rest of the paper.

⁵Note that these percentages are for new listings. Because fixed price listings generally remain active for at least 30 days, whereas the most common auction length is 7 days, a consumer who randomly sampled an active listing would see a smaller fraction of auctions as compared to fixed price listings.

2008, and were active for at least two years, their second-year listings were more than three percentage points more likely to be posted price than their first-year listings, after controlling for the overall aggregate growth of posted price.⁶

The data also reveal a robust but less obvious relationship between the choice of sales mechanism and listing outcomes. The overall sale probability of auction listings is considerably higher than the sale probability of posted price listings. In 2009, 49 percent of auction listings sold compared to 42 percent of posted price listing. At the same time, the average auction sales price is significantly lower than the average posted price sale, and this pattern of higher prices and lower sale rates for posted price listings holds for every major category of items.⁷ These statistics do not control for within-category differences in the items and sellers across the different sales mechanisms, but we will show below that the same pattern appears even when we compare identical items sold by the same seller.

2.2 Why Have Auctions Declined?

Figure 1 already documented the dramatic decline in the use of auctions over the last decade. Based on the patterns above, a natural candidate to explain the decline is that the marketplace itself has evolved: from an environment where consumers sold each other unique items such as dolls or collectibles, towards a retail channel for business sellers. There is evidence to support this shift. The share of “idiosyncratic” (i.e. non-duplicate) listings declined from 32 percent to 25 percent between 2003 and 2009. And even as late as 2005, business sellers accounted for just 16 percent of listings; by 2009 they accounted for 27 percent.

Nevertheless, and perhaps surprisingly, compositional changes in the types of sellers and products do not seem to have much power in explaining the shift toward posted prices. One way to see this is to decompose the overall growth in posted price transaction volume into changes in the share of volume by different types of sellers, changes in the shares of different product categories, and changes that occurred within seller-product classifications.

⁶This result comes from a regression in which the dependent variable is the share of posted-price listings for a seller in his second year on eBay minus the share in his first year. The right-hand-side is the corresponding difference on eBay overall over the same two years. The number we report is the constant term. As discussed below, there is little evidence of additional cohort-specific patterns after the second year.

⁷Except for “Phones” where auction sale rates and prices are both higher.

Because 2005 is the first year for which we have the data that allow us to construct the seller classification, we do this for 2005 to 2009, a period which captures the sharpest shift in sales formats.

Define $Z = 20.9\%$ to be the overall share of posted price transaction volume (in dollars) in 2005 and $Z' = 53.9\%$ the corresponding share for 2009. Let s index the seller classifications (four groups ranging from “occasional” to “business”) and c index eBay’s 33 product categories. Next, let $\sigma_{c,s}$ denote the share of volume in item category c and seller category s in 2005, let $\sigma_c = \sum_s \sigma_{c,s}$ denote the share of volume in item category c , and let $\sigma_{s|c} = \sigma_{c,s}/\sigma_c$ denote the share of volume in seller category s within item category c , also in 2005.

Using σ' s to define the corresponding quantities for 2009:

$$\begin{aligned} Z' - Z &= \sum_{c,s} Z'_{c,s} \sigma'_{c,s} - \sum_{c,s} Z_{c,s} \sigma_{c,s} \\ &= \sum_{c,s} Z_{c,s} \sigma_{s|c} (\sigma'_c - \sigma_c) + \sum_{c,s} Z_{c,s} (\sigma'_{s|c} - \sigma_{s|c}) \sigma'_c + \sum_{c,s} (Z'_{c,s} - Z_{c,s}) \sigma'_{c,s}. \end{aligned} \tag{1}$$

The first term captures the shift toward posted prices due to a change in the composition of products (e.g., from collectibles to electronics). The second term captures the shift due to changes in the sellers within product categories. The final term is the average change within seller-product groupings. Almost all of the shift is attributable to the final component. Of the overall 33 percentage point increase in posted price transaction volume, the final term is responsible for 27.7 percentage points.⁸ An increase in business sellers accounts for 2.8 percentage points of the change toward posted prices, and changes in product composition for 2.5 percentage points.

There is another way to look for compositional changes that also proves informative. Suppose that early sellers were primarily consumers offering idiosyncratic items, while later arrivals tended to be professionals offering retail items. Then later cohorts of sellers might be more likely to use posted prices. To examine this, we group sellers based on the year during which they first listed an item on eBay (all sellers who joined prior to 2002 are grouped together). Figure 4 then plots for each cohort the fraction of their listings that were posted price, and how this fraction has evolved over time. The figure shows that while there is

⁸The prevalence of posted prices rose in almost every seller-product category cell. Of the $4 \cdot 33 = 132$ cells, the posted price share has increased in 131 of them.

a modest “first year” effect (noted earlier), at any given point in time, different cohorts of sellers behave in remarkably similar fashion, with all cohorts evolving over time toward posted pricing.

These time-series patterns suggest that to understand why auction listings have declined, we should look at changes in seller incentives. We start in the next section by developing a simple theoretical model of the choice between sales mechanisms. We then develop an empirical strategy that looks at the trade-off between auctions and posted prices for a given seller and product, and examine how this trade-off has shifted over time.

3 A Simple Model

In this section, we develop a model of the choice between sales mechanisms. The model is designed to emphasize a trade-off between price discovery and convenience for buyers. It is extremely parsimonious. We focus on a single seller offering a single item, and discuss extensions later in the paper. Consistent with the evidence above, we try to capture how changes in buyer demand, or changes in competition that give rise to changes in residual demand, might affect the choice of sales mechanism.

A seller has a single item to sell, and a cost c of making the sale. We assume there are two or more buyers, each with the same value v for the item, and a common fixed reservation utility u . We assume that v is drawn from a log-concave distribution F , where the distribution (but not the realization) is known to the seller.⁹ The seller can choose between posting a price or running a second-price auction. If the item is sold via auction, there is a transaction cost that reduces the buyer valuation from v to a lower amount $v - \lambda$, where λ is assumed to be identical across buyers and known to the seller.

The seller’s problem is to choose a sales mechanism, and either a posted price or a reserve price. Suppose the seller opts for a posted price p . The item will sell if $v - u \geq p$, or with probability $Q_F(p) = 1 - F(p + u)$. So the expected profit is $\pi_F(p) = (p - c)Q_F(p)$. If instead

⁹In Appendix A of the online appendix, we discuss a more general version of the model in which there are $n \geq 1$ bidders, and the value of bidder i is given by $v_i = v + w_i$, where v is common across bidders and w_1, \dots, w_n are independent across bidders. Under appropriate distributional assumptions, all of the comparative statics discussed in this section continue to hold in this more general model.

the seller runs an auction, and sets a reserve price r , the item will sell if $v - u - \lambda \geq r$, or with probability $Q_A(r) = 1 - F(p + u + \lambda)$. If the item does sell, the auction price will be $v - u - \lambda$, so the expected price conditional on sale is $p_A(r) = \mathbb{E}[v - u - \lambda \mid v - u - \lambda \geq r]$. The seller’s expected profit is therefore $\pi_A(r) = (p_A(r) - c) Q_A(r)$.

Figure 5 illustrates the difference between the two sale mechanisms. The black line shows the posted price demand curve, i.e. for any price p , the probability of sale is $Q_F(p)$. The gray line shows the implied “auction demand curve”. That is, each possible reserve price r is converted to its implied sale probability $Q_A(r)$ and the expected price conditional on sale $p_A(r)$. The demand curves are drawn assuming that consumer values are distributed uniformly on $[0, 1]$, with $u = 0$ and $\lambda = 0.2$. Finally, the dashed gray line shows the probability of sale associated with different reserve prices, that is, the auction sale probability $Q_A(r)$ for each reserve price r .

An immediate observation is that the posted price demand curve is steeper than the auction demand curve.¹⁰ This is intuitive: the auction uniformly reduces the willingness to pay of all buyers including those with the highest value, but it also creates competition that increases the final price above the reserve price. The second effect is largest for low reserve prices (or high sale probability). Both effects are visible in the picture if we consider the auction sale curve. The vertical distance between the posted price demand curve and the auction sale curve is the transaction cost λ , while the distance between the auction sale curve and auction demand curve represents the effect of competition – the expected amount by which the auction price will rise above the reserve price.

This structure has relatively clear implications for seller incentives. We will not take the view that every seller chooses the optimal sales strategy for every listing. Indeed our empirical approach in the next section will be premised on the idea that sellers engage in considerable experimentation with alternative strategies. Nonetheless, it is useful to consider the profit-maximizing incentives. A profit-maximizing seller will want to choose a point on the upper envelope of the demand curves: using a fixed price if she aims to sell at a high price, and an auction if she aims to sell with high probability. For a fixed price listing,

¹⁰More generally, assuming F is log-concave, then so long as $\lambda \leq \mathbb{E}[v] - \underline{v}$, where \underline{v} is the lowest possible value of v , the fixed price demand curve will start above the auction demand curve and cross it exactly once to end below it.

the optimal price p^* maximizes $\pi_F(p) = (p - c)[1 - F(p + u)]$, and satisfies the first order condition: $p = c + \frac{1 - F(p + u)}{f(p + u)}$. For an auction, the optimal reserve price is $r^* = c$. Overall, the seller does best to use a posted price if and only if $\pi_F(p^*) \geq \pi_A(r^*)$.

The model gives rise to several comparative statics predictions. (1) All else equal, a higher level of c pushes the seller toward a posted price. For the example shown in Figure 5, an auction is optimal for $c \lesssim 0.317$ while a posted price is preferable for $c \gtrsim 0.317$. (2) All else equal, a higher level of u (i.e. lower buyer value relative to the next best alternative) also pushes the seller toward a posted price. For instance, if we fix $c = 0.3$ in the example of Figure 5, a posted price is optimal if and only if $u \geq 0.017$. (3) Finally, and most obviously, a higher level of λ also pushes the seller toward a posted price by directly making auctions less attractive.¹¹

Another prediction of the model is that (under some additional restrictions on λ) conditional on using a posted price, a seller will optimally choose a higher p and lower $Q_F(p)$, than if she sells by auction, in which case using the optimal reserve price $r^* = c$ implies a probability of sale equal to $Q_A(c)$ and expected price $p_A(c)$. Below we will see that this general pattern holds in the data: if we look at sellers who have tried using posted prices and auctions to sell a given item, their sale rate is higher with auctions but their expected price is lower. These observations suggest some simple tests of the model, and a way to disentangle possible explanations for the shift from auctions to posted prices over time.

4 Evidence from “Seller Experiments”

Our empirical analysis makes use of what we call “seller experiments,” by which we mean the common practice of individual sellers listing the same item, either simultaneously or sequentially, to use different sale formats or with different pricing parameters. Elfenbein et al. (2012) were the first to make use of matched listings, in their study of charity auctions. In Einav et al. (2013), we documented the widespread use of this approach by eBay sellers,

¹¹Reduced uncertainty about buyer values will also make posted prices more attractive. Specifically, if the seller has access to a signal z that is informative about v , this will increase the expected profit from posted price, but have no effect on the expected auction profit. In a related but somewhat different model, Wang (1993) also has shown that a reduction in the “dispersion” of F will favor price posting.

and described some of the features of eBay’s platform that favored it, and why it might be attractive to sellers who were uncertain about their optimal sales strategy. We also discussed how, by aggregating many thousands of moderately sized “experiments,” it was possible to use very large-scale heterogenous data but still control tightly for the most plausible sources of confounding or endogeneity.¹²

4.1 Data construction and summary Statistics

Our goal is to compare the performance of auction and posted price listings, holding fixed the seller and item. To do this, we start with all listings posted in 2003, 2005, 2007, and 2009 by a random 50% sample of all eBay sellers.¹³ We then identify every set of matched listings in this data, where a match means the listings have a common seller, a common ten-word item description and were posted in the same year. For data quality reasons, we drop a small number of sets (0.8% or less in each of the years) where the average posted price is above \$10,000, or where there is a outlier sale that occurred at a price more than five times the average. From the remaining matched sets, we select those that include at least four listings of each format, at least ten listings overall, and at least one auction and one posted price sale. These “seller experiments” constitute our core data.

Relative to our model of a seller making an optimal choice between sale formats, focusing on sellers who use both formats might appear odd. However, an overwhelming feature of the data is that sellers very commonly use both auctions and posted prices. Moreover, even if sellers favor one sale format, they commonly continue to use both over time, without converging to just one. Below we will discuss why using both sale formats together could be profit-enhancing, but our general view is that sellers tend to engage in considerable and persistent experimentation, even if they gravitate over time toward more profitable strategies.

In comparing auction and posted price listings of a given item, we have to address two measurement issues. First, the typical posted price listing lasts for up to 30 days. In contrast, the typical auction lasts one week. Second, posted price listings can allow for the sale of

¹²In Einav et al. (2012), we used seller experiments data from a single year, 2009, to study the market response to different auction listing strategies: alternative reserve prices, different shipping fees, and pre-emptive “buy it now” offers, and to look at price dispersion.

¹³The 50% sample is used to keep the size of the data manageable.

multiple units, and multi-unit listings became common in the later years of the sample. Auctions are for a single unit. To account for these differences, we define a posted price listing to be successful if it sells at least one unit within the first week.¹⁴ Therefore, a “success” for each format means a positive answer to the question: “Did the seller sell a unit within a week of listing?”

Finally, our analysis will rely on aggregating comparisons involving different items. To facilitate this, we define for each item a *reference value* equal to the item’s average price across all fixed price listings in the relevant experiment. Later, we sometimes work with normalized prices. The normalized sale price is the sale price divided by the item’s reference value. Similarly, the normalized start price in an auction is the start price divided by the reference value. So an item that lists on average for \$100 and then is offered by auction with a start price of \$50 has a normalized start price of 0.5. If the item then sells for \$95, its normalized sale price is 0.95.

Table 1 provides summary statistics on the data for 2003 and 2009. Relative to all of eBay, the data are of course a selected sample. They exclude the many items that are truly “one of a kind” and appear as one-off listings. As we noted above, these are overwhelmingly auction listings. We also have dropped hybrid auctions and other sales formats such as “best offer” as we did in Section 2. Nevertheless, the resulting data set is still quite large, and includes items from all product categories.

The statistics in Table 1 already indicate that one basic prediction of the model holds in the data: auction listings are more likely to sell than their matched posted price counterparts, but they result in lower prices on average.¹⁵ Figure 6 uses the 2009 sample to show that in fact this prediction holds in essentially every single product category. The figure is constructed by finding the differences in sale probability and price for every experiment in the sample, and then averaging these differences across the items in each category. With this construction,

¹⁴Due to the way the data is archived, we currently do not observe the exact timing of sales of multi-unit listings that are associated with more than one sale. We therefore use the distribution of the time until sale for single-unit posted-price listings to convert the overall sale rate of multi-unit listings to sale rate within a week. The underlying assumption is that the timing of sales (conditional on sale) is the same for multi-unit and single-unit posted-price listings.

¹⁵The one somewhat surprising exception to this is Jewelry, which in fact has a very high overall fraction of auction listings (see Figure 3), but where both the probability of sale as well the average price conditional on sale is higher for posted prices.

the differences in the sale rate and sale price are not due to some items being sold primarily by auction and others by posted price, but rather reflect differences between format outcomes holding the mix of items (and sellers) constant.

4.2 Changes in the “auction discount”

How do auction prices compare to posted prices for the same item offered by the same seller? To consider this, let i index sets of matched listings, and j index sales (i.e. successful listings) within an experiment. Let A_{ij} be an indicator variable equal to 1 if the sale resulted from an auction, and let p_{ij} denote the transaction price (posted price or final auction price).

We compare auction prices to posted sale prices using the fixed effects regression:

$$\log p_{ij} = a_i + bA_{ij} + e_{ij}. \quad (2)$$

In this specification, the coefficient b can be interpreted as the price premium (conditional on sale) for an auction listing relative to a fixed price listing of the same item by the same seller. Equivalently, the expected (log) auction price conditional on sale for item i is $a_i + b$, where a_i is the average (log) price of the posted price transactions of item i .

We estimate the regression separately for each year, and plot the estimated price differentials (the b coefficients) in Figure 7(a).¹⁶ As the figure makes clear, the price differential has widened dramatically. Sellers offering a given item by auction and posted price in 2003 received on average about 4.7% less from their auction sales. By 2009, the auction discount had widened to 16.5%. Interestingly, the average transaction price for posted price sales remained very similar over this period (despite turnover in the set of items).¹⁷ In our sample, the average log posted price was 3.1 in 2003 and 3.0 in 2009.

The price differentials in Figure 7(a) are estimated from two scenarios in the data that are potentially useful to distinguish. The first are situations where a seller conducts an auction while contemporaneously offering the same item by posted price. The second are situations

¹⁶The regression results underlying Figure 7 are available in Appendix Table A1(a).

¹⁷That is, the average (across items) of the a_i 's. To construct this average we use all items in the sample (which were all sold at least once in each format, given our sample restrictions), and weight each item by the number of sales in each format. We use the same average in Figure 7(a).

where the seller uses both formats to sell a given item but not at the same time. In principle, the former situation might offer a more direct look at buyer preferences (although because eBay’s search algorithm tends to separate matched listings, buyers often may not make a direct comparison). The latter arguably fits more cleanly with our theoretical model, where the outside option (u) is exogenous. In the Appendix, we replicate the analysis for these two different situations. The estimated auction discount (in either 2003 or 2009) is higher for auctions with a contemporaneous posted price. However, the main result that we emphasize — the growing auction discount over time — does not depend much on exactly how we set up the posted price comparison.¹⁸

It is also interesting to compare the success rates of matched auction and posted price listings. Again let i index a group of matched listings, and let j index listings within each group. We define q_{ij} to be an indicator equal to one if the listing results in a sale,¹⁹ and estimate a fixed effects linear probability regression of the form:

$$q_{ij} = \alpha_i + \beta A_{ij} + \varepsilon_{ij}. \tag{3}$$

In this specification, α_i is the average sale probability for posted price listings of item i , and β is the extra sale probability for auction listings. Again, we estimate the specification separately for each year. Figure 7(b) shows the results. The black line plots the average sale probability for a posted price listing, i.e. the average of the estimated α_i ’s for each year. The grey line plots the corresponding sale probability for auction listings, computed as the fixed price sale probability for that year plus the estimated β .

Two features stand out. As noted in Section 2, the average sale probability for both formats has declined. The decline has been substantial. The average posted price sale rate in our sample fell from 40.3% to 25% between 2003 and 2009. However, the difference in

¹⁸Appendix Table A1(b) and Appendix Figure A1 provide more details. For auctions with concurrent posted price listings, the auction discount has increased from 11.3% in 2003 to 18.2% in 2009, while for auctions with no concurrent posted price listing the discount has increased from 2.2% to 14.4%.

¹⁹As mentioned earlier, the typical auction duration on eBay is seven days, while the typical posted price listing lasts for 30 days (unless it sells earlier). For comparability, throughout the analysis in this section we define a successful posted price listing to be one that sells within the first seven days, although the results are not particularly sensitive to this choice.

success rates across formats (the estimated β), has not changed much.²⁰ It starts at about 8% in 2003, declines a bit, and then increases to 11% in 2009. A similar “participation” pattern arises if we look at the number of bidders in successful auctions. In 2003, the average number of unique bidders was 3.98 and 33% successful auctions had a single bidder. By 2009, the average number of bidders had fallen by 20% to 3.19 and the share of single-bid auctions had increased to 45%. Appendix Figure A2 provides more details.

Taken together, these results suggest two important changes over time: a secular fall in the success rate of item listings, and a fall in the relative demand for auctions that has reduced auction prices relative to comparable posted prices. These estimates do have one limitation. They do not account explicitly for possible changes in how sellers are setting auction reserve prices. In principle, lower auction prices could be explained by lower reserve prices. However, as we show in the Appendix, sellers in fact moved towards slightly *higher* reserve prices (see Appendix Figure A3). An implication is that the increase in the auction discount resulted from a fall in auction sale prices conditional on reserve prices (see Appendix Figure A4). We develop this point more fully in the next section.

5 Changes in Auction Demand

The growing auction discount provides a fairly strong indication that seller incentives have shifted. But as we noted, it does not cleanly separate shifts in demand — due to changes in buyer preferences, or in competition — from shifts in pricing behavior. In this section, we use variation in sellers’ pricing behavior to estimate the trade-off sellers face between probability of sale and expected price for both posted price and auction listings. These estimates can be interpreted as residual, or “listing-level” demand curves. We use the estimates at various points in time to map back to the parameters in the theoretical model that determine sellers’ incentives to use auction or posted price listings.

²⁰The use of category level or seller level fixed effects instead of experiment level fixed effects does not give rise to meaningful changes in the average success rate of posted price and auction listings.

5.1 Estimates of listing-level demand

To estimate listing-level demand, we exploit the within-experiment variation in auction start prices and posted prices. The variation in auction start prices is very large. The average within-experiment standard deviation in auction start price is 25% of the item’s reference value, as reported in Table 1. The variation in posted prices is not as dramatic. Nevertheless, the average within-experiment standard deviation in posted prices is still over 10% of the item’s reference value.

We consider the auction case first. We start with all auction listings in our sample, and estimate separately for each year the probability of sale as a function of the auction start price:

$$q_{ij} = \alpha_i + g(s_{ij}^n) + \varepsilon_{ij}. \quad (4)$$

In this specification, i indexes sets of matched listings, j indexes listings within an experiment, and q_{ij} is an indicator variable equal to one if the auction results in a sale. The key pricing variable is the normalized start price, s_{ij}^n . It is equal to the start price s_{ij} divided by the reference value for item i . The normalization allows for a comparison of items of different value. We specify the function $g(\cdot)$ flexibly by allowing separate dummy variables for different levels of the normalized start price.²¹ In this regression and those that follow, we weight each experiment by its total number of listings.

We estimate a parallel regression for the final auction price. Here we restrict attention to successful auctions and estimate for each year:

$$p_{ij}^n = a_i + h(s_{ij}^n) + e_{ij}. \quad (5)$$

The normalized price p_{ij}^n is the final auction price divided by the item’s reference value. Again, $h(\cdot)$ is estimated flexibly by allowing separate dummy variables for each level of the normalized start price.

The last step is to combine the regression estimates and trace out an auction demand curve for each year. For each value of the normalized start price s^n , we consider the locus

²¹Specifically, we categorize the (normalized) start price to the following bins: [0-0.3], (0.3-0.4], (0.4-0.5], (0.5-0.6], (0.6-0.7], (0.7-0.8], (0.8-0.9], (0.9-1], (1-1.1], (1.1-1.2], and >1.2.

$\bar{q}(s^n) = \bar{\alpha} + \hat{g}(s^n)$, $\bar{p}^n(s^n) = \bar{a} + \hat{h}(s^n)$, where $\bar{\alpha}$ and \bar{a} are the average fixed effects. (As before, the averages are weighted by the total number of listings in the corresponding experiments.) One point to emphasize is that we are constructing for each possible start price the probability of sale and the *expected* sale price (if there is a sale). Chen et al. (2012) have suggested that one reason sellers may prefer posted prices is to avoid the price variability that occurs with an auction, but we will not focus on this point here.

We also estimate posted price demand curves separately for each year. This exercise is conceptually more straightforward. We estimate the fixed effects linear probability model:

$$q_{ij} = k_i + m(p_{ij}^n) + \eta_{ij}. \quad (6)$$

Here q_{ij} is an indicator equal to one if the listing results in a sale, and p_{ij}^n is the normalized price of the listing (equal to the listing price divided by the item's reference value). To estimate $m(\cdot)$ we specify separate dummy variables for different levels of the normalized price.²² Note that with this flexible specification, the within-experiment variation in posted prices will limit the range over which we can obtain a demand estimate.

Figure 8 presents the estimated demand curves using the 2003 and 2009 samples.²³ One prediction of the theory is apparent. In both years, the posted price demand curve is steeper than the auction demand curve. Their relative positions, however, look quite different in 2003 and 2009. In 2003, the estimated auction demand lies mostly above the posted price demand. A seller facing these residual demand curves would be better off using an auction, unless her marginal cost was extremely high. In 2009, the auction demand lies below the posted price demand for a wide range of prices. For a seller facing these residual demand curves, the optimal sale format would depend on marginal cost.

The estimates in Figure 8 involve averaging over a range of heterogeneous items. It is informative to repeat the exercise focusing on more narrowly defined groups. To do this, we used the characteristics described earlier to rank the eBay product categories from

²²The pricing variation for posted-price listings is not as rich as it is for the auction listing. The specific (normalized) price bins we use are [0-0.7], (0.7-0.8], (0.8-0.85], (0.85-0.9], (0.9-0.95], (0.95-1.05], (1.05-1.1], and >1.1.

²³The underlying regression results are available in Appendix Tables A2 and A3 and Appendix Figure A5.

“most idiosyncratic” to “least idiosyncratic”.²⁴ We then selected the five most idiosyncratic and five least idiosyncratic categories from the largest fifteen categories. The first group of “idiosyncratic” categories includes Collectibles, Fan Shop, Toys, Jewelry, and Clothing. The second group of “commodity” categories includes Electronics, Computers, DVDs, Health, and Phones.

Figure 9 presents separate demand estimates for each year and group of categories.²⁵ The estimates for the idiosyncratic categories show that the average residual demand for auctions dominated posted prices in 2003. By 2009, however, the auction demand curve had shifted considerably, making posted prices attractive across a fairly wide price range. In contrast, the average demand estimates for the commodity categories suggest that posted prices were attractive even in 2003. By 2009, posted prices had become even more attractive, but only because both demand curves shifted inward by approximately the same amount. We return to the interpretation of these shifts below.

5.2 Identifying seller incentives

The estimates in Figure 8 show a clear shift in the relative demand for auctions over time. They also can be viewed as the empirical analogues of the theoretical demand curves in Figure 5. In this section, we tighten the connection. We use the demand estimates to calibrate the key parameters of the model. While the model is obviously too parsimonious to capture the heterogeneity of the many thousands of sellers and items in the data, the exercise is nonetheless useful to place some structure on how and why seller incentives have shifted over time. We do this both on average across all categories and for the idiosyncratic and commodity categories defined above.

The theoretical model derived item-level demand curves from three parameters: the

²⁴To construct the ranking, we used a large random sample of 2009 listings. For each category, we constructed the following variables: share of items that are classified in the eBay catalog, share of sales from multi-unit listings, share of listings with the word “new” in the title, and share of duplicate listings. We then used principal components to assign each category a one-dimensional score, with a high score corresponding to “less idiosyncratic” and a low score corresponding to “more idiosyncratic”. Appendix Table A4 provides more details. One point of caution is that although this classification seems quite natural, the items from the idiosyncratic categories that make it into our experiments sample are still retail items, in the sense that the seller is trying to sell multiple units. Hence we do not include the truly one-of-a-kind items that might come to mind in thinking about an idiosyncratic category such as Collectibles.

²⁵The underlying regression results are available in Appendix Tables A5(a), A5(b), and A5(c).

distribution of consumer valuations $F(\cdot)$, the utility from the next best alternative u , and the auction disutility λ . Each might have changed over time. Roughly, one might think of changes in the reservation utility u as capturing intensified competition, perhaps due to consumers being able to search more easily for alternative products. In contrast, changes in the auction disutility λ are a potential proxy for increased online entertainment options that may have made auction bidding less fun or exciting in comparison.

We begin by estimating λ . We already observed that in the theoretical model λ could be identified as the vertical distance between the fixed price demand curve and the auction sales curve. To see why, recall that given a posted price p , the probability of sale is $1 - F(p + u)$. Similarly, given an auction with start price r , the probability of sale is $1 - F(r + u + \lambda)$. So a given probability of sale q corresponds to a posted price $F^{-1}(1 - q) - u$ and an auction start price $F^{-1}(1 - q) - u - \lambda$. The difference is λ .

Figure 10 plots the estimated posted price demand curves for 2003 and 2009 (identical to Figure 8), along with the estimated auction sale curves from the auction sale regression (4). If the estimates were perfectly consistent with the theoretical model, the posted price demand curve for year t would be an exact upward translation of the auction sale curve. In practice, the shapes are similar but not identical, so to estimate λ_{2003} (corresponding to 2003) and λ_{2009} (corresponding to 2009), we compute the average vertical distance between the curves across the quantity range where we have nonparametric estimates of both curves.²⁶ We find that $\lambda_{2003} = 0.076$, and $\lambda_{2009} = 0.163$. That is, we estimate that an auction in 2003 involved a “hassle cost” equal to roughly eight percent of an item’s typical price (or reference value), and that this cost doubled between 2003 and 2009.

Next, we estimate changes in u by looking at how demand has shifted over time. Here, a normalization is required. Because consumer decisions are based on the incremental value $v - u$, it is not possible to separately identify u from the mean of the v distribution. To proceed, we normalize $u_{2003} = 0$. We further assume that $F(\cdot)$ remained constant over time.²⁷

²⁶The results we report use an unweighted average of the vertical distance between the two curves. We repeated the analysis using two sets of weights (one that weights by the reserve prices of auction listings in the data, and one that uses the empirical distribution of posted prices in the data), and the calibrated values for the parameters remain essentially the same.

²⁷In principle, we only need a known percentile of $F(\cdot)$ to have remained constant, but making the stronger assumption keeps things very simple. Note that if the distribution of consumer valuations remained constant, the posted price demand curves in Figure 10 would be exact translations. In practice they are similar in

It follows that to achieve sale probability q in year t required a posted price $F^{-1}(1 - q) - u_t$. So the vertical distance between the 2003 and 2009 posted price demand curves is $\Delta u = u_{2009} - u_{2003}$. We therefore estimate Δu by computing the average vertical distance between the estimated demand curves. This yields an estimate of $\Delta u = 0.144$. That is, between 2003 and 2009, the consumer reservation utility increased by around 14% of the reference value for an average item.²⁸

The final step is to evaluate how these changes in λ and u might have affected seller incentives. Here we return to the model and use some price theory. Let $p_F(q; \lambda, u) = F^{-1}(1 - q) - u$ denote the posted price a seller could charge and sell with probability q . With this notation, the seller's maximum profit from price posting is:

$$\Pi_F(\lambda, u) = \max_q q [p_F(q; \lambda, u) - c]. \quad (7)$$

From the envelope theorem,

$$\frac{\partial \Pi_F}{\partial u} = q_F^* \cdot \frac{\partial p_F}{\partial u} = -q_F^*, \quad (8)$$

where q_F^* is the optimal sale probability. Moreover, $\partial \Pi_F / \partial \lambda = 0$. Intuitively, a small change in u will not affect the optimal probability of sale, but will force the seller to lower the price by du to offset the reduction in residual buyer value.

Next, consider an auction seller. To sell with probability q requires a reserve price $r = F^{-1}(1 - q) - u - \lambda$. This leads to an expected price $p_A(q; \lambda, u) = \mathbb{E}[v \mid v \geq F^{-1}(1 - q) - u - \lambda]$. So the seller's maximum profit from an auction is:

$$\Pi_A(\lambda, u) = \max_q q [p_A(q; \lambda, u) - c]. \quad (9)$$

Again using the envelope theorem, $\partial \Pi_A / \partial \lambda = \partial \Pi_A / \partial u = -q_A^*$, where q_A^* is the optimal sale probability for an auction seller.

We are interested in the *difference* in posted price and auction profits. We have shown

shape, but not identical.

²⁸Note that the analysis of each year is based on a different set of seller experiments, so in order to compare demand and outcomes across different years, we implicitly assume that overall pricing patterns on the internet remained unchanged (e.g., we assume that markups online did not decline), so that the average posted-price transactions serve as a reasonable online price index.

that

$$\frac{\partial (\Pi_F - \Pi_A)}{\partial \lambda} = q_A^* \quad \text{and} \quad \frac{\partial (\Pi_F - \Pi_A)}{\partial u} = q_A^* - q_F^*. \quad (10)$$

The derivation indicates that a change in λ , which shifts only the auction demand curve, will have a bigger effect than an equivalent change in u , which shifts both demand curves. On the other hand, we have estimated a larger absolute change in u .

To sort out these competing effects, we approximate the effect of $\Delta\lambda$ on $\Pi_F - \Pi_A$, by $q_A^* \cdot \Delta\lambda$, and the effect of Δu by $(q_A^* - q_F^*) \cdot \Delta u$. This requires information on q_A^* and q_F^* . In 2003, the average sale probability in our data was 0.40 for posted price sales and 0.48 for auctions, so for this exercise we assume $q_F^* = 0.40$ and $q_A^* = 0.48$.

Table 2 shows the resulting calculation for the effects of u and λ on seller incentives. The calculation suggests that the increase in λ had more than three times the effect on the posted price profit differential as did the increase in u . If we set q_F^* and q_A^* using the 2009 sale probabilities, we obtain a somewhat smaller multiple, just under two. But either way the increase in λ seems considerably more important.

It is also informative to focus on particular product categories, as we did in the prior section. When we restrict attention to only the most idiosyncratic product categories, we estimate that $\lambda_{2003} = 0.06$ and $\lambda_{2009} = 0.186$. So $\Delta\lambda = 0.126$, while $\Delta u = 0.109$. As a result, by far the most important factor in these categories in explaining the move toward posted prices was a fall in the relative demand for auctions. The increase in λ accounts for well over 80% of the total change in relative profits.

We obtain somewhat different quantitative conclusions for the commodity categories. There we estimate $\lambda_{2003} = 0.186$ and $\lambda_{2009} = 0.235$. So $\Delta\lambda = 0.048$, while $\Delta u = 0.153$. As noted above, there appears to be a significant auction discount in these categories as early as 2003. As shown in Table 2, the implication is that the increase in λ and u had roughly similar effects on the profit differential between posted prices and auctions. To understand why the effects are equivalent despite the increase in u being much larger, recall that λ affects only the auction demand curve, whereas the increase in u reduces the profitability of both sale formats.

6 Changing Demand and Buyer Heterogeneity

Our analysis has highlighted the role of changing (residual) buyer demand in explaining the shift toward posted price listings. In this section, we ask whether the change was due to shifts in the composition of buyers, or rather in the behavior of buyers over time. We first show, parallel to our earlier analysis of sellers, that there is considerable heterogeneity in the cross-section of buyers. However, and also parallel to our seller findings, the time-series shifts in demand appear to have taken place within cohorts of buyers rather than through changes in the composition of buyers over time. Finally, we point out that the cross-sectional heterogeneity in consumers can help to explain why sellers might want to offer both types of sales mechanisms as a way to target different types of buyers.

As in our earlier classification of sellers, we divide buyers based on their purchases in a given year. We label buyers who purchase more than 100 items a year as “professional” and those who buy less than 5 as “occasional”, and include an intermediate category as well. Professional buyers are more likely to buy at auction. In 2003, the auction expenditure share of these buyers was 95%, compared to 80% for occasional buyers. The corresponding statistics for 2009 were 69% and 38%. Using the successful sales in the matched listings sample, we find that the buyer is 5.4% more likely to be a professional buyer if the item was listed for sale by auction.²⁹

An additional, and quite striking, feature of the data is that the more active buyers seem to get better deals. First, we have already documented an auction discount. So by doing more purchasing at auction relative to at posted prices, professional buyers are more likely to experience discounts. A more telling comparison, however, is to compare the prices paid at auction by professional and occasional buyers for equivalent items, using the experiments data. Using the 2009 data, we find that professional buyers on average obtain a 20% auction discount, compared to only 12% for the occasional buyers.³⁰

Despite these cross-sectional differences, however, the shift away from auctions over time

²⁹This is based on a linear probability regression (not shown), in which we use all the posted-price and auction sales and regress an indicator for professional buyer on an auction dummy variable and experiment fixed effects.

³⁰This estimate is based on regressing (log) price on experiment fixed effects, an auction dummy and a categorical variable for buyer’s types interacted with the auction dummy (not shown).

cannot be explained by an increase in occasional relative to professional buyers. The fractions of occasional and professional buyers on eBay have been relatively stable over time. The professional buyers account for just over 1% of the buyers, and the occasional buyers for around 60%. In fact, the purchasing share of professional buyers increased over the period we study. Professional buyers accounted for 14% of transaction volume in 2003 and 20% in 2009. Instead, the shift toward auctions has been within buyer groups, when buyers are stratified by purchasing intensity.

The decline of auctions also cannot be explained by differences between early cohorts of online shoppers and later cohorts. Figure 11 shows the purchasing propensities of different buyer cohorts, analogous to our earlier depiction of sellers. As can be seen in the Figure, buyers in their first year are more likely than average buyers on the site to buy using a posted price, but after the first year, different cohorts look very similar, and all cohorts have trended together toward posted price purchasing. Our interpretation, therefore, is that the documented shift in item-level demand curves is likely to come from changes in online attitudes and competition, rather than from changes in the composition of buyers.

7 Other Explanations and Extensions

We have provided a variety of evidence about how and why internet auctions declined over the last decade. In this Section, we briefly discuss three factors that did not play a major role in our analysis, but that potentially might have some explanatory power. The first is changes in the eBay platform that might have favored posted price listings. The second is the possibility that the basic empirical pattern can be explained by gradual seller learning. The third, which does not directly explain the fall in auction listings, but might help to explain their continued use, is the idea that sellers might use auctions and posted prices in combination as a form of price discrimination.

Changes in the eBay platform. Our analysis emphasized changes in the marketplace incentives facing sellers. There also have been some changes in eBay's platform that may have favored posted prices, especially beginning in 2008. In February 2008, eBay changed its search ranking algorithm. Rather than putting the soonest-to-end listing at the top — a

natural strategy for auctions with a fixed ending time, but less natural for posted price listings — it implemented a relevance ranking. Then in September 2008, eBay allowed 30 day posted price listings to be “rolled over” with automatic payment of the monthly listing fee. The mechanical effects of that change help explain the sharp concurrent rise in active posted price listings that can be seen in Figure 1. However, given that posted prices already had been on the rise for several years, it probably makes more sense to think of these platform changes as a response to changes in the market, rather than the impetus for evolution of sales format.

In addition to its search algorithm, eBay also has changed its fees on multiple occasions over the last decade. Many of these changes have not had clear implications for the choice of sales format, and their net effect is not entirely obvious. To take just one example, eBay in September 2008 lowered its posted price listing fee and raised the commission rate for successful sales. Posted price listings increased, despite the fact that in expectation total fees were slightly higher.³¹ More generally, it seems difficult to tell a story in which the steady decline of auctions in Figure 1 is driven by a series of occasional, discrete changes in the fee structure, only some of which have favored posted prices.

Gradual seller learning. Until the middle of 2002, eBay sellers did not have the option to use posted price listings, although they could use a hybrid format that allowed buyers to pre-empt an auction by purchasing at the “buy it now” price. One hypothesis we have heard is that once posted prices were introduced, they were from the start the better mechanism, but it took a long time for sellers to realize this and for posted prices to diffuse in the market.

Several pieces of evidence go against this hypothesis. First, posted prices are hardly an unfamiliar technology, so even if there is some stickiness in buyer and seller behavior, a decade of learning seems like a very long time. Second, our estimates of auction prices and demand indicate that in 2003, auctions were in fact a very attractive sale format. In particular, auctions offered a profitable way for sellers to increase sales probability in return for accepting a modest price discount.

³¹The effects here are confounded because sellers also gained the opportunity to roll over their listings. However, the fee changes were differential across categories, and the increase in posted price listings is not closely related to the differential changes in the fees.

Price discrimination strategies. In the previous section, we pointed out that more active buyers on eBay are more likely to buy at auction compared to occasional buyers. In principle, this heterogeneity might help to explain why sellers might want to offer both sale formats as a way to cater to different types of buyers. We illustrate this in a very simple extension of the model.

Imagine that there are two types of buyers: “deal shoppers” (denoted as type D) and “convenience buyers” (denoted as type C). They have the same intrinsic value v for a given item, drawn from $F(\cdot)$, but different (λ, u) . Assume that the convenience shoppers have a high λ_C so they are not interested in auctions, and a relatively low u_C that increases their willingness to pay. On the other hand, deal shoppers have a low λ_D and, because they are searching for deals, a relatively high u_D . Then, provided that $u_D - u_C$ is sufficiently large, and there are enough convenience shoppers, a seller who lists by posted price will sell to the convenience shoppers but to few (if any) deal shoppers. If she adds additional auction listings, she may be able to sell to the deal shoppers, without cannibalizing her existing market.³²

One point to note is that allowing for buyer heterogeneity as in this example might suggest a modification of our calibration exercise, because at least in this example auction and posted price demand would come from different types of buyers, so the lower auction sales curve could reflect the lower willingness-to-pay (higher u) of the deal shoppers as well as their λ . Trying to construct richer econometric models of e-commerce that incorporate different forms of heterogeneity across buyers, sellers, and items, and might help rationalize different types of price discrimination would be a worthwhile goal for future research.

³²In a more general version of this model, we might consider a more continuous distribution of types. In that case, one would have to take a stand on the extent to which buyers arbitrage between sale formats when both are offered at the same time. We do not focus on such a model because based on our estimates in Section 4.1 it is not clear that arbitrage between sale formats for a given seller/item is playing a very major role on the eBay marketplace (or else we would have expected to see very different pricing results for contemporaneous vs. non-contemporaneous auction and posted-price sales). One interesting feature of these models is that they also can rationalize the use of the hybrid auctions used by some sellers on eBay.

8 Conclusion

In this paper we document the dramatic shift of internet retail from auctions to posted prices, and explore the potential causes to it. Using rich data from eBay, we provide evidence that suggests that changes in the composition of buyers, sellers, or items is unlikely to be driving the decline in online auctions. Rather, we show that the shift can be primarily attributed to a decline in the relative demand for auctions, which operated within cohorts of buyers and item categories. We further explore whether the decline in auction demand is driven by increased retail competition or by a change in preferences toward convenience shopping. Using a calibration exercise of a stylized model, we find that the latter effect accounts for much of the overall effect, especially for item categories that include more idiosyncratic items.

As emphasized throughout the paper, our modeling approach and subsequent calibration exercise were purposely very parsimonious and stylized. A unique feature of internet commerce in general, and eBay in particular, is the extent of heterogeneity in the items that are being sold. One can easily imagine that the aggregate decline in online auctions, its precise timing, and the exact mechanism that led to it has operated slightly differently across sub-markets. Our goal in this paper was not to provide a story that fits every underlying item category, but rather to shed light on the primary factors that were in play. Indeed, trying to construct richer econometric models of e-commerce that would either focus on several item categories or try to incorporate the requisite heterogeneity across buyers, sellers, and items would be a worthwhile exercise, which would complement the evidence provided here.

Our analysis focus on eBay during the period of 2003 to 2009. While studying eBay – which in 2009 had ninety million active users and \$57 billion in gross merchandise volume – may be of interest in its own right, a natural question that may arise is to what extent our findings extend into later periods and more generally to other environments on the internet. Regarding the former, we suspect that the increased use of smart phones by eBay users may help exacerbate the decline in auctions; on smart phones, eBay browsing sessions are 25% shorter and cover 50% fewer page views. One might expect smart phone users to favor convenience even more. Regarding the latter, the recent decision of Prosper.com, in December 2010, one of the two leading social lending web sites, to abandon its original auction


model in favor of pre-set interest rates suggest that similar mechanisms may be at play elsewhere. Indeed, Prosper’s CEO explained the decision by arguing that pre-set rates were more efficient and that auctions were “too complicated” (Gonsalves, 2010). A more careful answer to this question, however, would require data we don’t currently have, for example on users’ internet browsing and usage and the extent to which other online entertainment options, such as Facebook or YouTube, have crowded out the early excitement that was initially associated with bidding in online auctions. We think that this is a promising avenue for further research.

References

- Ariely, Dan, and Itamar Simonson (2003). “Buying, Bidding, Playing, or Competing? Value Assessment and Decision Dynamics in Online Auctions.” *Journal of Consumer Psychology* 13, 113-123.
- Bajari, Patrick, and Ali Hortaçsu (2004). “Economic Insights from Internet Auctions.” *Journal of Economic Literature* 42(2), 457-486.
- Bauner, Christoph (2011). “Mechanism Choice and the Buy-It-Now Auction: A Structural Model of Competing Buyers and Sellers.” Working Paper, University of Massachusetts.
- Chen, Jong-Rong, Kong-Pin Chen, Chien-Fu Chou, and Ching-I Huang (2012). “A Dynamic Model of Auctions with Buy-it-now: Theory and Evidence.” *mimeo*.
- The Economist (2000), “In The Great Web Bazaar.” *The Economist*. February 24, 2000.
- Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan (2013). “Learning from Seller Experiments in Online Markets.” NBER Working Paper No. 17385.
- Elfenbein, Daniel W., Ray Fisman, and Brian Mcmanus (2012). “Charity as a Substitute for Reputation: Evidence from an Online Marketplace.” *Review of Economic Studies*.
- Elfenbein, Daniel W., Ray Fisman, and Brian Mcmanus (2013). “Market Structure, Reputation and the Value of Quality Certification.” Working Paper, University of North Carolina.
- Gonsalves, Atone (2010). “Social Lender Prosper.com Drops Auction Model.” *InformationWeek*, December 20, 2010.
- Hall, Robert E. (2002). *Digital Dealing: How eMarkets are Transforming the Economy*. W. W. Norton.
- Hammond, Robert G. (2010). “Comparing Revenue from Auctions and Posted Prices.” *International Journal of Industrial Organization* 28(1), 1-9.

- Kultti, Klaus (1999). "Equivalence of Auctions and Posted Prices." *Games and Economic Behavior* 27, 106-113.
- Malmendier, Ulrike, and Young Han Lee (2011). "The Bidder's Curse." *American Economic Review*.101(2), 749-787.
- Lu, Xiaohua, and R. Preston McAfee (1996). "The Evolutionary Stability of Auctions over Bargaining." *Games and Economic Behavior* 15, 228-254.
- Lucking-Reiley, David (2000). "Auctions on the Internet: What's Being Auctioned, and How?" *Journal of Industrial Economics* 48(3), 227-252.
- Surowiecki, James (2011). "Going, Going, Gone: Who Killed The Internet Auction?" *Wired Magazine*, May 17, 2011.
- Sweeting, Andrew (2012). "Dynamic Pricing Behavior in Perishable Goods Markets: Evidence from Secondary Markets for Major League Baseball Tickets." *Journal of Political Economy*, forthcoming.
- Wang, Ruqu (1993). "Auctions versus Posted-Price Selling." *American Economic Review* 83(4), 838-851.

Figure 1: Auction Share on eBay over Time



For each month, the figure shows the average daily share of active eBay listings (black) and transaction revenues (gray) from pure auction listings out of all pure auction and posted price listings. Less common formats, such as hybrid auctions, are not included. The sharp drop in Fall 2008 coincides with a decision in September 2008 to allow "good till canceled" posted price listings (see Section 7).

Figure 2: Google Search Volume for Online Auctions and Online Prices




Figure presents results from "Google Trends" for search terms "online auctions" and "online prices." The y-axis is a Google generated index for the weekly volume of Google searches for each of the two search terms, which should make the weekly volume figures comparable over time and across the two search terms.

Figure 3: Posted Prices are more Prevalent for “Less Idiosyncratic” Items

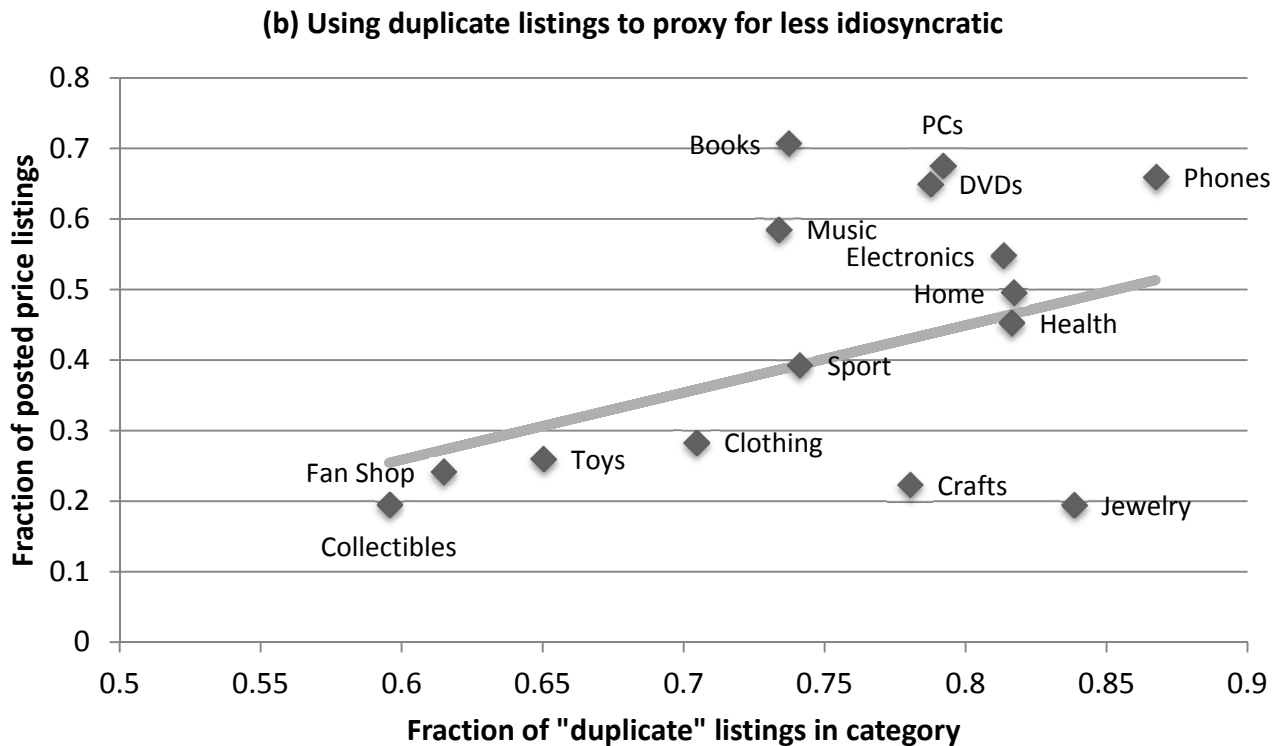
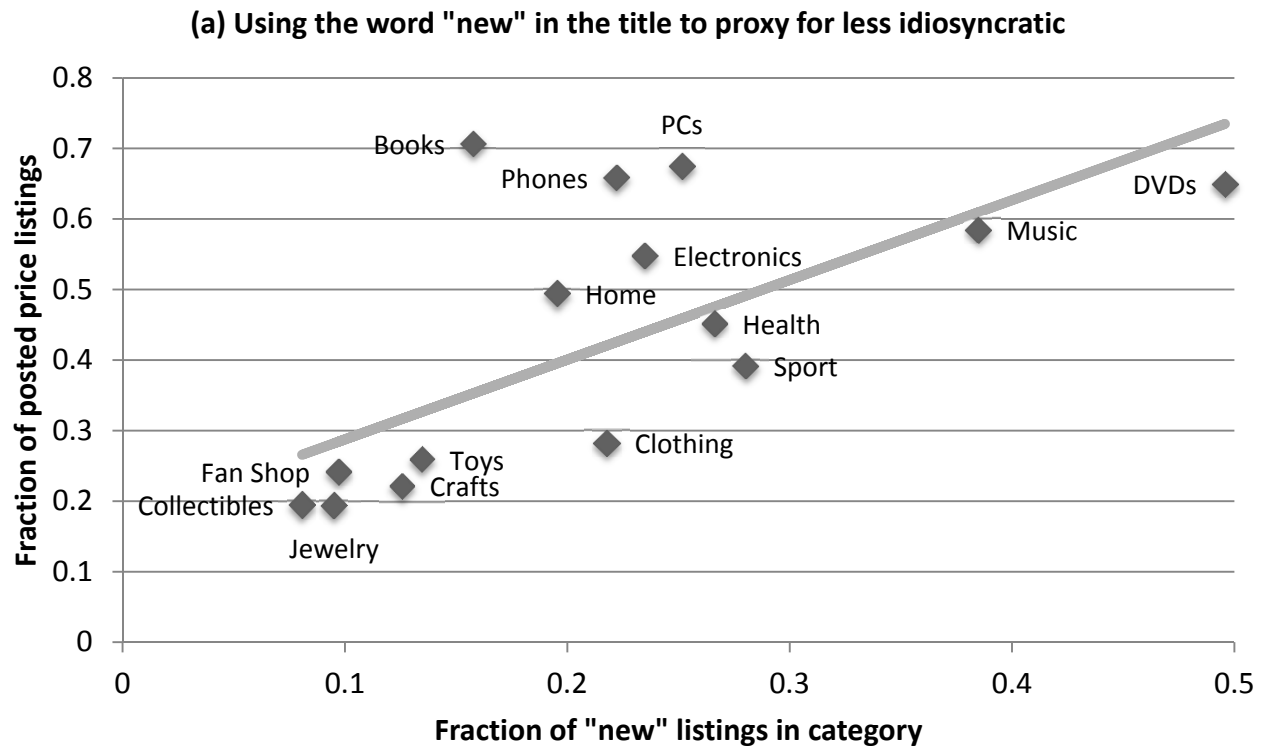


Figure presents the share of posted price listings in the largest 15 item categories, plotted against (a) the share of category listings that have the word "new" in the listing title, and (b) the share of listings that have at least one “duplicate” (i.e. another listing by the same seller with the same title). Both figures show a listing-weighted regression line, and are based on all eBay listings in 2009 that were either posted price or pure auction listings.

Figure 4: Posted Price Use, by Seller Cohort

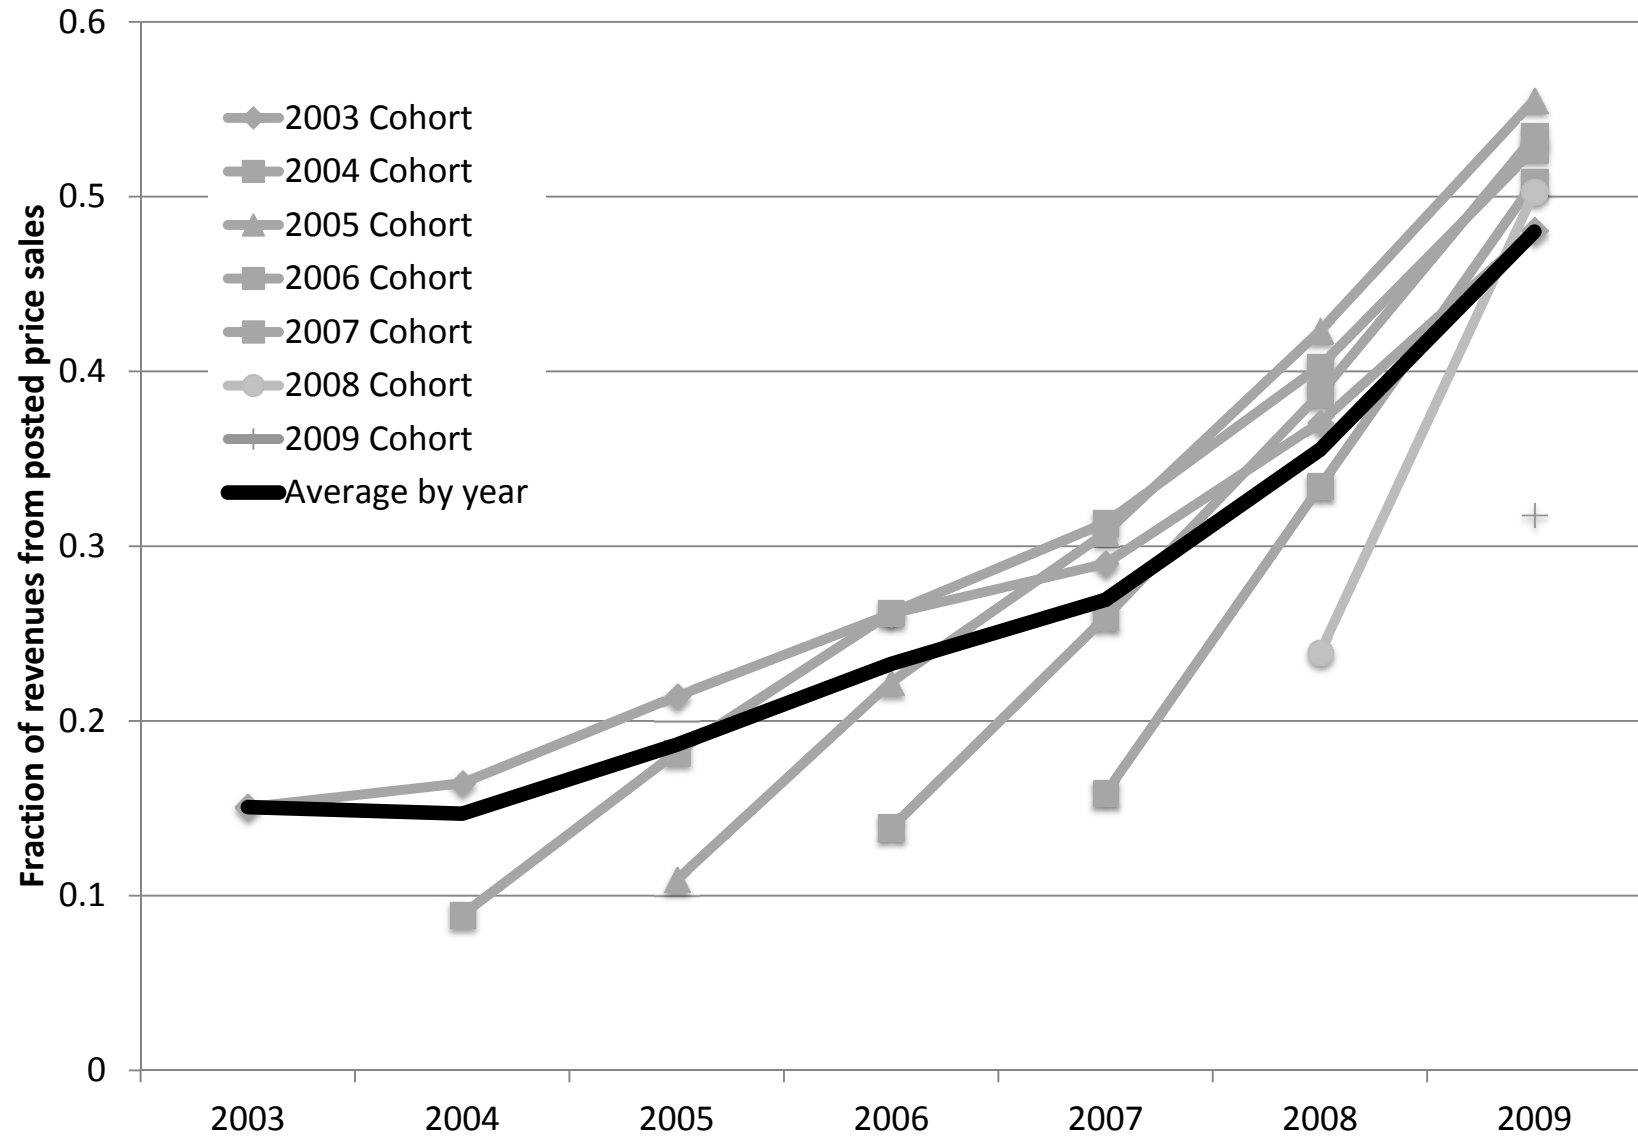


Figure presents the evolution of the annual fraction of eBay sellers' revenue from posted price listings (out of revenues from posted price and "pure" auction listings), separately for each cohort of sellers. A seller is assigned to a cohort based on the calendar year in which the seller's first ever listing was posted on eBay. The thick black line presents the overall platform average for each year.

Figure 5: Theory – The Auction and Posted-Price Demand Curves

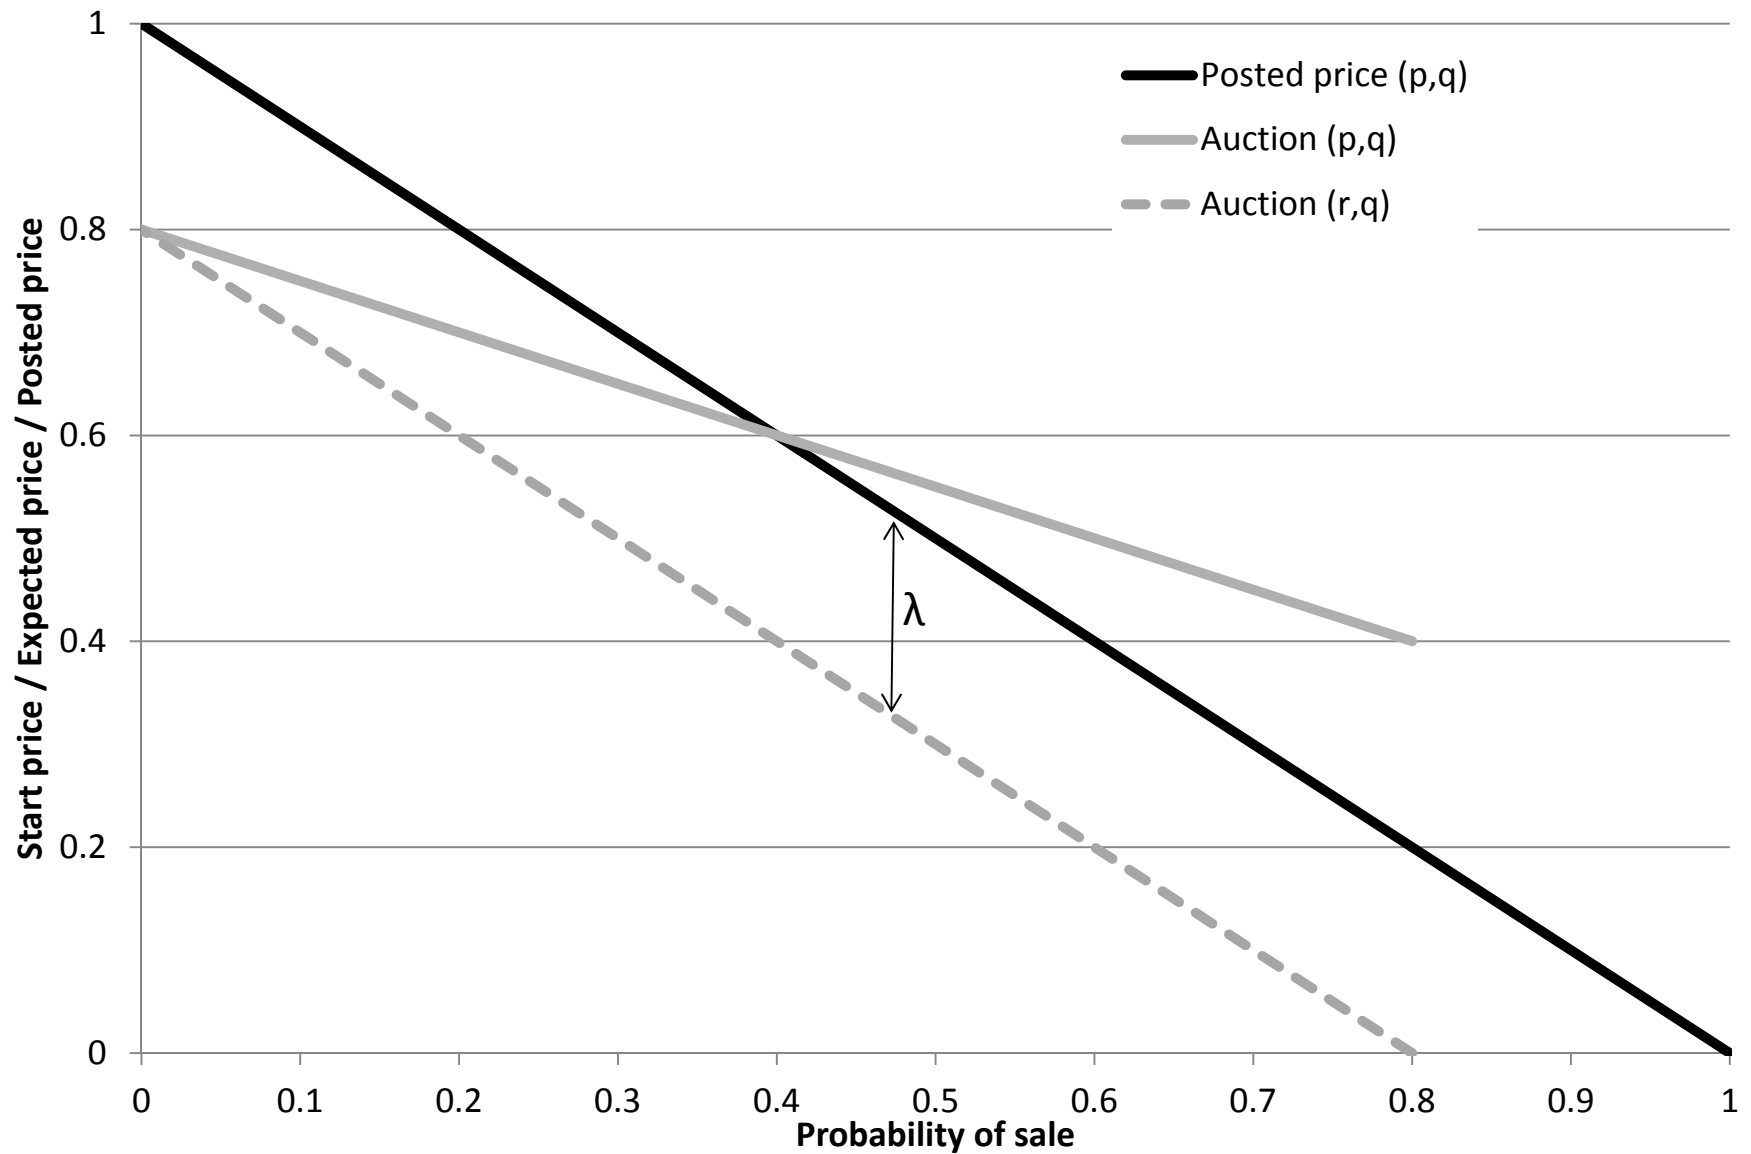


Figure is based on the model of Section 3, assuming that values are drawn from a uniform distribution on $[0,1]$, with $u=0$ and $\lambda=0.2$. The black solid line is the posted-price demand curve. The dashed gray line is the auction sale curve, which shows the probability of sale as a function of the auction start price. The solid gray line is the “auction demand curve,” which plots the probability of sale from a given start price against the expected sale price from that start price.

Figure 6(a): Sale Rates of Posted Price and Auction Listings, by Category




Figure 6(b): Average Sale Price of Posted Price and Auction Listings, by Category

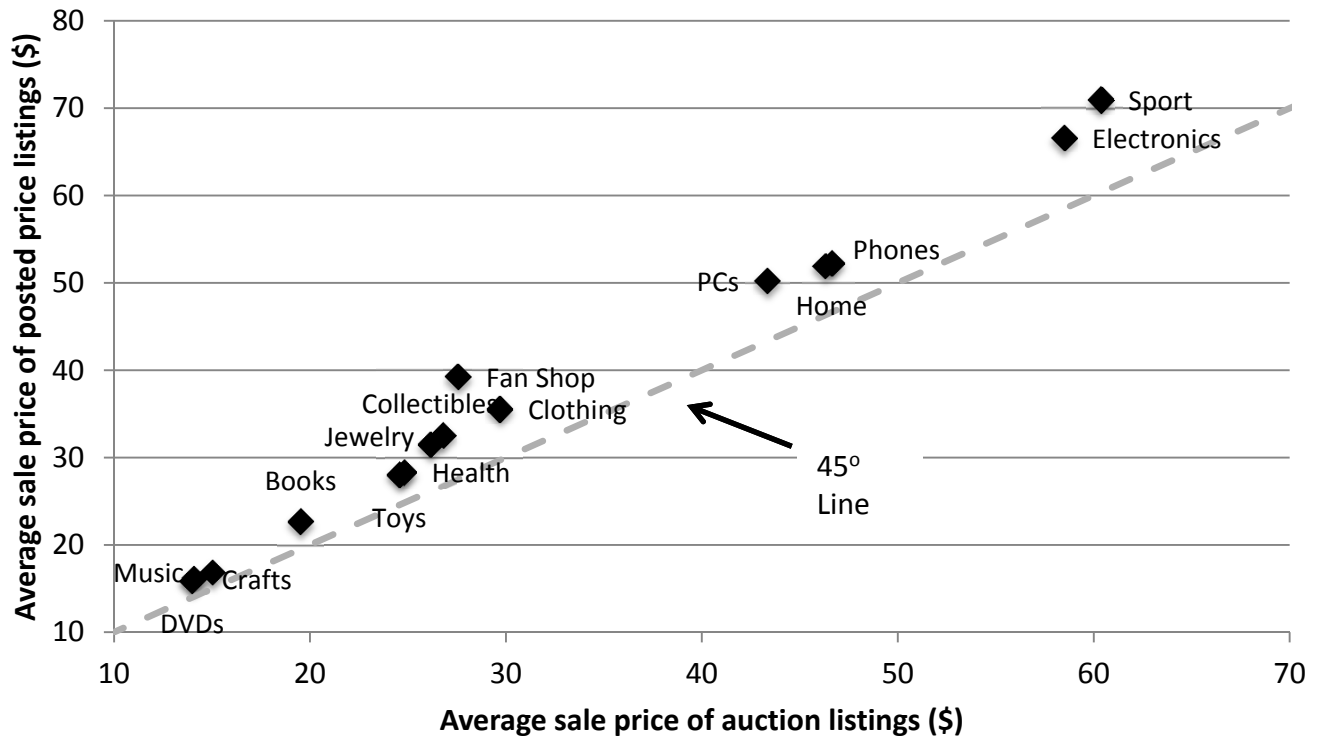


Figure is based on the 2009 seller experiments sample described in Section 4 and reports (a) the average sale rate and (b) the average sale price in each category, by listing format. For each item (experiment) in the sample, we compute the sale rate and price in each listing format, and then average (unweighted) across all items within each category.

Figure 7(a): Trends in The "Auction Discount" Based on Seller Experiments

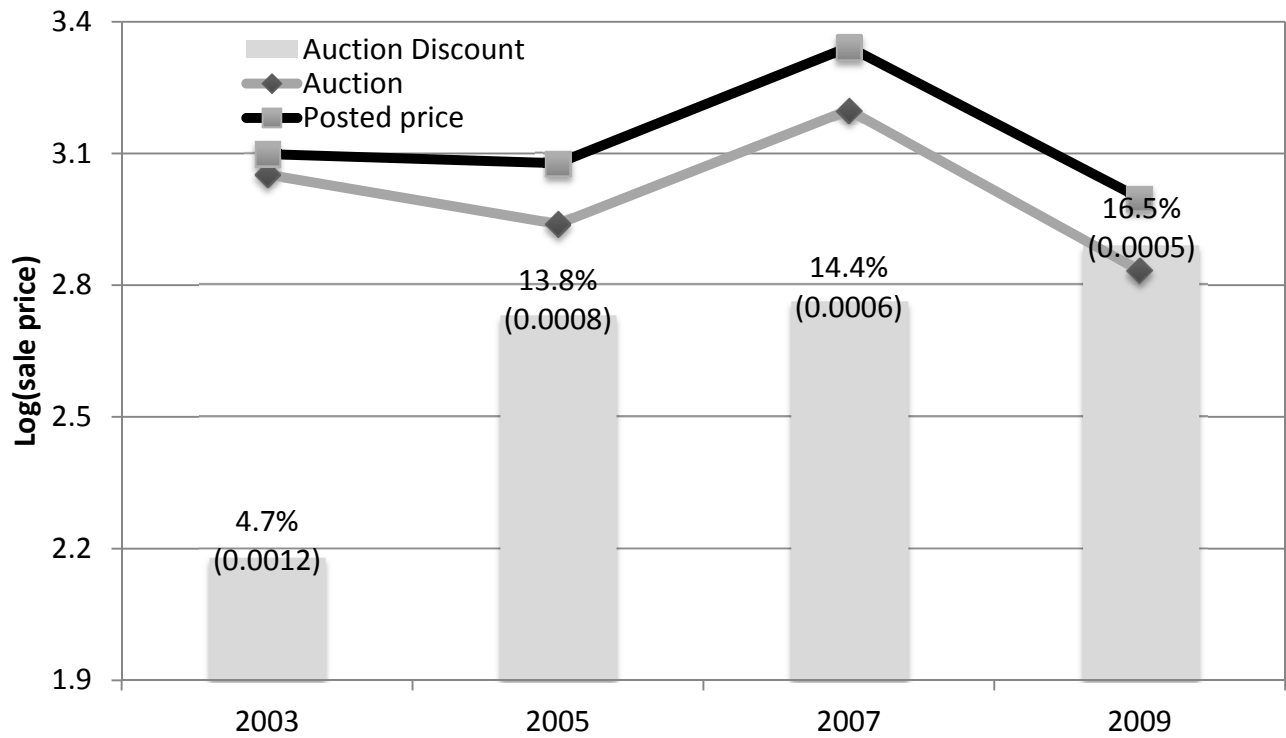


Figure 7(b): Trends in Sale Rate Based on Seller Experiments




Figure is based on the seller experiments sample described in Section 4. In (a) the lines represent the predicted value from a regression of (log) sale price on an indicator that is equal to one for an auction listing and on seller experiment fixed effects. The regression is estimated for each year separately. The bars represent the estimated coefficient on the auction indicator (which is the same as the vertical difference between the two lines). Robust standard errors are in parenthesis. In (b) the lines represent the predicted value from a linear probability regression of a sale indicator on an indicator that is equal to one for an auction listing and on seller experiment fixed effects, again estimated separately for each year.

Figure 8(a): The Relative Demand Curves in 2003

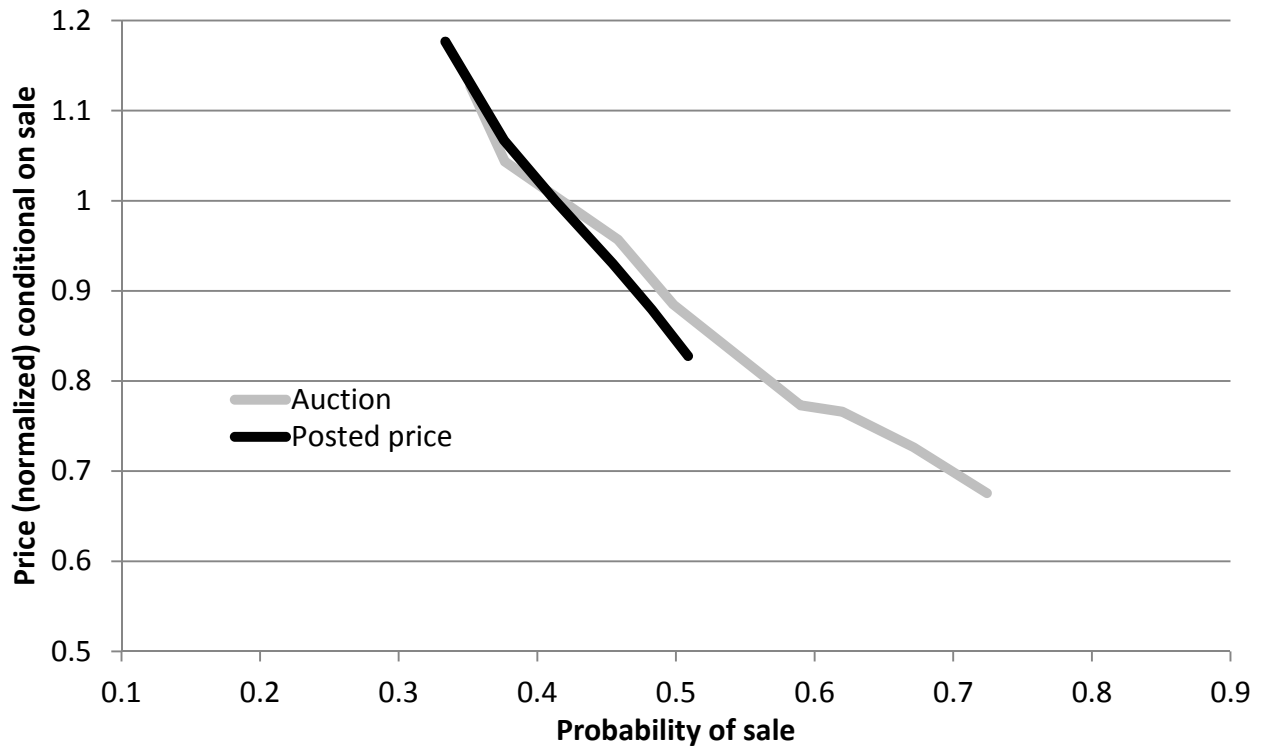



Figure 8(b): The Relative Demand Curves in 2009



Figures are based on the 2003 and 2009 samples described in Section 4. The posted price demand curve is based on estimating a linear probability model of a sale indicator on the posted price and experiment fixed effects. The auction demand curve is based on estimating a similar sale equation (sale indicator on start price and experiment fixed effects) and a separate price equation (normalized sale price on start price and experiment fixed effects), and combining the estimates to construct an auction demand curve as described in the text, and shown in Figure 5.

Figure 9(a): The Relative Demand Curve of Idiosyncratic Categories




Figure 9(b): The Relative Demand Curve of Commodity Categories




Figure is based on the 2003 and 2009 seller experiments sample described in Section 4. Figure (a) restricts estimation to Electronics, Computers, DVDs, Health, and Phones. Figure (b) restricts attention to Collectibles, Fan Shop, Toys, Jewelry, and Clothing. Figures are otherwise equivalent to Figures 8(a) and (b) combined.

Figure 10: Calibrating λ and u from Posted-Price and Auction Sale Curves




Figure illustrates the way by which we calibrate the values of λ and u , and presents the empirical analog to the theoretical Figure 5. The black lines are the estimated posted price demand curves in 2003 and 2009 (shown in Figures 8(a) and 8(b)). The gray lines are the auction sale curves (that is, our estimates – based on estimating equation (4) – of the effect of start price on the probability of sale for auction listings). The vertical distance between the two graphs is an estimate of λ , and the vertical distance between the two posted price demand curves is Δu .

Figure 11: Posted Price Use, by Buyer Cohort




Figure presents the evolution of the annual fraction of buyers' expenditure on posted price purchases (out of expenditures on posted price and "pure" auction purchases), separately for each cohort of buyers. A buyer is assigned to a cohort based on the calendar year in which the buyer's first ever purchase was made. The thick black line presents the overall platform average for each year.

Table 1: Summary Statistics of Experiments Samples

	2003 Experiments Data			2009 Experiments Data		
	(23,057 experiments and 1,168,033 listings, of which 748,545 (64%) are auction listings)			(83,685 experiments and 5,924,448 listings, of which 4,185,969 (71%) are auction listings)		
	Mean	25th pctl	75th pctl	Mean	25th pctl	75th pctl
Number of listings in experiment	50.7	18	50	70.8	19	68
Auction listings in experiment	32.5	8	29	50.0	9	45
Posted-price listings in experiment	18.2	6	19	20.8	6	18
Share of auction listings	0.56	0.38	0.74	0.61	0.43	0.81
Experiment avg. sale rate for auctions	0.51	0.22	0.80	0.44	0.16	0.73
Experiment avg. sale rate for posted prices (within 7 days)	0.44	0.20	0.67	0.37	0.17	0.50
Difference in sale rate (auction minus posted price)	0.07	-0.10	0.23	0.08	-0.12	0.25
Experiment avg. sale price (normalized) for auctions	0.96	0.83	1.04	0.86	0.75	0.98
Experiment avg. sale price (normalized) for posted prices	0.97	0.97	1.00	0.98	0.97	1.00
Difference in sale price (auction minus posted price)	-0.01	-0.15	0.08	-0.13	-0.24	0.00
Experiment avg. sale price (reference value)	40.3	8.2	32	42.7	10.0	40.0
Std. deviation of (normalized) auction start prices	0.25	0.14	0.34	0.19	0.09	0.27
Std. deviation of (normalized) posted prices	0.10	0.00	0.15	0.08	0.02	0.11

Table 2: Calibration Results

	All categories	Idiosyncratic	Commodity
<u>Panel A. Calibrated values</u>			
λ_{2003}	0.08	0.06	0.19
u_{2003}	----- Normalized to zero -----		
λ_{2009}	0.16	0.19	0.23
u_{2009}	0.14	0.11	0.15
<u>Panel B. Success rates (q*)</u>			
Posted price success (2003)	0.40	0.34	0.49
Auction Success (2003)	0.48	0.38	0.56
Posted price success (2009)	0.25	0.25	0.29
Auction Success (2009)	0.37	0.31	0.41
<u>Panel C. Relative effects of λ and u</u>			
Change in λ	0.09	0.13	0.05
Change in u	0.14	0.11	0.15
Relative importance of λ vs. u (2003 quantities)	3.60	10.22	2.52
Relative importance of λ vs. u (2009 quantities)	1.92	6.11	1.12

The top panel presents estimates of u and λ from the estimated demand curves in 2003 and 2009. The middle panel shows success rates estimated from the linear probability regression in Section 4.1. The bottom panel derives the relative profit effects due to the calibrated increase in u and λ , as explained in Section 4.3. Specifically, we compute the increase in posted price profit minus auction profit due to λ and also due to u and report the ratio in the Table.