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

Liran Einav, Chiara Farronato, Jonathan Levin and Neel Sundaresan†

July 2012

PRELIMINARY AND INCOMPLETE.
COMMENTS ARE VERY WELCOME.

Abstract. Consumer auctions played a major role in the early days of internet commerce, but today's online environment has shifted toward 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. We develop a simple model that captures the trade-off between auctions and posted price sales mechanisms, and use a parsimonious specification to assess its quantitative implications. The results suggest that buyer demand has evolved to favor posted prices, perhaps due to greater competition for online consumer attention, and this can explain a significant fraction of the shift toward price posting. The remainder can be explained by narrower retailer margins, perhaps due to more intense competition or improved consumer search. We also discuss why sellers may continue to use both mechanisms as a form of price discrimination.

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. Instead, most online commerce nowadays 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 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 nearly everyone nowadays shops online, but because the choice of sales mechanism, and competition between sales mechanisms, is one of the classic problems in microeconomic theory.

We combine very simple economic theory 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. Another possibility is a change in con-

---

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

[2]The 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.

sumer tastes. 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. Yet another possibility is that 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 a-priori expected: auctions are favored for used and idiosyncratic items, and by less experienced or professional sellers. Other patterns are more surprising, such as the failure of compositional changes to explain the decline in auctions. Taken together, the patterns provide a useful starting point for theoretical modeling. In Section 3, we develop a simple 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. Modeling the trade-off allows us to match the basic patterns in the data, including a robust pattern that auction listings tend to have higher sale probabilities but lower prices. The model is sufficiently simple that 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.

In Section 4, we use the model to develop an empirical analysis of the choice of sales mechanism. Our empirical strategy builds on earlier work reported in Einav et al. (2011), in which we documented the prevalence of "seller experimentation" on eBay — that is, the very high frequency with which sellers post identical listings either simultaneously or over time while varying their sale format or other pricing parameters. Here we use a large dataset of seller experiments spanning from 2003 (when auctions were the dominant sales mechanism) to 2009 (by which time posted prices had overtaken auctions). An immediate finding from this data is the gradual growth of an "auction discount" — the difference between what a seller receives if she offers an item by auction as compared to the posted price at which she is selling the (same) item — from less than five percent in 2003 to sixteen percent by 2009.

We then use variation by sellers in their pricing and reserve price policies to estimate non-parametric listing-level demand curves for both auctions and posted prices in 2003 and in 2009. We show that success rates have fallen for both formats, and that the demand for

auction listings relative to posted price listing has fallen considerably, shifting seller incentives toward using posted prices. The qualitative features of the estimated demand curves also match our theoretical model, allowing us to quantify the key parameters of the model. We use the calibrated model to decompose the change in seller incentives into a reduction in residual demand (e.g., due to increased competition) and a reduction in the relative demand for auctions (e.g., due to a greater convenience preference, perhaps due to competition from other forms of online entertainment). We find that the latter could explain 40% or more of the shift in the seller incentives.

In the final section of the paper, we discuss some reasons why the two sales mechanisms might continue to coexist. One obvious reason is the heterogeneity of goods traded online and the diversity of sellers. It is natural that idiosyncratic and hard-to-price items, particularly valuable ones, should be offered by auction, or that inexperienced sellers might enjoy some benefit from letting the market determine the appropriate price for their product. More interesting is the possibility that experienced sellers of retail items would use auctions in combination with posted prices, a pattern that is also quite typical. We show that in a simple extension of our model, such a strategy can be a desirable form of price discrimination, and provide evidence consistent with this hypothesis.[3] From this perspective, some of the current use of online auctions, at least on eBay, simply may be analogous to the types of couponing and marketing strategies commonly seen in retail.

Our paper relates to a rich theoretical literature on the choice of sales mechanisms, and to a considerable number of papers studying online auctions and online commerce. The choice between posted prices and auctions has been studied theoretically by, among others, Wang (1993), Lu and McAfee (1996), and Kultti (1999). The literature on online auctions is surveyed by Bajari and Hortacsu (2004). Hammond (2010) and Bauner (2011) focus specifically on the choice between auctions and price posting. Hammond (2010) analyses online CD sales and finds that sellers with a higher opportunity cost tend to choose fixed price over auction mechanisms. In his analysis, it is seller heterogeneity that explains the coexistence of auctions and fixed prices for similar items. Bauner (2011) uses eBay sales of

---

[3]The same extension can also explain the use of eBay's hybrid format that allows sellers to run an auction but also to post a price at which a buyer may "preempt" the auction by buying the item immediately. We discuss this below.

Major League Baseball tickets to disentangle factors behind sellers' choices of sale formats: he finds that not only opportunity costs, but also heterogeneity across buyers, broadly divided to "fixed price lovers" and "neutrals," help explain the coexistence of auction and fixed price sales for the trade of identical goods. Ariely and Simonson (2003) and Lee and Malmendier (2011) both analyze auction prices relative to comparable posted prices, and present results that relate to our estimates of the auction discount. Finally, our empirical strategy of exploiting near-identical listings by a given seller was first used by Elfenbein et al. (2011) to study charity auctions, and more broadly described in our related work (Einav et al., 2011).

2   Setting, Data, and Preliminary Evidence

Our analysis focuses on eBay, which started in 1995 exclusively as an online auction site. Many auction sites were active at the time, but eBay rapidly became the most successful (Lucking-Reiley, 2000). Its marketplace transaction volume reached 9.3 billion dollars in 2001. Initially, eBay was known for its vast diversity of 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" or BIN format).[4] 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.

---

[4] A 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.

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 that auctions would be used mainly 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.

One clear pattern in the data is the greater prevalence of auction listings 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 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 [5] 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 qualitatively just what one would expect.

Another way to identify more idiosyncratic items starts from the observation that many listings (about 75% 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 where the seller has multiple copies, whereas unique or truly idiosyncratic items (e.g. a lock of teen heartthrob Justin Bieber's hair) would show up as individual listings. Of the individual listings that appeared in 2009, about 75% are auctions, compared to 57% of the duplicates.

---

[5] In 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.

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.

Similar patterns can be identified for sellers. Here, we classify sellers, based on their past transaction volume on eBay, into experienced ("business") sellers, inexperienced ("occasional") sellers, and sellers of 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 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 a-priori 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] This 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.

---

[6]The dependent variable in this regression is the share of posted-price listing of a seller in his second year on eBay minus the share in his first year. The right-hand-side is the corresponding difference in 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.

[7]Except for "Phones" where auction sales rates and prices are both higher.

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, one 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 44 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, and changes in the shares of different product categories, and changes that occurred within seller-product classifications. 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 format.

Define $Z = 20.9\%$ to be the overall share of posted price transaction volume 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 defining the corresponding quantities for 2009, we have the following decomposition:

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

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. In the data, 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. An increase in business sellers accounts for 2.8 percentage points of the change toward posted prices, and changes in product composition for only 2.5 percentage points.[8]

There is also another way to look for compositional changes that 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 a modest "first year" effect (noted above), at any given point in time, different cohorts of sellers behave in remarkably similar fashion, with all cohorts evolving over time toward posted pricing.

In short, the time-series patterns suggest that to understand the decline in auctions relative to price posting, we should look at seller incentives rather than changes in the nature of the items or the identify of the sellers. This is how we proceed, first by developing a simple theoretical model of the choice between sales mechanisms, and then by using an empirical strategy that looks at the trade-offs between auctions and posted prices for a given seller and product.

3   A Simple Model

In this section, we develop a model of the choice between sales mechanisms. The model is extremely parsimonious, and is developed to emphasize a trade-off between price discovery and convenience for buyers. We consider a single seller offering a single item, and discuss

---

[8]In fact, the prevalence of auctions fell in almost every seller-product category cell. Of the $4 \cdot 33 = 132$ cells, the posted price share has increased in 118 of them.

richer 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.

Consider a seller with a single item to sell, and a cost c of making the sale. We assume there are two or more buyers indexed by i, each with the same value v for the item, and a common fixed reservation utility u.[9] We assume that v is drawn from a log-concave distribution F, where the distribution (but not the realization) is known to the seller. 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 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 $p_A(r) = \mathbb{E}[v - u - \lambda \mid v - u - \lambda \geq 0]$. 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, of 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 reduces the willingness to pay of all

---

[9]It is straightforward to allow for individual values to have an idiosyncratic component, an extension that we are likely to add in future drafts.

[10]The more general version of this observation is that assuming F is log-concave, and that $\lambda \leq \mathbb{E}[v] - \underline{v}$,

buyers including those with the highest value, but it also creates competition that increases the final price above the reserve price, and this effect is largest for low reserve prices (or high sale probability). These effects are both 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 demand curve represents the effect of competition – the expected amount by which the auction price will rise above the reserve price.

This structure means that the model has relatively clear implications for seller incentives. While we will not take the view that every seller chooses the optimal sales strategy for every listing — and indeed our empirical approach in the next section will be premised on the idea that sellers engage in considerable experimentation with alternative strategies — it is nonetheless useful to consider the profit-maximizing incentives. For a fixed price listing, 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 simple, but useful, 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

---

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.

[11]Reduced 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.

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 posted prices over time.

4   Evidence from "Seller Experiments"

Our primary 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, using different sale formats or with different pricing parameters. In Einav et al. (2011), we documented the widespread use of this approach by eBay sellers, and described some of the features of eBay's platform that favored it, and also why it might be attractive to sellers who were uncertain about their optimal sales strategy or how it might be evolving over time. 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]

To make use of this general approach here, we start with all listings posted in 2003, 2005, 2007, and 2009 by a random 50% sample of every seller to ever post on eBay.[13] We identify every collection of "matched" listings by these sellers, where a "match" means the listings have a common seller, a common ten-word item description and were posted in the same year. Given the way we subsequently use the data, we restrict attention to collections of matched listings ("experiments") that have at least ten listings overall and drop (mostly for data quality reasons) a small number of experiments (0.8% or less in each of the years) for which the average (posted) price is above $10,000, or for which there is an auction that resulted in a price more than five times the average fixed price.

Because we want to compare the performance of auction and posted price listings, holding

---

[12]In Einav et al. (2011), we used seller experiments data from a single year, 2009, to study the market's 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.

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

fixed the seller and item, we limit ourselves to sets of matched listings that include at least one auction and one posted price sale and at least four listings of each format. Relative to our model of a seller making an optimal choice between sale formats, focusing on sellers who use both formats might appear odd. An overwhelming feature of eBay's data, however, is that although most sellers do use either mostly auctions or mostly posted price listings, many experiment with the other format. That is, sellers do not engage in a single type of sale strategy for a given item, or even converge to one eventually. Though we will discuss below why using both sale formats could be an optimal form of price discrimination, our general view is that while sellers may gravitate toward more profitable strategies, they engage in considerable and persistent experimentation.

Finally, throughout this section, we need to make several additional measurement adjustments in order to make the outcomes of auction and posted-price listings comparable to each other. First, because the typical posted-price listing lasts for 30 days (unless it sells earlier) while the typical auction duration is a week, we consider a posted-price listing to be successful only if it sells within the first seven days. Second, especially in the later years of the sample, multi-unit posted-price listings become common (relative to auction listings that list one unit at a time). We define a (posted price) multi-unit listing to be successful if it sells at least one unit within the first week.[14]

Table 1 provides more details and some summary statistics on the resulting data. A useful point to emphasize is that the data are, by construction, a selected sample. They exclude the many items on eBay that are truly "one of a kind" and appear as one-off listings. As we noted above, these are overwhelmingly auction listings throughout our sample period. As we have done throughout, we drop hybrid auctions and other sales formats such as "best offer." The resulting data set is still quite large, and covers items from all product categories, so we think that the qualitative results reported in this section are not too "special."

Figure 6 uses the 2009 sample to illustrate that the basic predictions of the model from the previous section hold in the data: auction listings are more likely to sell across all cate-

---

[14]Due 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.

gories (except for "Jewelry"), but they sell for lower prices on average. Interestingly, three of the categories with the highest share of auction listings ("Jewelry," "Collectibles," and "Fan Shop") are qualitatively associated with the worst auction performance: in these categories the auction sale rate is only slightly higher while the price discount is relatively large. Importantly, given the sample design, the figure is constructed using the experiments data so the sale rate and sale price are compared for identical set of items that are sold via both formats.

4.1   Changes in the "auction discount"

As a first analysis, we use the experiments data to estimate the consequences of using one sales mechanism rather than the other holding fixed the seller and the item. To do this, let i index sets of matched listings, and j index transactions (i.e. successful listings) within an experiment. Let A_{ij} be an indicator variable equal to 1 if a given transaction resulted from an auction, and let p_{ij} denote the transaction price (posted price or final auction price).

We start by comparing the transactions prices from the two sales mechanisms, using a fixed effects regression of the form:

$$\log p_{ij} = a_i + bA_{ij} + e_{ij}. \tag{2}$$

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 8(a). 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 a much larger 16.5%. Interestingly, the average transaction price for posted price sales remained very similar over this period (despite there bring turnover in the set of

items).[15] In our sample, the average log posted price was 3.1 in 2003 and 3.0 in 2009.

The price differentials in Figure 8(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 where the seller use 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 process tends to separate matched listings, buyers very often may not make a direct comparison). The latter arguably fits more cleanly with our theoretical model, where the outside option (u) is exogenous. It turns out, however, that the estimated auction discount (in either 2003 or 2009) is not systematically different when we split the sample to look at these situations separately, so the main feature of the data — the growing discount in auction prices over time — does not depend much on exactly how we set up the posted price comparison.

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,[16] 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 i, and β the extra sale probability for auction listings relative to posted price. Again, we estimate the specification separately for each year. Figure 8(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 [17]. The grey line plots the corresponding sale probability for auction listings,

---

[15] That 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.

[16] As 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.

[17] To construct this average we use all items i in the sample, and weight each item by the number of total listings (in both formats).

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 – in our sample, the average posted price sale rate in the sample has fallen from 40.3% to 25% between 2003 and 2009. However, the difference in success rates across formats (the estimated β), has not changed much over time. It has stayed relatively constant, starting at about 8% in 2003, declining a bit, and then increasing 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 3.21 and 43% successful auctions had a single bidder. By 2009, the average number of bidders has fallen by 7% to 2.98 and the share of single-bid auctions had increased to 48%.

Taken together, the results in this section 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. Of course, these estimates have an important limitation in that they do not account explicitly for possible changes in how sellers are setting reserve prices relative to posted prices (something that in principle could affect auction prices relative to posted prices). In the next section, we use within-experiment variation in auction reserve prices and posted prices to estimate (residual or listing-level) demand curves for auction and posted price listings, which will allow us to link the time trends more precisely to changes in market demand.

4.2   Changes in auction demand

The growing auction discount provides a fairly strong indication that seller incentives to use auctions have shifted. But as we noted, it does not cleanly distinguish shifts in demand — from changes in buyer preferences, or changes in competition — from changes in pricing behavior. In this section, we use variation in sellers' pricing for individual items to estimate residual (or listing-level) auction and posted price demand curves. We then use the estimated demand curves, and how they have shifted over time, to map back to the specific parameters in the theoretical model that determine sellers' incentives to use auction or posted price listings.

To estimate listing-level demand, we exploit the significant within-experiment variation

in auction start prices and posted prices.[18] Consider the auction case first. Starting with all auction listings in our experiments data, we first estimate, for each year of data, the probability of sale for an auction listing as a function of its start price:

$$q_{ij} = \alpha_i + g\left(s_{ij}^n\right) + \varepsilon_{ij} \tag{4}$$

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, which is equal to the start price s_{ij} divided by a reference value for item i that allows us to make item values comparable. To define item i's reference value, we use the average price across matched fixed price listings. With this normalization, an item selling for \$100 with a start price of \$50, and an item selling for \$30 with a start price of \$15 both have a normalized start price of 0.5. We specify the function $g(\cdot)$ flexibly by allowing separate dummy variables for different levels of the normalized start price.[19]

We estimate a parallel regression for the final auction price, again for each year, restricting attention to successful auctions:

$$p_{ij}^n = a_i + h\left(s_i^n\right) + e_{ij}. \tag{5}$$

Here p_{ij}^n is again a normalized price, equal to the final auction price divided by the item's reference value, and h is also 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. To do this, we take each value of the normalized start price $s^n \in [0, 1.2]$, and consider the locus $\overline{q}^n\left(s^n\right) = \overline{\alpha} + \hat{g}\left(s^n\right), \overline{p}^n\left(s^n\right) = \overline{a} + \hat{h}\left(s^n\right)$, where $\overline{\alpha}$ and $\overline{a}$ are the average fixed effects, where, as before, the average is weighted by the total number of listings in the corresponding sample.

---

[18] Einav et al. (2011, Table 5) document the remarkable degree to which sellers vary their auction start prices for a given item.

[19] Specifically, 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.

We also estimate posted price demand curves separately for each year. This exercise is conceptually more straightforward, although the extent of price variation is not as dramatic, so that the range over which we can obtain a highly flexible estimate of the demand curve is not as broad. In this case, we proceed by estimating a fixed effects linear probability model where the dependent variable is again an indicator equal to one of the listing results in a sale,

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

In this specification, p_{ij}^n is the normalized price of the listing (equal to the price of the listing divided by the item's reference value). To estimate $m(\cdot)$ we specify separate dummy variables for different levels of the normalized price.[20]

Figure 8 presents the results. The top panel plots the estimated demand curves using the 2003 sample, while the bottom panel repeats the same exercise for 2009. The first observation is that, both in 2003 and 2009, a main qualitative prediction of the theory holds in the data: the estimated posted price demand curves are steeper than the auction demand curves. The second observation is the difference between 2003 and 2009. In 2003 the estimated auction demand curve mostly (except for high prices) lies above the posted price demand curve, suggesting that a seller facing this average residual demand curve would be better off using an auction, unless her marginal cost is extremely high. In 2009, however, this is not the case. As in the theory (Figure 5), the two curves cross, so that for a seller facing these average residual demand curves, the optimal sale mechanism would depend on the seller's cost.

A few points are worth mentioning. First, as the figures make clear, there is more variation in auction pricing than in posted pricing. This shows up in the fact that we are able to trace out demand nonparametrically for auctions over a greater range. Because the auction "price" is also an expected price for a given probability of sale, it is also the case that in the data, the variation in auction sale prices for a given item is much larger than the variation across posted price transactions. We exploit this below in discussing differences across buyers. Finally, we note that our estimated demand curves involve averaging over

---

[20]The pricing variation for posted-price listings is not as rich as it is for the acution 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 for the 2009 sample, and more limited (because there is not as much variation) for 2003: [0-0.975], (0.975-1), 1, and >1.

heterogeneous items. It is also possible to repeat the exercise focusing on more narrowly defined groups of items.

4.3   Identifying Seller Incentives

[Note: Results/estimates in this subsection are still preliminary]

The estimated demand curves 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 6. In this section, we tighten the connection by using the demand estimates to calibrate the key parameters of the model. While the model is obviously far too parsimonious to capture the heterogeneity of the thousands of sellers and millions of item listings in the data, we view the exercise as nonetheless useful to provide some simple calculations of how and why seller incentives may have shifted over time.

The theoretical model derived item-level demand curves from three parameters: the 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 posted price p, the probability of sale is $1 - F(p + u)$, while the probability of sale at auction given start price r 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 $F^{-1}(1 - q) - u - \lambda$, with the difference between them being λ.

Figure 9 plots the estimated posted price demand curves for 2003 and 2009 (identical to those in Figure 8), along with the estimated auction sales curves from the auction sales regression (4). If the estimates were perfectly consistent with the model, the posted price demand curve for year t would be an exact upward translation of the auction sales curve.

In practice, the shapes of the curves 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. We find that $\lambda_{2003} = 0.076$, and $\lambda_{2009} = 0.163$. That is, we estimate that running an auction in 2003 involved a transaction cost equal to roughly eight percent of an item's typical price (or reference value), and that this cost doubled between 2003 and 2009.

We can also 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 $v - u$, so u cannot be separately identified from the distribution of v's. To proceed, we normalize $u_{2003} = 0$, and we further assume that $F(\cdot)$ remained constant over time, which is a necessary but strong assumption given that the set of sold items has changed over time. (In principle, we only need a known percentile of $F(\cdot)$ to have remained constant, but making the stronger assumption keeps things very simple.)

Now suppose we compare the posted-price demand sales curves in 2003 and 2009. To achieve sale probability q with an auction in year t required a start price $F^{-1}(1 - q) - u_t - \lambda_t$. Therefore the vertical distance between the 2003 and 2009 auction demand curves is $\Delta u + \Delta \lambda$ (where $\Delta u = u_{2009} - u_{2003}$ and $\Delta \lambda = \lambda_{2009} - \lambda_{2003}$), and the vertical distance between the two posted-price demand curves would be Δu. Again, the demand curves would be exact translations if the data exactly fit the theory. In practice they are again similar in shape but not identical, so to estimate the vertical distance between them we average over their points of common support. This calculation leads us to estimate $\Delta u + \Delta \lambda = 0.224$. From above, we estimated $\Delta \lambda = 0.087$, so it follows that $\Delta u = 0.137$. That is, between 2003 and 2009, the consumer reservation utility increased by roughly 14% of the reference value for an average item.[21]

Our final step is then to evaluate whether these changes in λ and u can account for the decline in auctions, and the extent to which each of these parameters contribute to it.

---

[21]Note 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.

To do this, we use the estimated relationship between the 2009 auction start price and the probability of sale to back out the distribution of valuations $F(\cdot)$, using the relationship $q = 1 - F\left(r + u_{2009} + \lambda_{2009}\right).$[22] With this final step, we can calculate the optimal profits a seller could make from auctions or posted prices given the calibrated $F(\cdot)$ and various values of λ, u, and marginal cost c. This is shown in Table 2.

The table shows that – as the theory suggests – if the marginal cost is sufficiently low, selling the item by auction is more profitable for any combination of the calibrated values of λ and u (top panel of the table). That is, even in 2009, when the relative auction demand is low, auctions are more profitable for sufficiently low values of c. For higher values of c, however, the optimal sale mechanism would depend on the parameter values (middle panel): it is preferable to sell by auction for the 2003 parameter values, but more profitable to sell by posted price for the 2009 parameter values. For higher values of c, posted price would dominate (bottom panel)

We also use the table to provide an illustrative quantification of the relative importance of the changes in λ and in u. The right-most column of the table reports the difference in profits from auction vs. posted-price listing under different values λ, u, and c. The relative importance of the change in λ is higher for sellers with higher values of c. For low values ($c = 0.3$), the change in λ accounts for about 40% of the overall change in the relative profitability of posted-prices, while u is responsible for the other 60%. For higher values ($c = 0.5$), the change in λ alone would account for about 75% of the overall difference, while for high cost values ($c = 0.9$) λ is even more important.

4.4 Sensitivity and Robustness Checks

[Subsection to be added in the future; including discussion/checks of the zero pass-through assumption]

---

[22]The relationship here is given by $q = 1 - F\left(r + 0.32\right)$. We estimate $1 - F(\cdot)$ only over the range from $q = 0.08$ to $q = 0.77$, so we extrapolate by setting $v_{min} = 0$, $v_{max} = 2$, and assuming uniformly distributed values on $[0, F(0.08)]$ and on $[F(0.77), 1]$.

5   Changing Demand and Buyer Heterogeneity

Our analysis so far has argued that changes in buyer demand, rather than a changing composition of sellers or items, has driven the shift toward posted price listings. In this section, we ask whether the change in buyer demand is due to changes in the composition of buyers, or rather due to shifting buyer preferences. We show that parallel to our seller analysis, there is considerable heterogeneity in the cross-section of buyers, but that the time-series shifts in demand appear to have taken place within cohorts of buyers rather than through differences in the buyers who started shopping online later in the sample period. 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.

5.1   Buyer Heterogeneity

Buyers on eBay are heterogeneous, ranging from occasional buyers who make purchases on eBay infrequently, to more professional buyers who make purchases on eBay more regularly, perhaps in order to turn around and sell these items at profit via other retails channels. Parallel to eBay's classification of sellers, we describe the differences in purchasing behavior between "professional" and "occasional" buyers. We divide buyers on the basis of their number of purchases in a given year. Professional sellers are those who buy more than 100 items, while occasional purchase less than 5.

The fractions of occasional and professional buyers on eBay have been relatively stable over time, from with professional buyers accounting for just over 1% of the buyers, and occasional buyers for around 60%. Yet, the purchase volume of professional buyers has increased over time, accounting for 14% of eBay transaction volume in 2003 and 20% in 2009. At a given point in time, "professional" buyers are more likely to buy at auction relative to "occasional" buyers. In 2003, the auction expenditure share of professional buyers has been 95%, compared to 80% of occasional buyers. This difference has increased to 69% vs. 38% in 2009. Using the seller experiments sample, conditional on selling an item, we find that the buyer is 5.4% more likely to be a professional buyer if the item was listed for sale by

auction.[23]

Professional buyers are also able to get better deals. First, we have already documented an auction discount, so by being more likely (as just described) to purchase at auction rather than at posted prices, professional buyers are more likely to enjoy this auction discount. Furthermore, even conditional on buying at auction, professional buyers get the item for a better price, enjoying a 20% discount.[24]

However, as illustrated in Figure 10, it is not the case that the site was originally populated by auction-lovers, and later arrivals are convenience buyers. Similar to our analysis of sellers, the figure presents the expenditure share on auction listings relative to posted prices for different cohorts of buyers (defined based on the calendar year at which they made their first purchase). As can be seen, buyers in their first year are more likely than average buyers on the site to buy using a posted price, but after first year, different cohorts look very similar, and all cohorts have trended together toward posted price purchasing.

Our preferred 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.

5.2 Price Discrimination

Above we provided evidence that at a given point in time, it is possible to identify groups of buyers who appear to have distinctly different preferences for auctions and posted prices. Here we note that if we extend our model to take buyer heterogeneity into account, it becomes relatively easy to rationalize why sellers might want to offer both posted price and auction listings as a way to cater to different types of buyers.

We illustrate this in a very simple extension of the model. We assume there are two types of buyers: active "deal shoppers" (denoted as type A) and "convenience buyers" (denoted as type B). For simplicity, assume that all buyers have the same intrinsic value v for a given item, drawn from $F(\cdot)$, but different (λ, u). The convenience shoppers have a sufficiently high

---

[23]This is based on a linear probability regression, 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.

[24]This estimate is based on regressing (log) price on experiment fixed effects and an indicator for a professional buyer.

22

λ_B that they are not interested in auctions, and a relatively low u_B so they have relatively high willingness to pay. On the other hand, deal shoppers have a low λ_A and a relatively high u_A (because they are looking hard for deals), which means a relatively low willingness to pay.

Now imagine a seller who lists her items at a posted price p. Provided that $u_A - u_B$ is sufficiently large, she will sell to the deal shoppers but to few (if any) convenience shoppers. On the other hand, if she lists by auction, she will sell to the deal shoppers, but not to the convenience shoppers (because their λ_B is very high). So she can expand her market, without cannibalizing her own sales, by offering both posted price and auction listings.

In 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 have worked on models of this sort, some of which also endogenize the level of u by allowing for competition between multiple sellers, but we do not focus on them 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.

One point to note is that allowing for buyer heterogeneity, even as in our simple example above, might suggest a modification of our calibration exercise, because at least in the example here auction posted price demand are coming 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 the requisite heterogeneity across buyers, sellers, and items would be a worthwhile exercise but is not our goal in this particular paper.

6 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 buyer, 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 about 40% or more of the overall effect.

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 submarkets. 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.

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. Finally, 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 more generally. The recent decision of Prosper (from 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.

The Economist (2000), "In The Great Web Bazaar." *The Economist*. February 24, 2000.

Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan (2011). "Learning from Seller Experiments in Online Markets." NBER Working Paper No. 17385.

Elfenbein, Daniel W., Ray Fisman, and Brian Mcmanus (forthcoming). "Charity as a Substitute for Reputation: Evidence from an Online Marketplace." *Review of Economic Studies*, forthcoming.

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.

Lee, Young Han, and Ulrike Malmendier (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.

Wang, Ruqu (1993). "Auctions versus Posted-Price Selling." *American Economic Review* 83(4), 838-851.
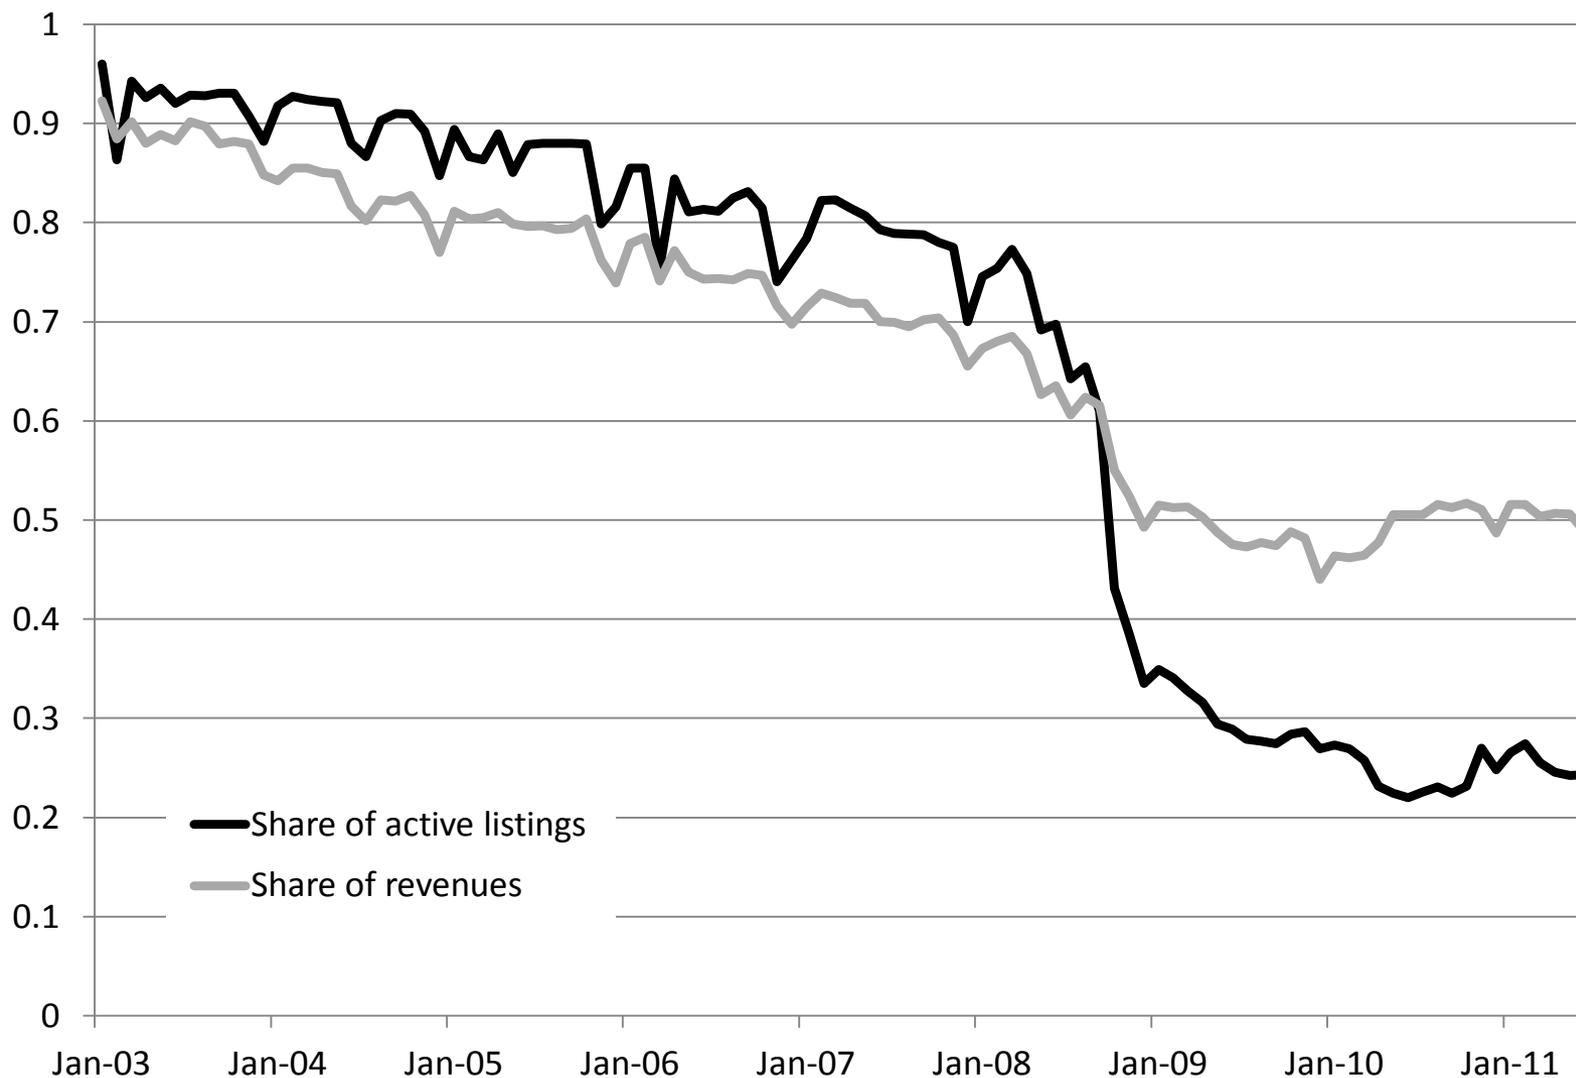
Figure 1: Auction Share on eBay over Time

Figure presents the decline in auctions on eBay. The unit of observation in the figure is a calendar month. For each month, the figure shows the average daily ratio of new eBay listings (black) and transaction volume (gray) from pure auction listings out of all pure auction and posted price listings. Other, less common listing formats, such as hybrid auctions, are not used in generating the figure. We note that the sharp drop in the share of new auction listings (in fall 2008) is primarily driven by the fact that in September 2008 eBay reduced posted-price listing fee and started allowing "good till canceled" listings.

Figure 2: Google Search Volume for Online Auctions and Online Prices
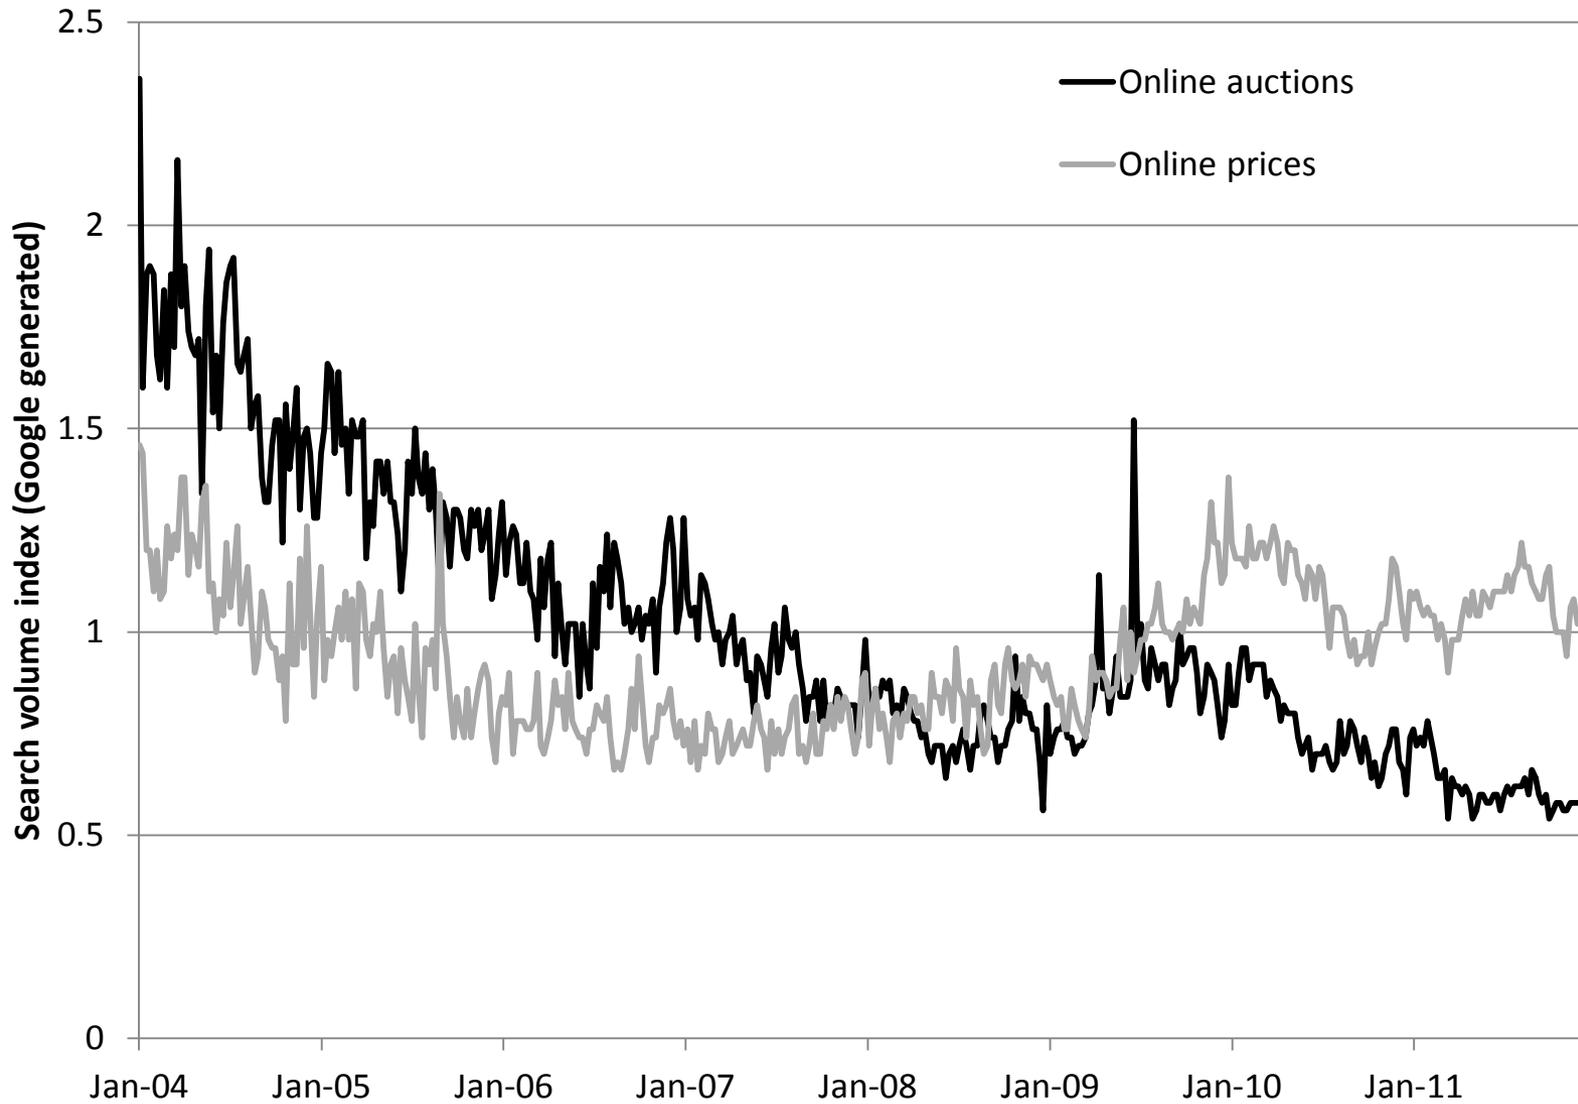
Figure presents results from "Google Trends" for search terms "online auctions" and "online prices." The y-axis is a Google generated index for the weekly volume of Google searches for each of the two search terms, which should make the weekly volume figures comparable over time and across the two search terms.

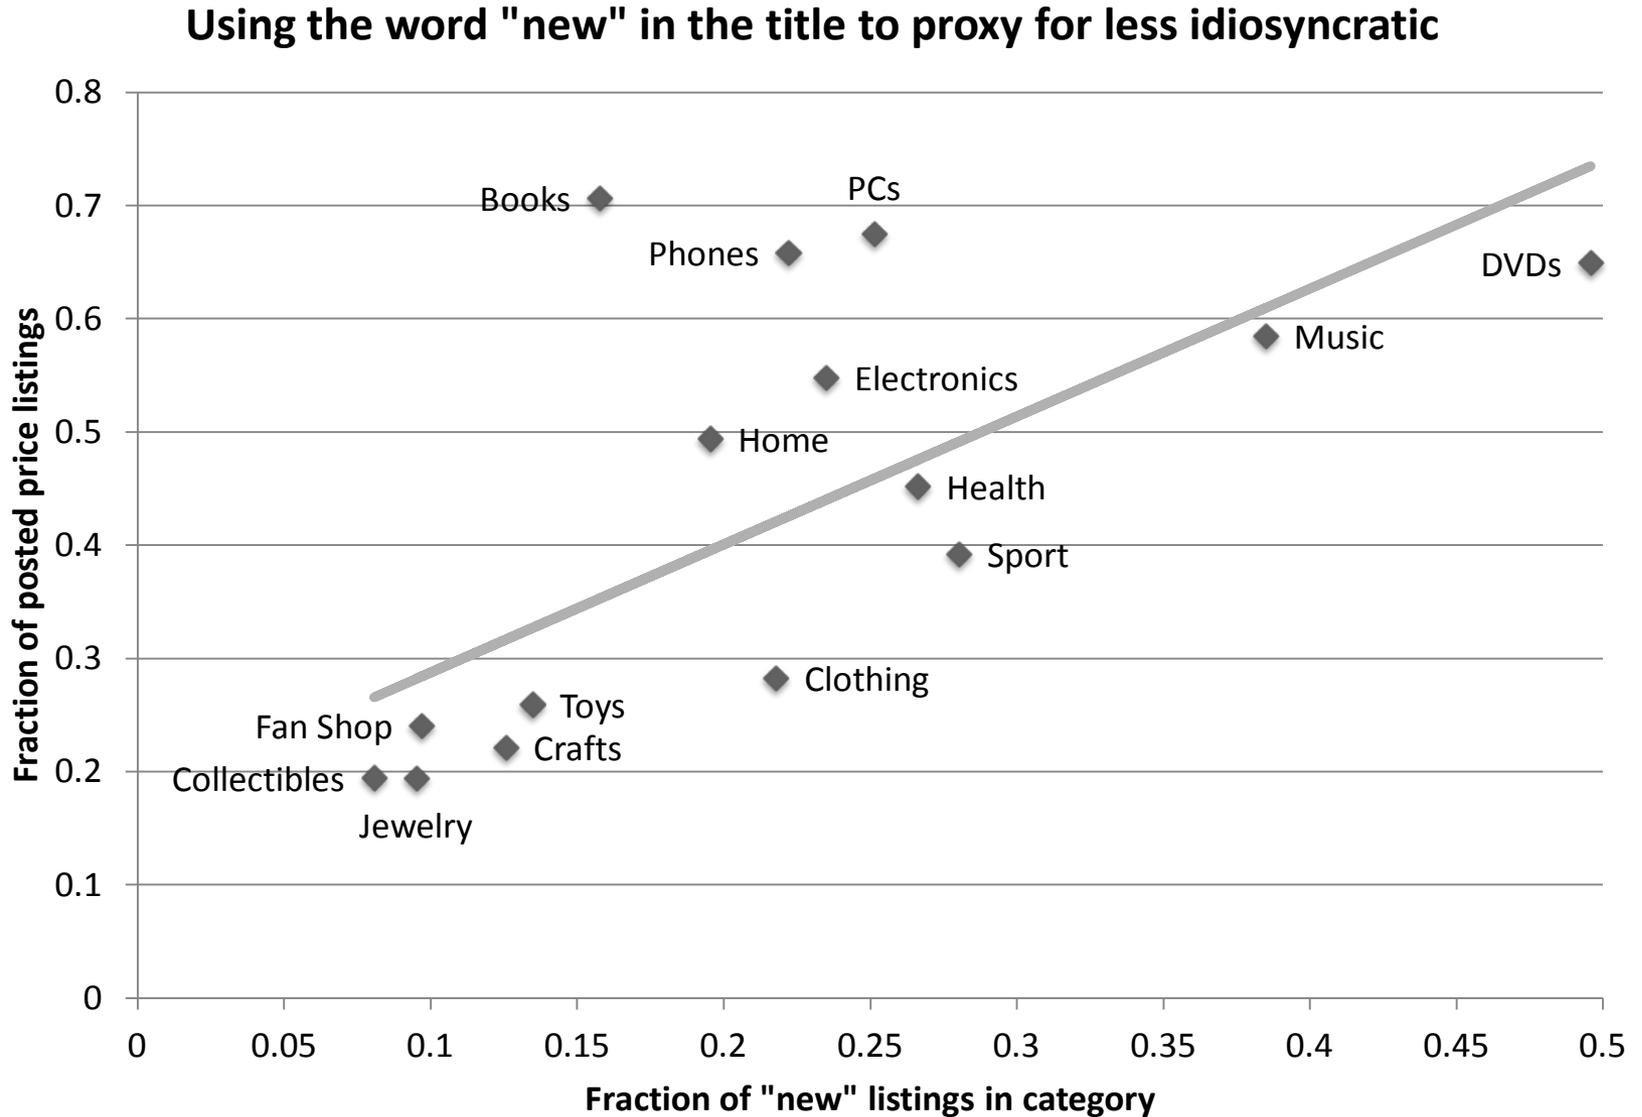# Figure 3(a): Posted Prices are more Prevalent for "Less Idiosyncratic" Items

Using the word "new" in the title to proxy for less idiosyncratic

Figure presents the share of posted price listings in each of the largest 15 (out of 33) item categories, plotted against the share of category listings that have the word "new" in the listing title. The figure also presents a fitted (weighted) regression line. The figure is based on all eBay listings in 2009 that were either posted price or "pure" auction listings .

Figure 3(b): Posted Prices are more Prevalent for "Less Idiosyncratic" Items
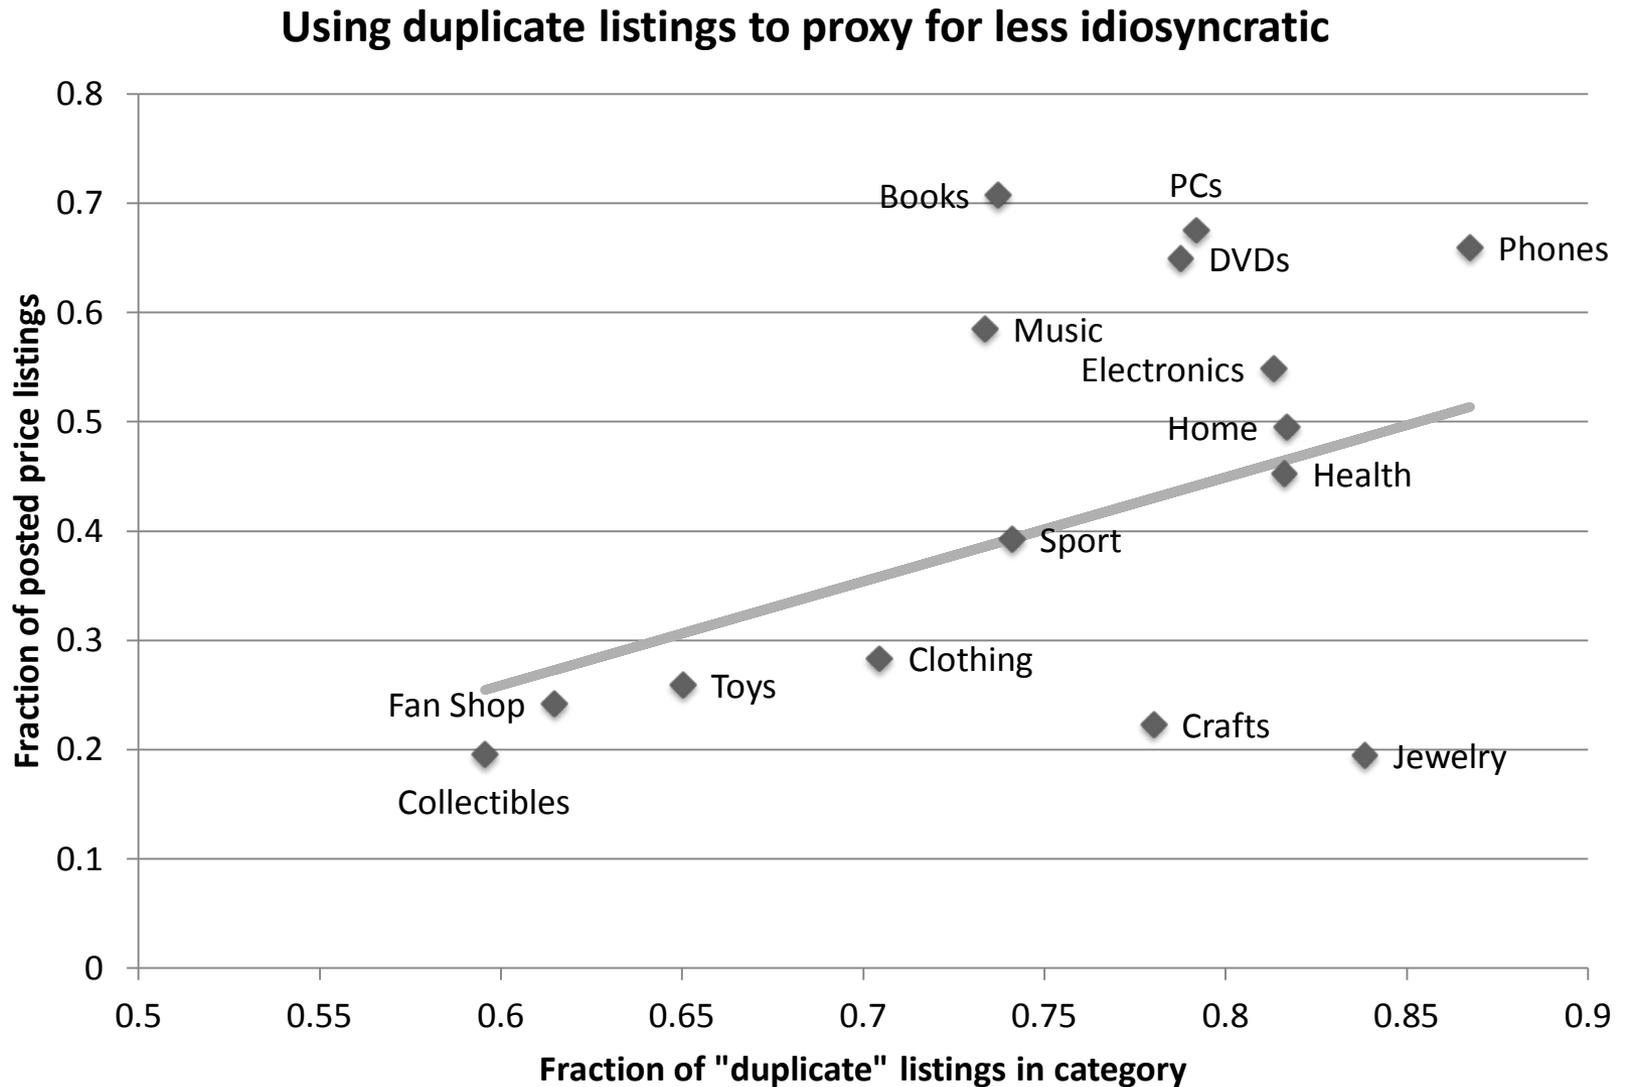
Using duplicate listings to proxy for less idiosyncratic

Figure presents the share of posted price listings in each of the largest 15 (out of 33) item categories, plotted against the share of category listings that are part of a "duplicate" listings (a listing is part of a duplicate if there is at least one other listing of the same seller with an identical category and title). The figure also presents a fitted (weighted) regression line. The figure is based on all eBay listings in 2009 that were either posted price or "pure" auction listings .

Figure 4: Posted Price Use, by Seller Cohort
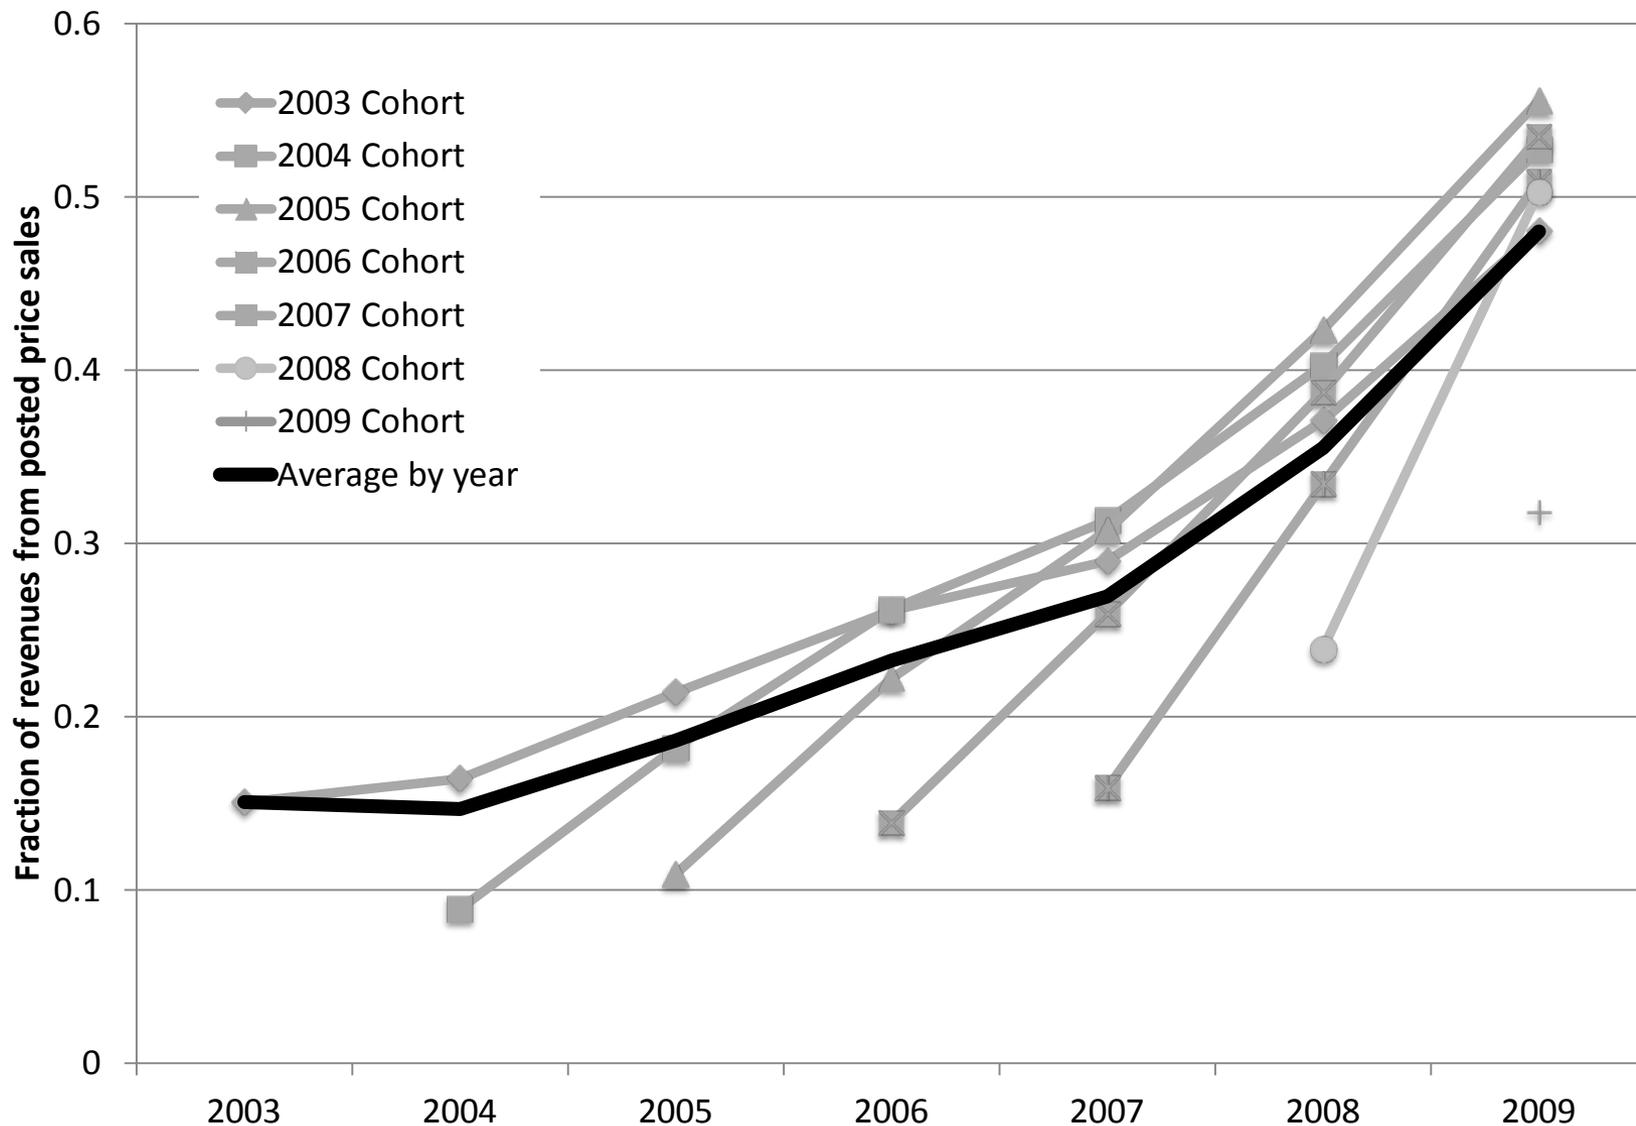
Figure presents the evolution of the annual fraction of eBay sellers' revenue from posted price listings (out of revenues from posted price and "pure" auction listings), separately for each cohort of sellers. A seller is assigned to a cohort based on the calendar year in which the seller's first ever 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
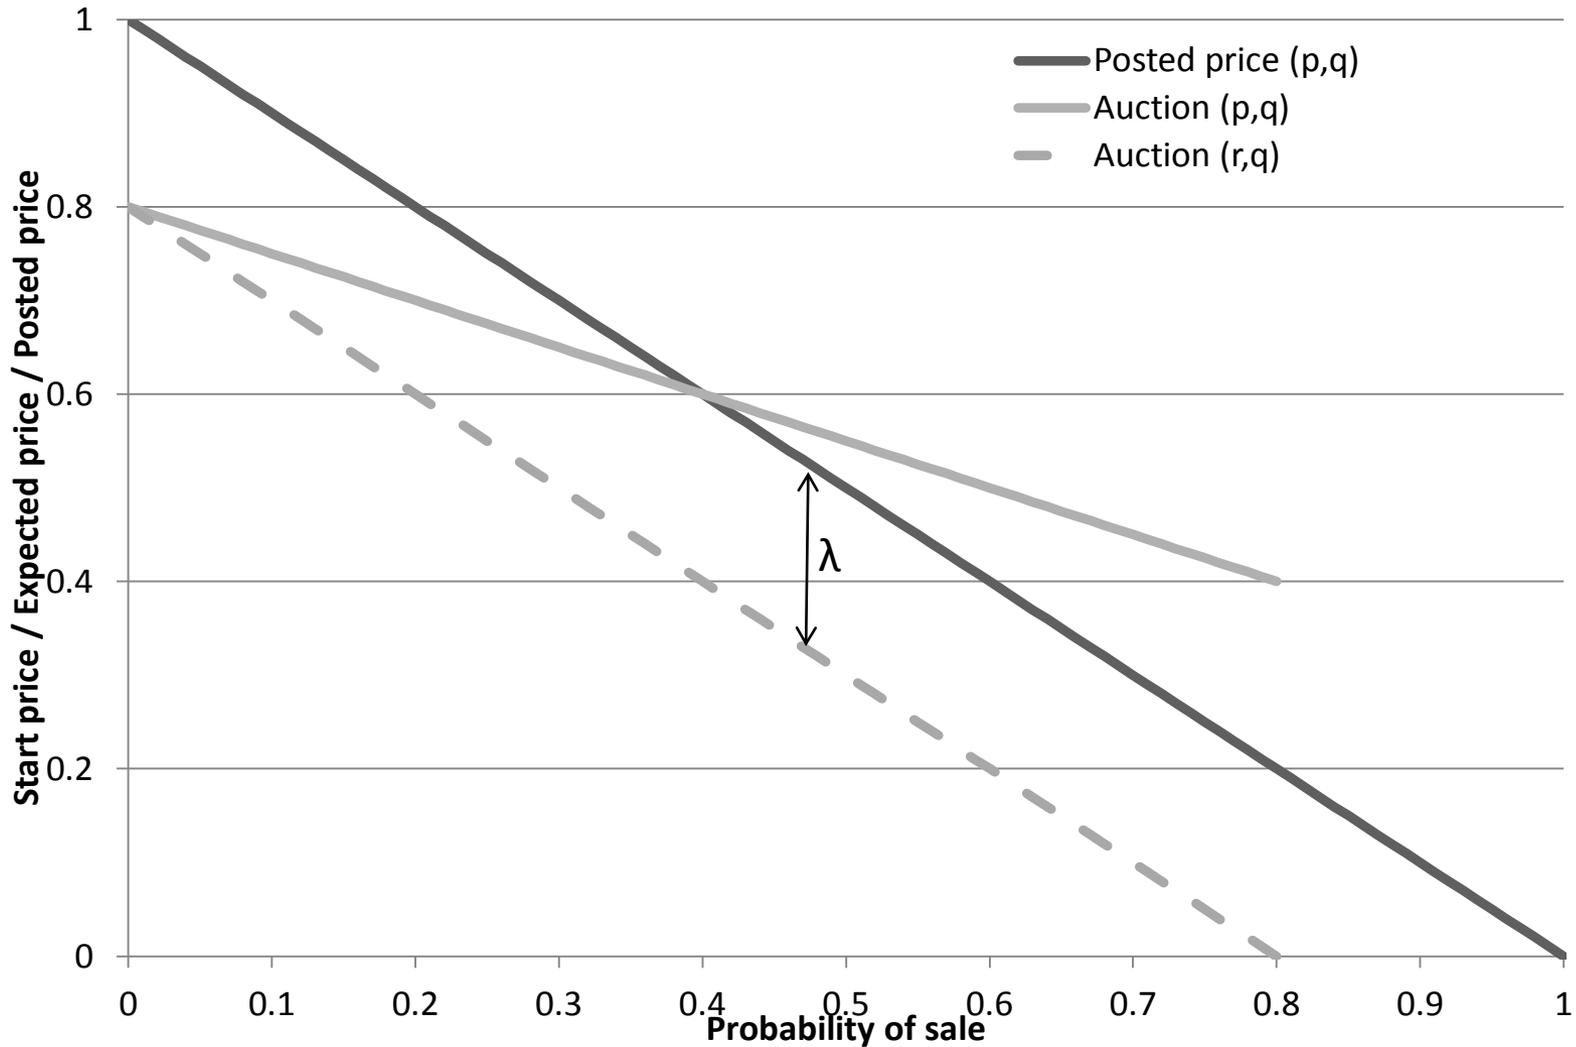
Figure is based on the theory model of Section 3, assuming that values are drawn from a uniform distribution over [0.1], that the reservation utility u=0 and that the disutility from auction λ=0.2. The black solid line represents the "standard" posted-price demand curve. The dashed gray line is the auction sale curve, which graphs the probability of sale as a function of the auction start price; given our assumptions, this line is the a parallel shift down (of size λ) of the posted-price demand curve. Finally, the solid gray line is the "auction demand curve," which graphs the probability of sale that would result from a given start price against the expected sale price that would result from the same start price.

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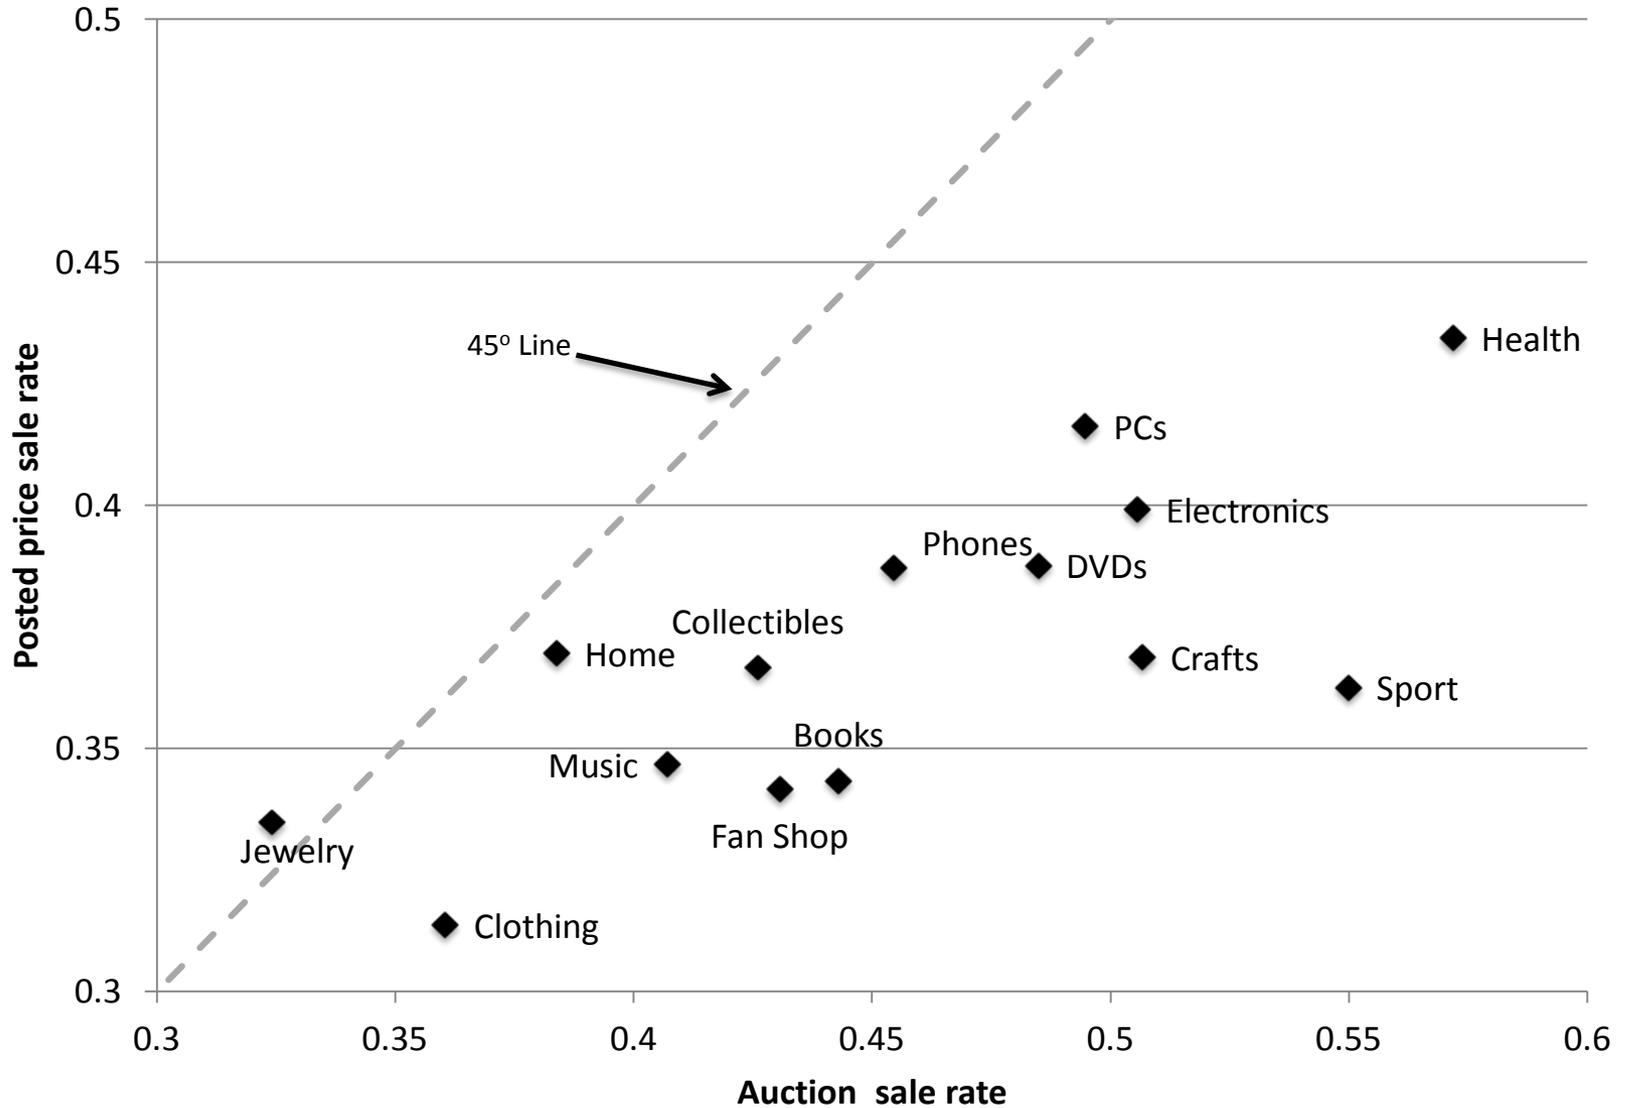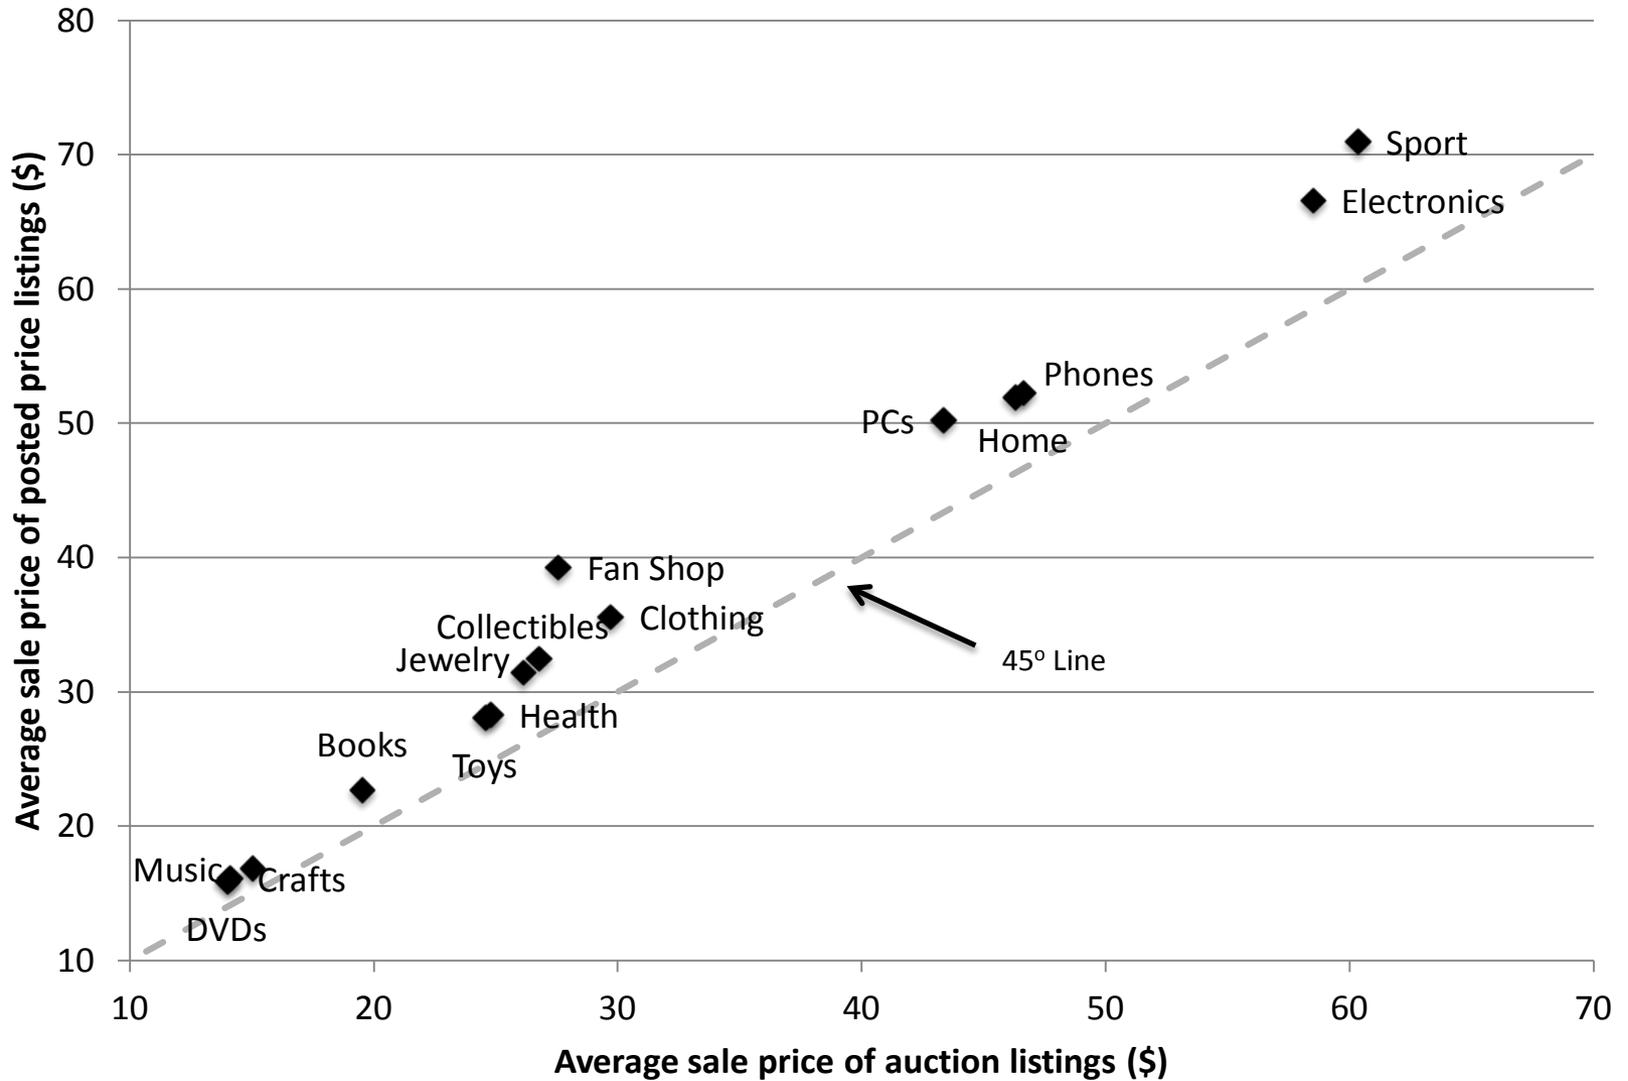

Figure is based on the 2009 seller experiments sample described in Section 4 and reports the average sale rate in each category, by listing format. For each item (experiment) in the sample, we compute the sale rate in each listing format, and then average (unweighted) across all items within each category.

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 the average sale price in each category, by listing format. For each item (experiment) in the sample, we compute the average sale price in each listing format, and then average (unweighted) across all items within each category.

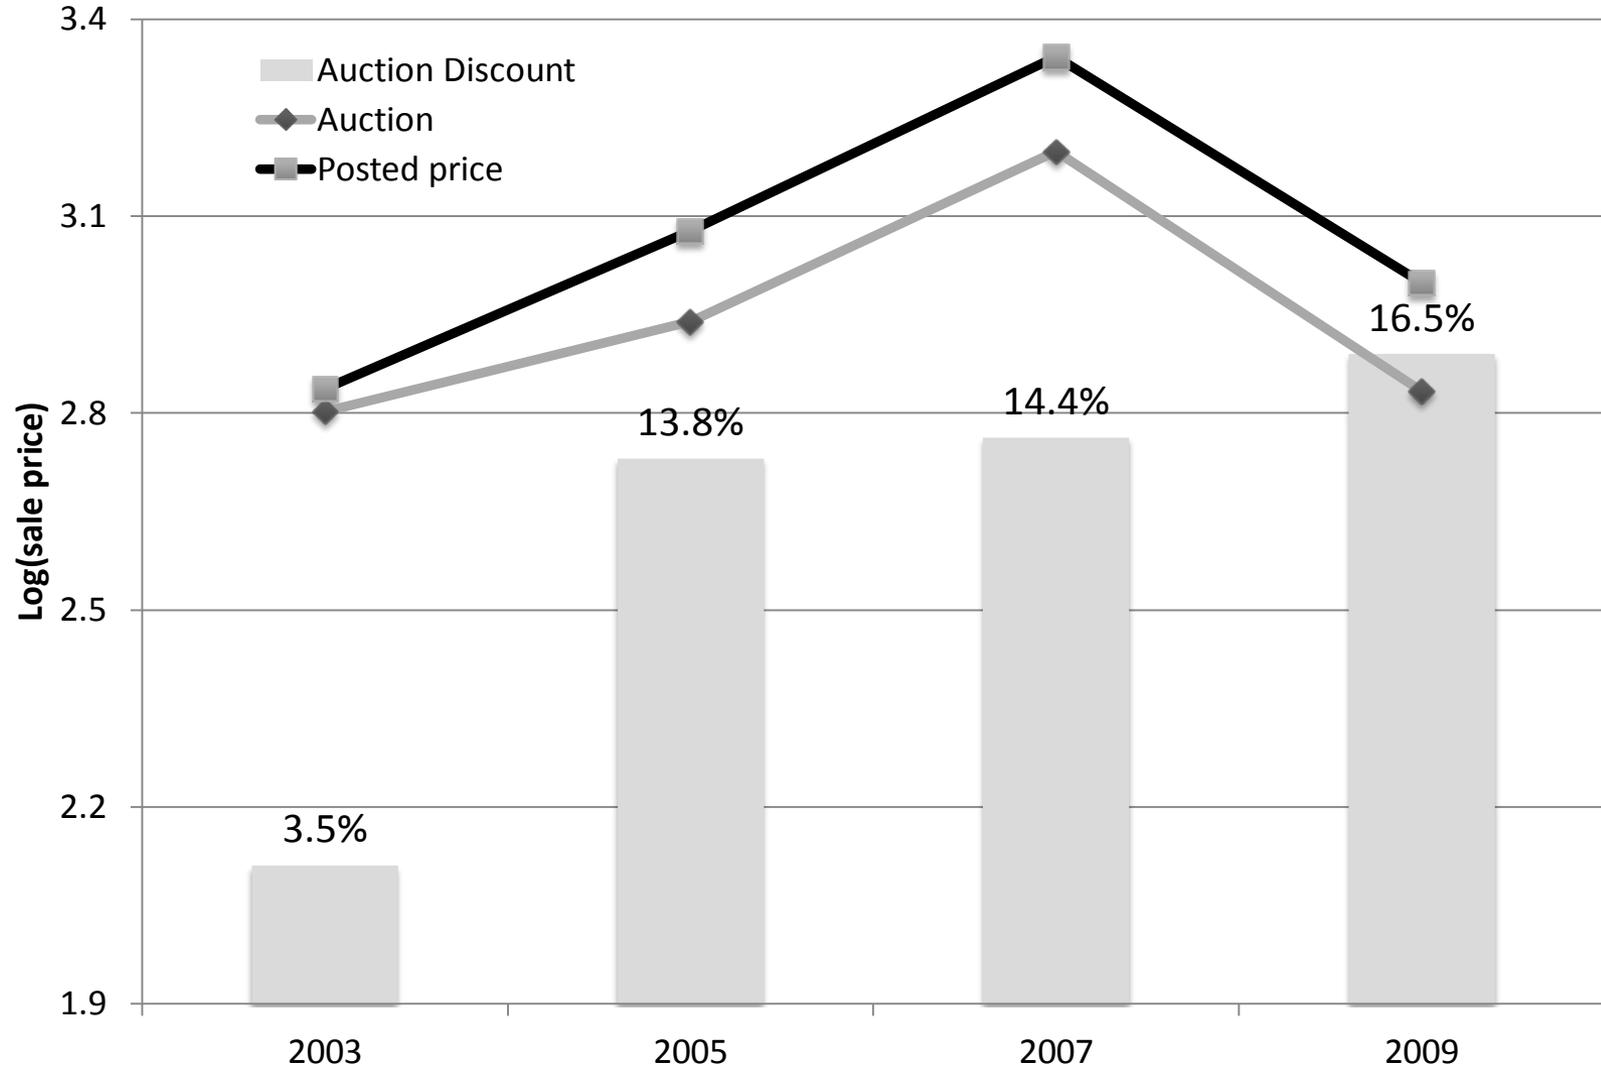# Figure 7(a): Trends in The "Auction Discount" Based on Seller Experiments

Figure is based on the seller experiments sample described in Section 4. 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). The sample is the year-by-year seller experiments sample, as described in Section 4, and these regression only use listing that resulted in successful sales.

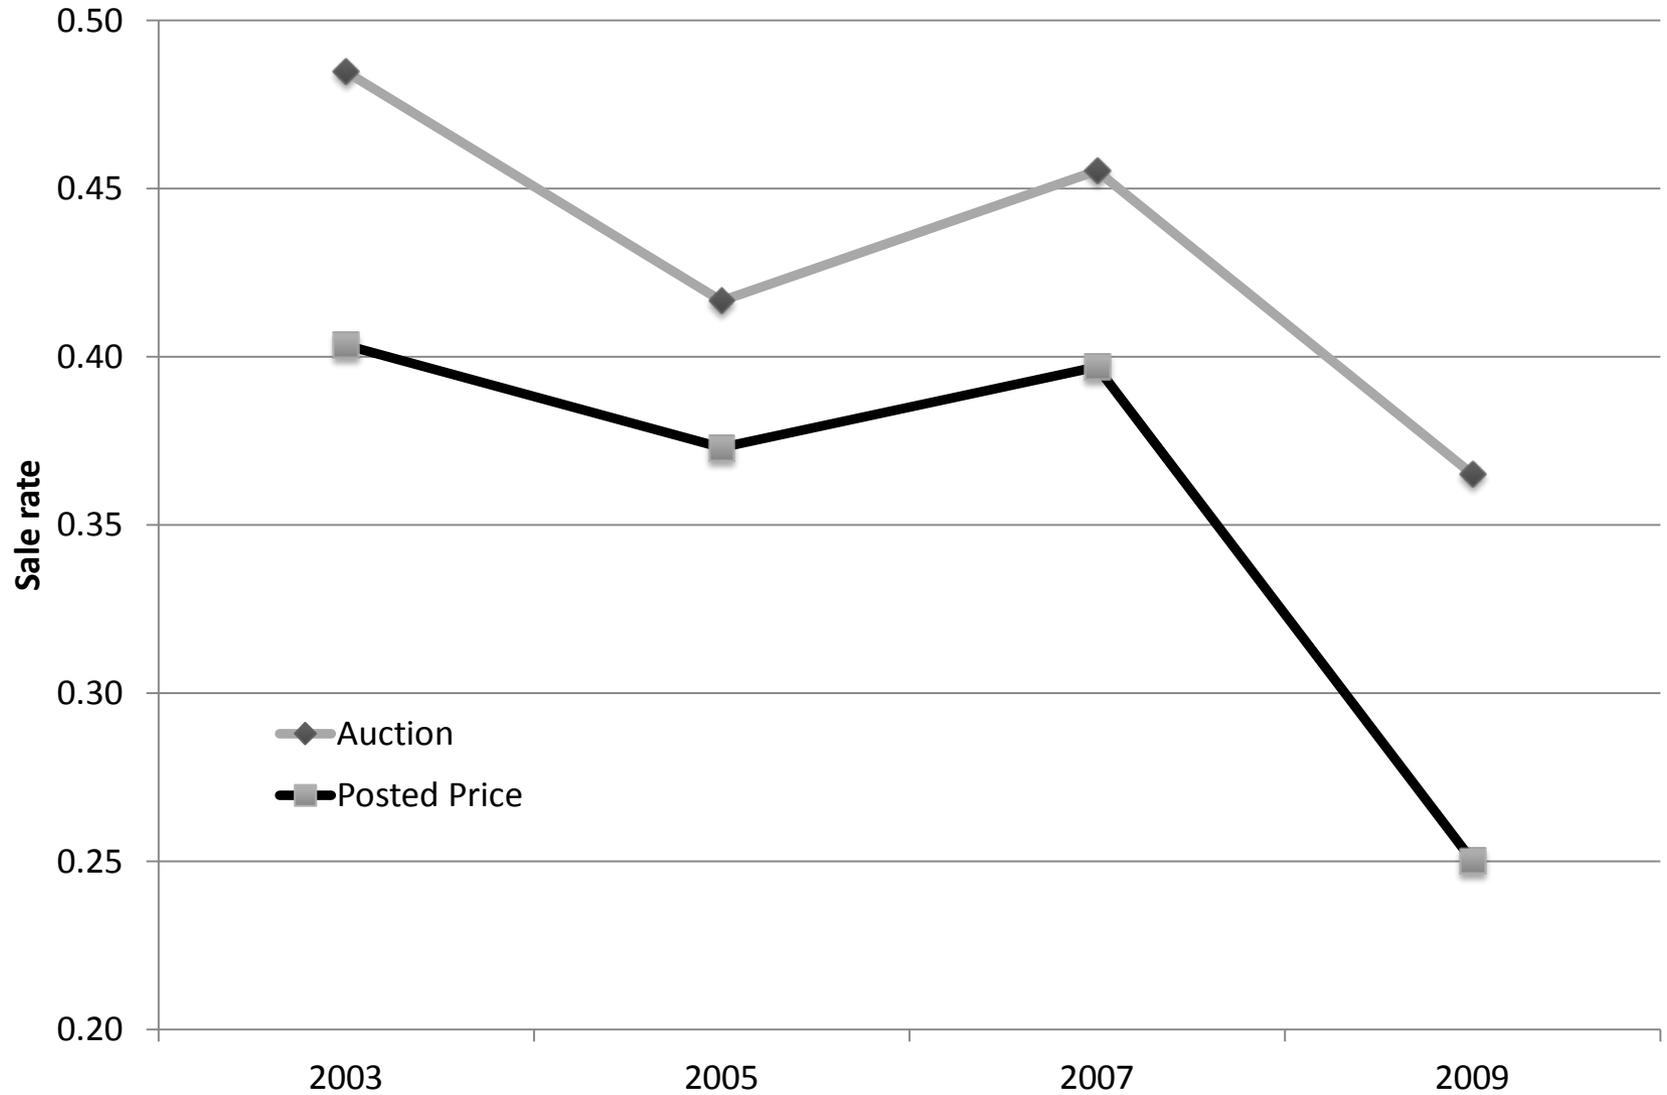# Figure 7(b): Trends in Sale Rate Based on Seller Experiments

Figure is based on the seller experiments sample described in Section 4. 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. The regression is estimated for each year separately. The sample is the year-by-year seller experiments sample, as described in Section 4.

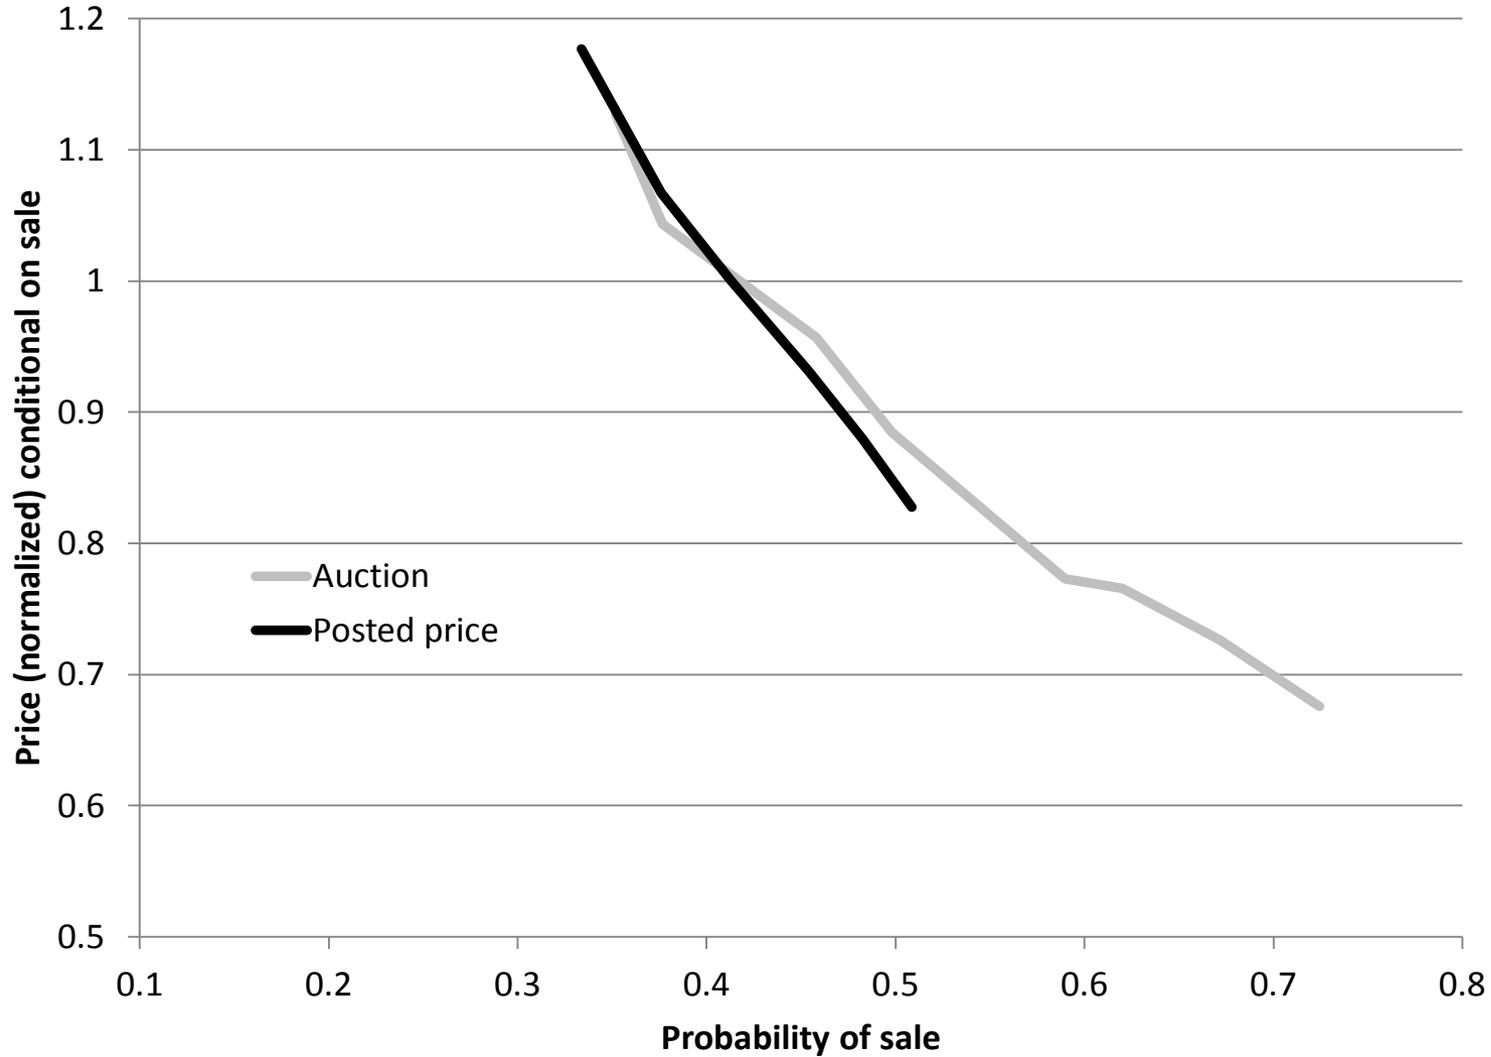# Figure 8(a): The Relative Auction Demand Curve in 2003

Figure is based on the 2003 seller experiments sample described in Section 4. Graphs are based on within experiment and within sale format variation in prices (posted prices or start prices, depending on the sale mechanism). 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 (with reserve price as the key regressor) and a separate price equation (log sale price on reserve price and experiment fixed effects), and combining the estimates to construct the auction demand curve in a way that is parallel to the theoretical auction demand curve constructed in Figure 5.

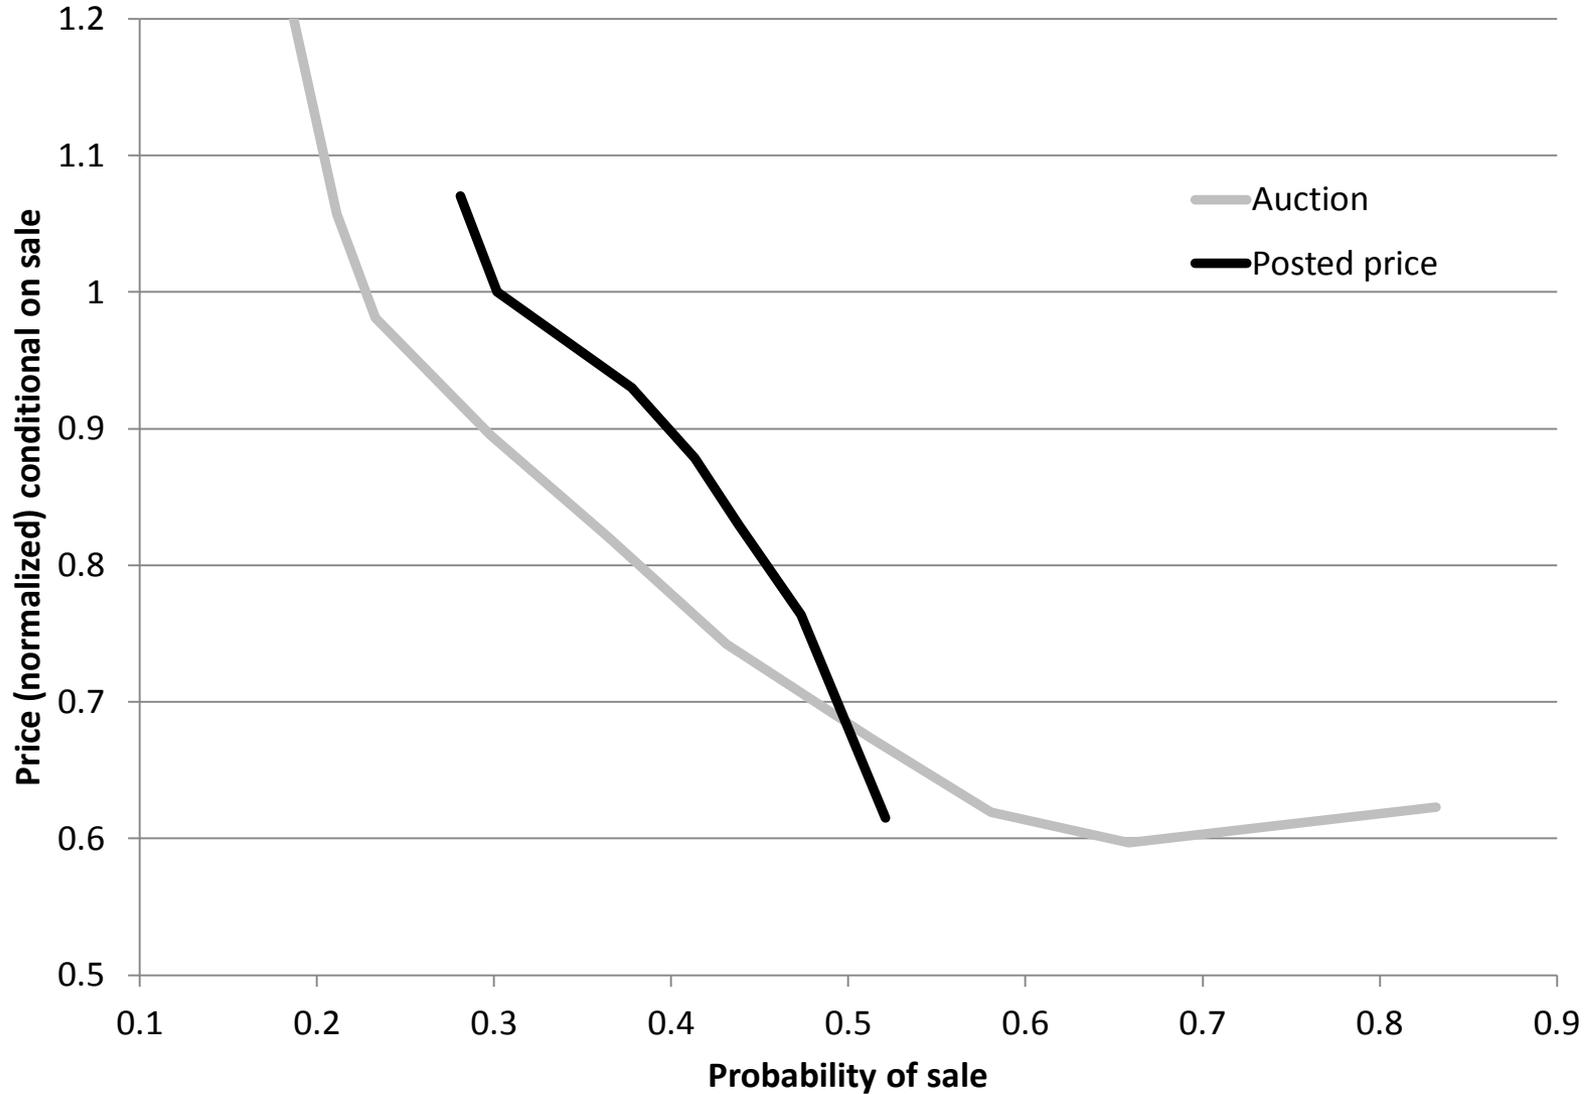# Figure 8(b): The Relative Auction Demand Curve in 2009

Figure is based on the 2009 seller experiments sample described in Section 4. Graphs are based on within experiment and within sale format variation in prices (posted prices or start prices, depending on the sale mechanism). 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 (with reserve price as the key regressor) and a separate price equation (log sale price on reserve price and experiment fixed effects), and combining the estimates to construct the auction demand curve in a way that is parallel to the theoretical auction demand curve constructed in Figure 5.

Figure 9: Calibrating λ and u from Posted-Price and Auction Sale Curves
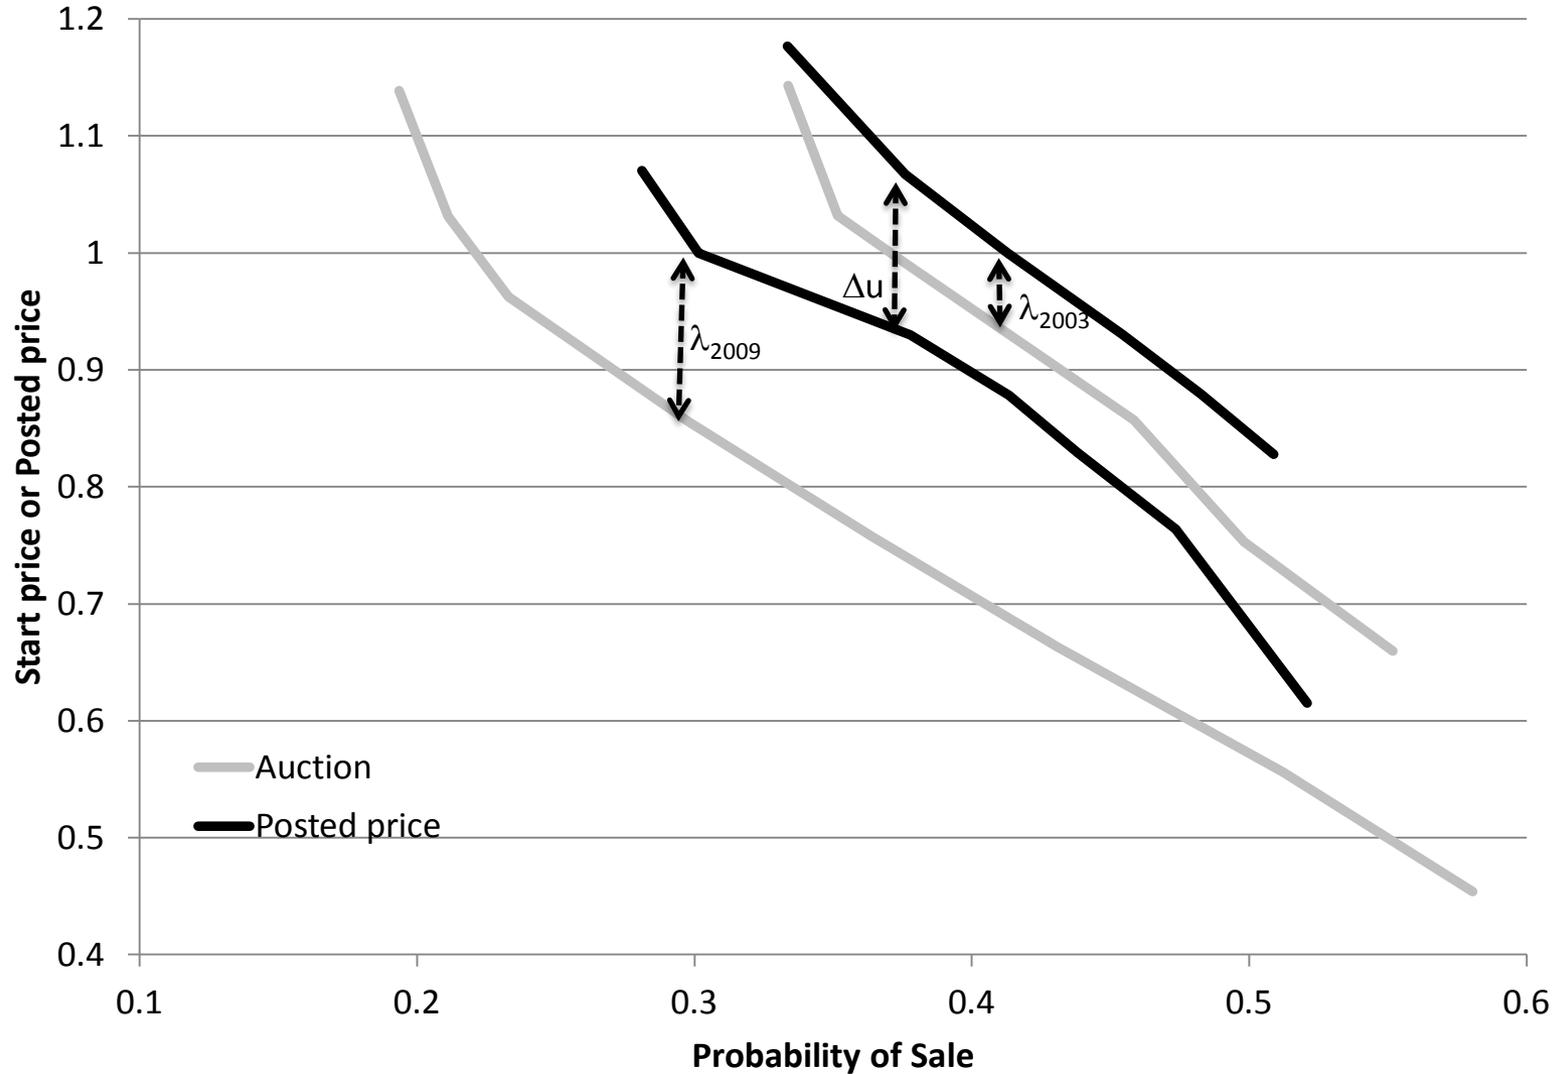

Figure illustrates the way by which we calibrate the values of λ and u, and presents the empirical analog to the theoretical figure (Figure 6). The black lines are the estimated demand curves in 2003 and 2009 (identical to the curves presented in Figures 8(a) and 8(b), respectively), while the gray lines 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). Based on our simple model, 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 10: Posted Price Use, by Buyer Cohort
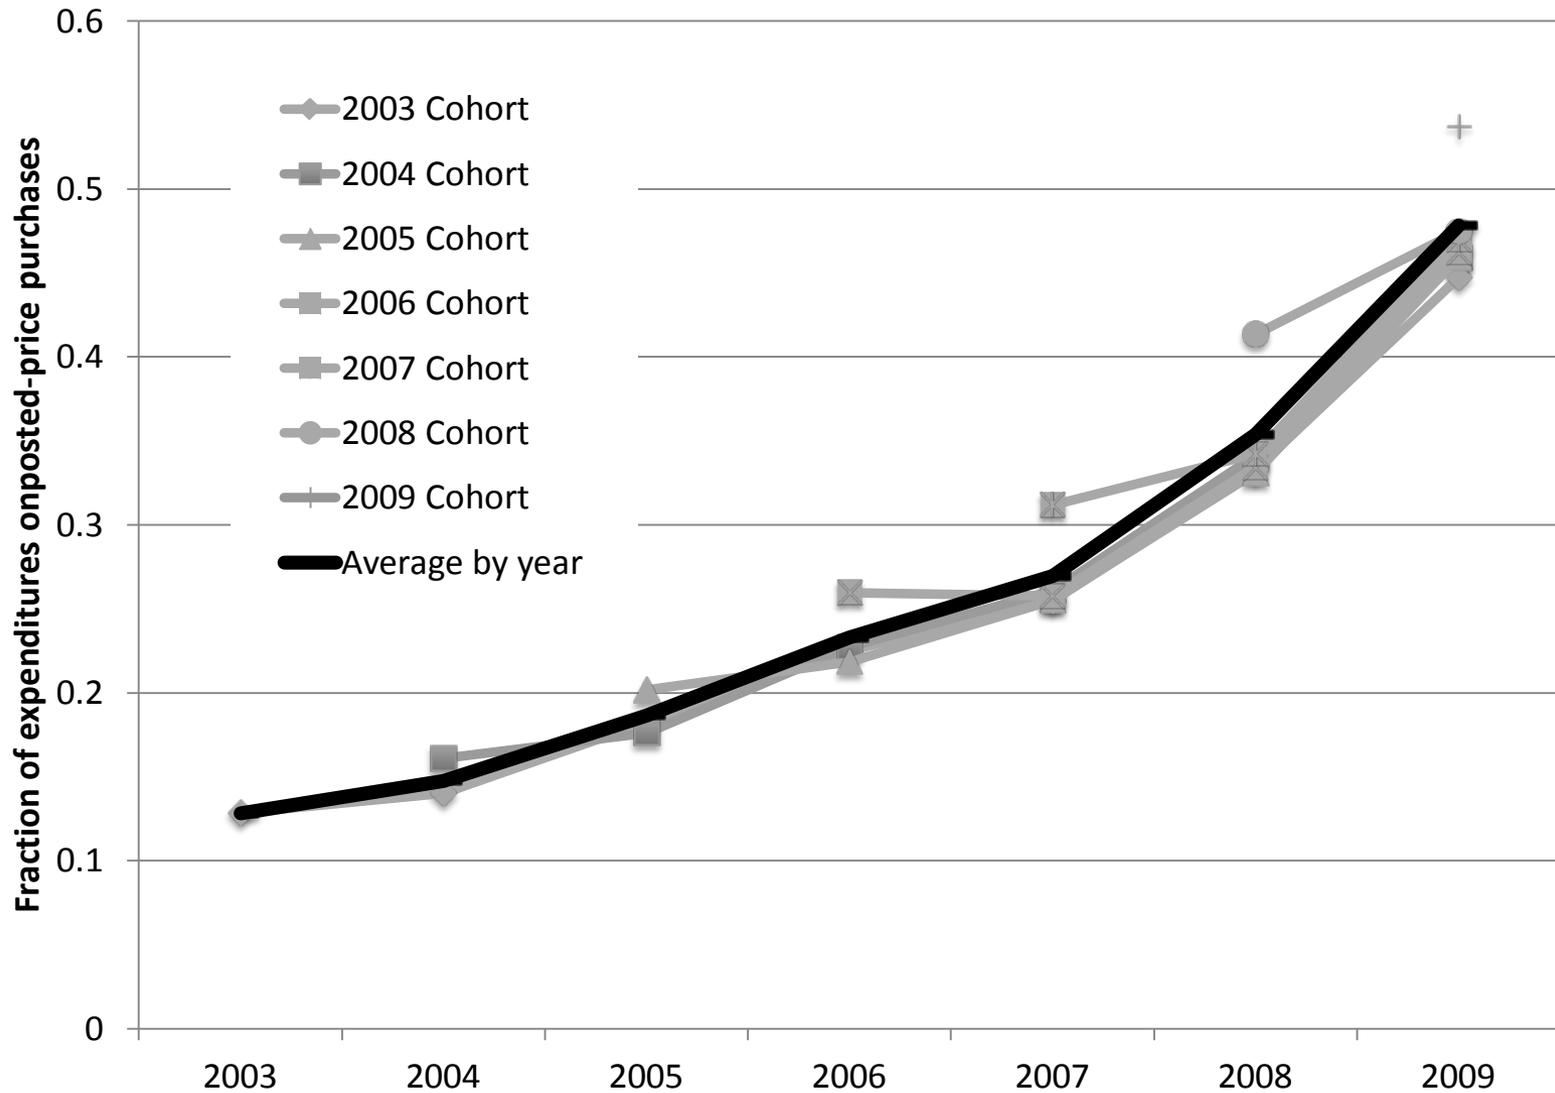

Figure presents the evolution of the annual fraction of buyers' expenditure on posted price purchases (out of expenditures on posted price and "pure" auction purchases), separately for each cohort of buyers. A buyer is assigned to a cohort based on the calendar year in which the buyer's first ever purchase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

Table 1: Summary Statistics of Experiments Samples

	2003 Experiments Data (23,057 experiments and 1,168,033 listings, of which 748,545 (64%) are			2009 Experiments Data (83,685 experiments and 5,924,448 listings, of which 4,185,969 (71%) are		
	Mean	25th pctile	75th pctile	Mean	25th pctile	75th pctile
Number of listings in experiment	50.66	18	50	70.79	19	68
Auction listings in experiment	32.46	8	29	50.02	9	45
Psoted-price listings in experiment	18.19	6	19	20.77	6	18
Share of auction listings	0.560	0.378	0.739	0.615	0.433	0.814
Experiment average sale rate for auctions (within 7 days)	0.507	0.222	0.800	0.442	0.165	0.727
Experiment average sale rate for posted prices (within 7 days)	0.438	0.200	0.667	0.365	0.167	0.500
Difference in sale rate (auction minus posted price)	0.069	-0.100	0.229	0.077	-0.119	0.255
Experiment average sale price for auctions	0.959	0.833	1.042	0.858	0.746	0.979
Experiment average sale price for posted prices	0.968	0.969	1.000	0.983	0.972	1.004
Difference in sale price (auction minus posted price)	-0.009	-0.148	0.079	-0.126	-0.237	-0.003
Experiment average sale price (reference value)	40.3	8.2	32.0	42.7	10.0	40.0
Sandard deviation of (normalized) auction start prices	0.249	0.135	0.335	0.190	0.090	0.271
Standard deviation of (normalized) posted prices	0.102	0.001	0.151	0.078	0.018	0.107

Table presents summary statistics for the earliest (2003) and latest experiments sample. The unit of observation for each entry of the table is an experiment, that is a grouped set of listings that had identical seller, category, and title, and were listed in the same calendar year.

Table 2: Calibration Results

λ	u	Cost	Optimal Auction Profits	Optimal Posted-Price profits	ΔΠ
λ_{2003}	u_{2003}	Cost = 0.3	0.459	0.287	-0.172
λ_{2003}	u_{2009}	Cost = 0.3	0.338	0.228	-0.111 (44%)
λ_{2009}	u_{2003}	Cost = 0.3	0.382	0.287	-0.095 (56%)
λ_{2009}	u_{2009}	Cost = 0.3	0.262	0.228	-0.034 (100%)
λ_{2003}	u_{2003}	Cost = 0.5	0.283	0.203	-0.079
λ_{2003}	u_{2009}	Cost = 0.5	0.162	0.156	-0.007 (74%)
λ_{2009}	u_{2003}	Cost = 0.5	0.206	0.203	-0.003 (78%)
λ_{2009}	u_{2009}	Cost = 0.5	0.137	0.156	0.018 (100%)
λ_{2003}	u_{2003}	Cost = 0.9	0.109	0.111	0.002
λ_{2003}	u_{2009}	Cost = 0.9	0.087	0.089	0.002 (0%)
λ_{2009}	u_{2003}	Cost = 0.9	0.095	0.111	0.016 (100%)
λ_{2009}	u_{2009}	Cost = 0.9	0.074	0.089	0.016 (100%)

Table shows results based on the calibration exercise from Figure 9. Each row presents the optimal profits that would result from an auction listing and posted-price listings, if the underlying parameters were the calibrated ones from 2003, from 2009, or from a combination of the two. We present these results for three representative item cost: of 0.3, 0.5, and 0.9 (normalized by the item's reference value). The right-most column presents the implied difference in profits between auction and posted-price listing, and the fraction for which the change in each of the two parameters accounts for.