Learning from Seller Experiments in Online Markets*

Liran Einav, Theresa Kuchler, Jonathan Levin and Neel Sundaresan†

September 2012

Abstract. The internet has dramatically reduced the cost of varying prices, displays and information provided to consumers, facilitating both active and passive experimentation. We document the prevalence of targeted pricing and auction design variation on eBay, and identify hundreds of thousands of experiments conducted by sellers across a wide array of retail products. We use the data to measure the dispersion in auction prices for identical goods sold by the same seller, to estimate nonparametric auction demand curves, to analyze the effect of "buy it now" options and other auction design parameters, and to assess consumer sensitivity to shipping fees. We also investigate the robustness of the results by isolating different types of identifying variation, as well as the heterogeneity of the estimates across item categories. We argue that leveraging the experiments of market participants takes advantage of the scale and heterogeneity of online markets and can be a powerful approach for testing and measurement.

*We thank John Asker, Susan Athey, Eric Budish, Preston McAfee, Hal Varian and Glen Weyl for helpful comments. We acknowledge support from the National Science Foundation, the Sloan Foundation, the Stanford Institute for Economic Policy Research, and the Toulouse Network on Information Technology. The data for this study were obtained under a consulting agreement between the Stanford authors (Einav, Kuchler, Levin) and eBay Research.

†Department of Economics, Stanford University (Einav, Kuchler, Levin), NBER (Einav and Levin), and eBay Research Labs (Sundaresan). Email: leinav@stanford.edu, tkuchler@stanford.edu, jdlevin@stanford.edu, and nsundaresan@ebay.com.

1 Introduction

The internet has dramatically reduced the cost of changing prices, displays and information provided to consumers, and of measuring the response to these types of changes. As a result internet platforms, retailers and advertisers increasingly can customize and vary their offers. One effect of this flexibility is to facilitate learning. Google, for instance, conducts thousands of experiments each year to refine its search platform (Varian, 2010), and Microsoft constantly experiments with its advertising platform (Athey, 2011). Our goal in this paper is to describe and illustrate another benefit: what amounts to large-scale experimentation by market participants can be used to address traditional economic questions about consumer behavior and market outcomes.

Our analysis focuses on eBay, the largest e-commerce platform and a primary sales channel for tens of thousands of retailers. We define a "seller experiment" on eBay to be a case where a given seller lists a given item multiple times while varying pricing or auction parameters. This practice — analogues of which can be observed in other internet markets, such as for sponsored search or display advertising — is extremely common. Of the hundred million listings appearing on eBay on a given day, it is possible for more than half to find a near-duplicate listing of the same item by the same seller, with modified sale parameters. Drawing on a single year of listings, we assemble a dataset consisting of hundreds of thousands of seller experiments conducted across thousands of diverse sub-markets.

We apply this data to analyze consumer behavior and the effects of alternative pricing strategies. Our empirical strategy is straightforward. For each application, we identify in the data a large number of experiments where the seller has varied the relevant pricing parameter. We pool these experiments, many of which are modest in size, and use fixed effect regressions to identify average effects. The scale of the data is sufficiently large, and the variation sufficiently pervasive, that we can consider how effects vary across product categories and price points, and use different data construction strategies (e.g. matching only contemporaneous listings, or sequential listings) to examine whether our estimates might suffer from endogeneity or sample selection biases. We take this approach to four main analyses.

First, we estimate the variability or dispersion in auction prices, holding fixed both the product and the seller. In an environment where physical search costs are extremely low, one might expect auction prices for a given item sold by a given seller not to vary much and, if the seller also offers the item at a posted price, to be capped above by the posted price. Instead, we find that auction prices vary substantially. The average coefficient of variation is 10-15%, when we compare equivalent auctions in the same calendar month. At the same time, we find that auction prices generally do not rise above equivalent (i.e. same seller, same item) posted prices, an event that was more common a decade ago (Malmendier and Lee, 2011). We reconcile these findings by showing that on average auction prices are well below equivalent posted transaction prices.

Second, we estimate auction demand using variation in auction start prices. As an auction seller raises her start (or reserve) price, she lowers the probability of sale but raises the expected final price conditional on selling. Variation in the start price therefore traces out a familiar demand curve in price-quantity space. Our nonparametric demand curve estimates have a rather unexpected feature: they are highly convex, so their associated marginal revenue curves are not downward sloping. An implication is that very low and very high start prices should be preferred to intermediate ones. Consistent with this, we show that the observed distribution of start prices is bimodal. We also use the same start price variation to examine a behavioral hypothesis of Ku et al. (2006) and Simonsohn and Ariely (2008) that low start prices can create bidding "escalation" that ultimately leads to higher final prices. We find some patterns in the data that are consistent with this effect, but the evidence is weak.

Third, we analyze the effect of "buy-it-now" options in consumer auctions. A buy-it-now option allows a buyer to preempt the auction by purchasing the item at a posted price set by the seller. In principle, this can allow a seller to discriminate between impatient but possibly high value buyers, and bargain hunters who are willing to wait for the auction. We find that the effect of offering a buy-it-now option depends a great deal on how the buy price is set. At the typical level used by sellers, the effect on revenue is negligible. Consistent with the price discrimination theory, however, sellers do generate additional revenue by setting a relatively high buy price. We also evaluate the behavioral hypothesis that buy prices might act as a reference point in subsequent bidding.

Fourth, we revisit a finding of Tyan (2005), Hossain and Morgan (2006) and Brown, Hossain and Morgan (2010) that consumers appear to underweight shipping fees relative to regular prices. This application illustrates how our empirical strategy allows us to exploit the scale of internet data. We expand from the five specific items studied by Tyan, and the 20 CD and Xbox titles, and two specific iPod models in the latter papers, to analyze targeted shipping fee variation for over six thousand distinct items. In this large sample, we estimate that moving from a small shipping charge to "free shipping" increases the expected auction price by more than \$2. We also confirm the earlier finding that once fees are positive, consumers do not fully internalize increases. We estimate that every \$1 increase in the shipping fee reduces the auction sale price by only around \$0.82.

The empirical strategy we pursue in this paper, while very straightforward, differs from almost all prior studies of eBay and other internet markets. Prior work has relied mainly on two approaches. The first is to use observational data, selecting a small number of narrowly defined products and attempting to control for quality variation across sellers and listings by using observed covariates (e.g., Bajari and Hortacsu, 2003). The second is to use field experiments in which a researcher sells a small number of identical items while varying one or a few sale parameters (e.g., Lucking-Reiley, 1999). Either way, the analysis typically is limited to a handful of items and tens or hundreds of sales. In contrast, we aggregate evidence from thousands of items and tens or hundred of thousands of sales.

The approach we take is one possible response to a tension in analyzing large-scale internet data. The tension is between, on the one hand, leveraging the vast scale of the data, and on the other hand, obtaining plausible identification of economic effects. We elaborate on this trade-off in Section 4. We demonstrate how biases can arise in large sample estimates that pool listings across heterogeneous sellers and items, even within narrow product categories. At the same time, we show that effects can vary greatly across products, limiting the conclusions that can be drawn from small-scale experiments. Another advantage of our approach relative to researcher-conducted field experiments is that it can be applied retrospectively to study how consumer behavior or pricing incentives have changed over time.

The main concern with relying on seller-induced variation in pricing and auction design is the potential for endogeneity or sample selection biases. We discuss below why certain features of eBay's platform make it both easy and desirable for sellers to experiment or vary their listings in ways that replicate conscious experimentation. Still, sellers may vary sale parameters for "non-experimental" reasons, and one significant concern is that our results are driven by sellers changing start prices or shipping fees or auction design choices in response to changes in consumer demand, or only after an initial strategy has failed. To evaluate these potential biases, we also report estimates using a variety of more stringent criteria to match listings, so that for instance we rely only on variation in pricing across contemporaneous listings of the same item by the same seller. The estimates are very similar using alternative definitions of an experiment.

A few prior studies have taken approaches related to ours. Ostrovsky and Schwarz (2009) study a platform-wide field experiment in which reserve prices in Yahoo!'s search advertising market were changed for thousands of individual keywords. Elfenbein, Fisman and McManus (2012) study the effect of charity contributions by eBay sellers, using a matched listings approach that is essentially the same as the one we employ here. They do not remark on either the prevalence of duplicate listings or the opportunity for using them as a broader research tool. Finally, in Einav, Farronato, Levin and Sundaresan (2012) we use the approach developed here, but apply it to data from multiple years, to explain why sellers on eBay have moved over time from selling by auction toward posted prices.

The remainder of the paper proceeds as follows. Section 2 describes the use of duplicate listings and "experiments" by retail sellers on eBay, our data construction, and summary statistics. Section 3 uses the experiments data to analyze the problems described above: price variability, auction demand, buy-it-now prices, and shipping fees. Section 4 compares the use of seller experiments with using more heterogeneous observational data, and also shows why results from a limited set of products may not be representative. In Section 5 we conclude by discussing why sellers vary their pricing parameters so often and so widely. A lengthy appendix provides many additional analyses that address various potential endogeneity and selection biases. We replicate all the results using a range of samples and specific approaches to matching listings, showing that the results are highly consistent across these alternatives.

2 Background, Data, and Empirical Strategy

2.1 Background and Empirical Challenge

The e-commerce platform eBay had approximately ninety million active users and $57 billion in gross merchandise volume in 2009, the year of our data. The site includes large and active sub-markets for collectibles, electronics, clothes, tickets, toys, books, jewelry and art, both new and used. Products are offered by thousands of professional retailers, and millions of individual users. The platform's scale, and the ease of collecting data and running experiments, has made it a focal point for research on online markets.[1]

Sellers on eBay have considerable flexibility in designing a sales strategy. Sellers select a listing title and picture of their product, a longer item description, a shipping fee, and a sales mechanism. Traditionally, most sellers have used ascending auctions. This means specifying an auction duration, a start price, and perhaps an additional secret reserve price, or a buy-it-now price at which a bidder can purchase the item before an initial bid is made. Sellers also can use regular posted prices. Nowadays, posted price transactions account for more than half of eBay's sales volume. It is easy for sellers to change these sale parameters from listing to listing.

The diversity of selling strategies creates an opportunity to learn about how consumers respond to different pricing and sales mechanisms, to gain insights into consumer search, the efficiency and competitiveness of the market, and to test hypotheses about consumer behavior. At the same time the diversity of the sellers and products, a common feature of online platforms, poses a challenge for researchers. We illustrate this point and how it motivates our empirical strategy in Figure 1.

Figure 1(a) shows the eBay listings displayed following a search for "taylormade driver" (a type of golf club) on September 12, 2010.[2] The market for even this narrowly defined product is large (over 2500 listings) and heterogenous. The products are differentiated (different models and sizes, new and used), as are the sellers (by location, reputation score, whether

[1] Bajari and Hortacsu (2004) and Hasker and Sickles (2010) review dozens of papers using data from eBay.

[2] Consumers shopping on eBay find items either by typing in search terms or browsing through different categories of products. Products are displayed as listings similar to Figure 1(a), and can be sorted in various ways. The default sort is based on a relevance algorithm. Consumers then click on individual listings to see more detailed item information, place bids, or make purchases.

they are "top-rated"), and the sales mechanisms (posted prices, auctions, buy-it-now auctions, all with different end times), and the shipping arrangements and fees. Because listings vary along many dimensions, it is challenging to attribute consumer responses to specific sales strategies, despite observing thousands of contemporaneous listings in a narrow product category. This problem has motivated the use of field experiments in which researchers post a small number of listings, say fifty or a hundred, which vary on only one or two pricing dimensions.

Ideally one would like an empirical strategy that preserves the type of variation in the field experiment approach, but that can be scaled to study the larger marketplace. The key idea of this paper is to make (and subsequently exploit) a simple observation, which is that sellers frequently "experiment" with a given listing by varying their pricing or choice of sales mechanism. Figure 1(b) provides an example. It shows a subset of thirty-one listings located by the search query above. They are all for the same item, and have been listed by the same retailer (with the user name *budgetgolfer*). However, they are not completely identical. Eleven of them offer the driver for a fixed price of \$124.99, while the other 20 are auctions scheduled to end within the next week. Also, the listings have different shipping fees, either \$7.99 or \$9.99. So this group of listings can play the role of a small field experiment to identify the dispersion in auction prices, and their relationship to posted transaction prices, or to assess whether auction prices fully adjust to account for shipping fees.

As we describe below, the behavior of posting near-identical listings with varying prices, fees and sales mechanisms — either contemporaneously or over time — is extremely common. We discuss below several reasons for this, but one factor is simply mechanical. Auctions on eBay are for a single unit, so a retailer who wants to sell multiple units must post multiple listings. Once a retailer is making multiple listings, there is little cost and some informational benefit to trying different approaches, even concurrently given that eBay's search algorithm will typically spread the listings across multiple pages of results rather than in head-to-head competition.[3] The next sections describe how we search eBay's data to identify "seller

[3] The advice to experiment with different strategies is common on websites and discussion boards that cater to eBay sellers. For instance, in a post picked somewhat at random from the reviews.ebay.com site, the user *cjackc* advises that sellers review historical data on the best day to end an auction, "... and then experiment with your own unique listing to see if you can find even more success...." because "... your items are unique and what works for others may not work best for you." (http://reviews.ebay.com/Is-Sunday-

experiments" and our approach to aggregating them.

2.2 Experiments data

We construct our data from the universe of eBay.com listings in 2009. We exclude only auto and real estate listings, which have a somewhat different institutional structure. We look for matched sets of listings that involve the same seller offering the same item. Because most eBay listings do not include a well-defined product code, we use the listing title and subtitle to identify products.

Specifically we identify all sets of listings that have an exact match on four variables: seller identification number, item category, item title and subtitle. We then drop single listings that have no match. This leaves around 350 million listings, grouped into 55 million matched sets. We refer to each set of matched listings as a *seller experiment.* As an example, the listings in Figure 1(b), together with any additional matched listings that were active before or after the day of the screenshot, comprise one experiment.[4]

Our empirical strategy relies on variation within experiments in sale parameters and outcomes. In this paper, we focus primarily on auction listings and outcomes, which leads us to refine the data in several ways. In particular, we restrict attention to experiments that include at least two auction listings and at least one successful posted price listing. The former is necessary to have within-experiment auction comparisons. The latter, as we explain below, provides a useful way to normalize prices in order to make experiments comparable and compute average treatment effects. Finally, we include only those experiments where the listings have a non-empty subtitle. This is a convenient way to reduce the size of the data to make it manageable, while focusing on more professional retailers who tend to use subtitles. In the Appendix, we also report all our results for a random 10% subsample of the experiments that meet our initial criteria.

This generates our baseline dataset: 244,119 experiments with a total of 7,691,273 list-

Really-the-Best-Listing-Day?ugid=10000000008235490)

[4]Note that by using title and subtitle to identify items, we exclude cases in which a seller might have offered the same item with varied listing titles. On the other hand, it is also possible that we might include certain cases in which a seller offered different items under the same title or used different photos for the same item, although we manually checked a random sample of the data and did not find any examples of this, so we suspect that such instances are not common.

ings. The data include cases in which a seller posts multiple overlapping auctions and in which a seller runs multiple non-overlapping auctions, as well as combinations thereof. Table 1 presents summary statistics, along with corresponding statistics for the entire "seller experiments" data and for a large random sample of eBay auction listings. In the baseline data, just over a third of the listings result in a sale, with an average price around $67.

By construction, the items in our sample are less unique and idiosyncratic than many items sold on eBay, and the sellers relatively professional. This is reflected in Table 1 in the fraction of items that are "catalogued," the experience of the sellers, and their tendency to use "sophisticated" sale strategies such as a "But-It-Now" (BIN) option. It also shows up in the distribution of items across product categories. Relative to the rest of eBay, our sample includes more cell phones, video games and electronics, and less clothing, jewelry and collectibles. Essentially we are looking at professional and semi-professional retailers, while eBay as a whole also includes a vibrant consumer-to-consumer market.

Table 2 provides summary statistics at the experiment level. The average experiment in our baseline data has 32 auction listings. About 70 percent of the experiments have at least one sale. Figure 2 shows the distribution of experiment sizes. Roughly 45 percent of the experiments have four or fewer listings, but there are also many (much) larger experiments. The typical experiment includes multiple listings that occur over a relatively short time period, just under two months on average.

Our empirical strategy relies on the fact that when sellers post multiple listings for the same item, they regularly vary different sale parameters. The amount of variation in the data is large. Table 3 reports the number of experiments that contain variation in different sale parameters of interest. The first column shows that of the 244,119 experiments in the baseline sample, more than 140,000 have variation in the auction starting price, more than 17,000 have variation in the shipping fee, more than 90,000 have variation in the BIN option, and more than 92,000 have variation in the auction duration. The remaining columns show that we can find large numbers of experiments with variation in a given sale parameter even if we condition on other sale parameters being held fixed. We rely on this below to construct samples of experiments in which we seek to pinpoint specific pricing effects.

2.3 Empirical Strategy

Our data includes a large number of experiments involving distinct items and sellers. We aggregate the information in these experiments using fixed effects regressions. Let i index experiments, t index listings within experiments, and z_{it} denote a listing parameter whose effect we want to know. For a given outcome of interest y_{it}, we estimate regressions of the form:

$$y_{it} = \alpha_i + f\left(z_{it}\right) + \varepsilon_{it}, \tag{1}$$

where α_i is an experiment fixed effect and ε_{it} is an error term assumed to be mean-independent of z_{it} within experiments.

There are at least two reasons to pool experiments as in our specification. First, many experiments are small, so pooling provides much greater statistical power. Second, it seems easier and more digestible to report an average effect rather than thousands of distinct effects for individual items. That being said, we break out estimates by item value, and discuss heterogeneity across item categories in Section 4.

One challenge in aggregating effects is that the experiments involve items of different value. A \$1 increase in the auction reserve price may be important for a \$5 item but not so important for a \$500 item. To address this, we define a reference value for each item, and evaluate price changes for an item relative to its reference value. Specifically, we define each item's reference value v_i as the average price across posted price transactions of that item.[5] Then when we consider auction sales, we focus on the normalized price $p^n_{it} = p_{it}/v_i$ rather than the auction price p_{it}. Similarly, in studying auction reserve prices, we use the normalized reserve price $s^n_{it} = s_{it}/v_i$ rather than the dollar reserve price s_{it}. A more general alternative would be to estimate treatment effects of the form $f\left(s_{it}, v_i\right)$ rather than $f\left(s_{it}/v_i\right)$ but we find, rather surprisingly that there seems to be little gain from doing this.

We rely on two further assumptions to identify average treatment effects. The first is that the idiosyncratic effects of each experiment denoted by α_i enter in an additive and separable

[5] Recall that in selecting experiments into our baseline sample, we required each experiment to have at least one successful posted price listing. Note that we use posted price transactions and not listings so that the reference value is not affected by excessively high posted prices that never sell. We also experimented with modifications to this definition, for example using the median transaction price or trimming outliers before taking averages, and the results (not reported) remain virtually the same.

way. The second is that sale parameters within each experiment are not correlated with factors that bear directly on auction outcomes. This assumption deserves some discussion, which we turn to next.

2.4 Threats to Identification

In our baseline analysis, we group listings of a given item by a given seller over a period of up to a year, and estimate treatment effects using variation in auction parameters within these matched listings. An obvious concern is that sellers may be changing their sale parameters as a reaction to changing demand. Then our estimates will suffer from a standard endogeneity bias. Our estimates also could be tainted by various forms of selection bias, for instance if sellers change their sale parameters only after an initial strategy has failed.

One way to address these concerns is to vary the definition of an experiment, by using more stringent criteria for matching listings. For instance, we can match listings only if they occur within a short time window, or overlap in time, or even start and end on the same day. As we strengthen the criteria, we reduce any potential variation in demand, as well as any possible variation in the seller's information at the time of posting. Doing this, however, has the drawback that it throws out a great deal of potentially valuable data, such as from sellers who are experimenting with their sale parameters in a persistent and ongoing way.

Relying on short time windows, or contemporaneous listings, also raises a different concern. It is a concern that also could be raised about many field experiments. The concern is that nearly-identical listings may compete with one another. The structure of eBay's search algorithms, which spread duplicate listings over many pages of closely related listings, arguably alleviates this concern. Nevertheless, it suggests that an alternative approach of grouping near-identical listings that *do not* overlap in time also may be valuable.

A virtue of our general approach is that it is straightforward to replicate our estimates using any of these alternative definitions of an experiment. In the Appendix, we report replications of all our results using a variety of different samples and specifications. In particular, we report results where listings are grouped only if they occur in overlapping fashion, only if they occur contemporaneously, only if they occur sequentially, and only if they occur sequentially but within a thirty day period. We also report results for a subsample of large experiments, and for a sample in which all auctions occurred in the presence of a parallel posted price listing.

The results remain strikingly similar across all these exercises. For this reason, the main text proceeds in straight-ahead fashion, without continual references to the Appendix.

3 Learning from Seller Experiments

In this section, we use the experiments data to analyze selective questions about auction design, consumer behavior, and market outcomes. Because one of our goals is to illustrate the scope of the approach, we consider several different questions, focusing on ones that have been relatively central to discussions of internet commerce. We relate our results to prior findings along the way. In ongoing work (Einav et al., 2012), we show how seller experiments can be combined with other theoretical and empirical approaches to explore more fully some of the findings about consumer preferences and draw out the implications for optimal seller and market responses.

3.1 Price Dispersion and "Excessive Bidding"

We start by reporting some large-sample findings about the variability in auction prices. The first is that auction prices for identical items sold by the same seller vary substantially, by around 10-15%, even if one focuses on auctions that occur close in time. The second is that auction prices generally do not rise above equivalent posted prices. A third finding that reconciles the first two is that auction prices, on average, are significantly below equivalent posted prices.

These findings relate to an ongoing debate about price dispersion and consumer search in online markets. In principle, the low physical search costs on the internet should limit price dispersion. Yet studies by Bailey (1998), Brynjolfsson and Smith (2001), Baye, Morgan and Scholten (2004) and Ellison and Ellison (2009), all report substantial dispersion in posted prices, even on structured price comparison websites.[6] Recent work by Malmendier and Lee

[6] Ambrus and Burns (2012) provide a recent state-of-the-art theory of rational bidding behavior when bidders are not fully focused on an auction, and show that a wide range of price outcomes can be consistent with equilibrium.

(2011) also provides a striking “failure” of consumer search. They document an episode on eBay in 2004 in which a particular board game was available from two sellers for $129.95, while other sellers offered the game for auction. Malmendier and Lee find that auction prices exceeded the posted price more than 40 percent of the time, often by 10 dollars or more. They argue that this is inconsistent with rational search and that a significant number of consumers are irrationally over-bidding.

A complicating factor in existing studies is that prices are compared across retailers, and the prices being compared are typically posted prices rather than transaction prices. This makes it difficult to disentangle differences in retailer attractiveness from frictions in consumer search, or in some cases to rule out the possibility that consumers mostly ignore high-priced alternatives. Our experiments data allow us to identify the transaction price variability across auctions by a single seller, both on average and for different types of sellers and products. In addition, by focusing on auctions that overlap with the presence of an equivalent (same seller, same item) posted price listing, we can examine the Malmendier and Lee over-bidding hypothesis for a sample of hundreds of thousands of items.

We report our basic findings on price dispersion in Table 4. We use as our metric the coefficient of price variation, or the standard deviation of a group of auction prices divided by the mean price. We compute the coefficient of price variation for each experiment, and for a refinement in which we partition each experiment by calendar month. The average coefficient of price variation is 0.11 (0.10 with the finer partition of each experiment). The degree of variation climbs to 0.15 if we also consider matched sets of auction listings (same seller, same item) that do not have a matched fixed price sale. In contrast, there is less variability for experienced sellers, or when the seller uses a BIN option or a higher reserve price. Overall, however, 10% to 15% price variation across equivalent auctions appears to be a pervasive feature of the market.

Next, we report on how auction prices compare to equivalent posted prices. Recall that we defined an item’s reference price or “value” v_i to be the average price across posted price sales of the item by the same seller. For a successful auction with price p_{it}, define p_{it}/v_i to be the relative price. Figure 3(a) plots the distribution of relative auction prices for items with values less than $10, between $10 and $30, between $30 and $100, and between $100 and \$1,000. Our data also include a few goods that sell for posted prices above \$1,000, but they are sufficiently rare that we drop them to focus the analysis.

Auction prices are strikingly low compared to equivalent posted prices. The average relative price is around 0.84, and the median is around 0.87. So around half of the auction sales we observe occurred at a discount of 13 percent or more relative to the posted price. We also can examine the prevalence of excessive bidding in which the auction price exceeds the reference value. This is relatively atypical. Less than 20 percent of auction prices exceed the reference price, and most of these episodes involve very small overpayments. To see this, Figure 3(b) plots the analogous distribution of $p_{it} - v_i$, the absolute (dollar) difference between the auction and reference price. Of the 1,178,855 successful auctions in our sample, only about 5 percent result in prices more than \$10 above the item's posted price.

To be consistent with the subsequent analysis in the paper, Figure 3 compares auction prices to the average posted sale price of the same item over the course of the year. If one is looking for over-bidding, a more apt comparison might be to a concurrent posted price offered by the same seller, should one exist. In the appendix we repeat the analysis, limiting attention to auctions for which there was a matched posted price offer available at the auction close (when most bidding occurs). Our data includes 98,536 successful auctions that meet this criteria. Interestingly, when we replicate Figure 3 for this smaller sample, the results are nearly the same, with the vast majority of auction sales occurring below the posted price and very few meaningfully above (see Appendix Figure G.3).

To summarize, we have used hundreds of thousands of matched auction listings to document significant price variation across sales of identical goods by identical sellers. The same approach indicates that auction prices exceed their matched posted price rather infrequently, and on average are well below. The latter finding suggests that consumers who pay the posted price, rather than getting a discount by avoiding auction fever, are paying extra for the convenience of an immediate guaranteed purchase. We explore this issue, and the implications for sellers in deciding whether to offer items by auction, posted price or both in Einav et al. (2012).

3.2 Auction Start Prices and Demand Curves

In this section, we show how variation in auction start prices (or reserve prices) can be used to test some basic principles of auction theory and to trace out nonparametric auction demand curves. In a standard private value auction model, an increase in the reserve price lowers the probability of a successful sale, but raises the price conditional on sale. The price increase occurs because increasing the reserve price from s to s' either eliminates sales that would occur at prices between s and s' or forces their price up to s'. Conditional on the auction price increasing above s', the distribution of sale prices is the same whether the reserve price was s or s'.

There are other models of auctions in which reserve prices can have more nuanced effects. These include models with entry or bidding costs, or with common value elements, or behavioral models. For instance, Ku et al. (2006) argue that bidders may exhibit escalating commitment so that lower start prices increase the odds of a sale and also the price conditional on sale. They present supporting evidence of this based on a sample of Persian rug and digital camera auctions on eBay. Simonsohn and Ariely (2008) found that while lower start prices did not necessarily increase the price conditional on sale, they did increase the price conditional on it rising above the higher start price – again consistent with a "bidding frenzy" theory. In contrast, other researchers (Kamins et al., 2004; Reiley, 2006; Lucking-Reiley et al., 2007) found that lower start prices generally led to lower prices conditional on sale, without testing the upper tail.

We take a large-sample approach to these hypotheses using our experiments data. There are 142,653 experiments in our baseline sample with variation in the start price. To limit the variation in other auction parameters, we restrict attention to listings with free shipping, no secret reserve price, and no BIN option. This leaves 19,777 experiments with start price variation, encompassing a total of 494,170 listings, or about 25 listings on average per experiment. As above, we normalize start and sale prices by the items' reference values, so a start price of 0.35 means that a particular auction started at a price that was 35% of the item's posted price.

There is a stunning amount of variation in start prices. The top panel of Table 5 shows the overall distribution of start prices for items of different values. The bottom panel summarizes the within-experiment price variation. For the latter, we find the minimum and maximum start price for each experiment, and cross-tabulate the experiments according to these numbers. It is quite common for a seller to auction the same item multiple times with widely different start prices. For instance, of the 3,262 experiments that contain at least one very low start price ($p_{it}^n < 0.05$), 1,401 (43 percent) have at least one listing with a start price of $p_{it}^n > 0.85$, and several hundred have at least one start price of $p_{it}^n > 1$. As we discuss below, there are fewer intermediate start prices, but still enough to obtain robust estimates.

We use this variation to estimate fixed-effects regressions where the dependent variable is either an indicator for a successful sale or the price conditional on sale. We allow the start price to have a flexibly estimated non-linear effect by using a set of indicator variables for different start price levels. The regression results are presented in Table 6, and in Figure 4.

Figure 4(a) (top left panel of Figure 4) plots the effect of the (normalized) start price on the probability of sale. A sale is almost guaranteed when the start price is very low, but the sale probability drops to less than 0.2 for high start prices. The figure shows separate sales curves for each of our four value categories. These come from separately estimated regressions, so that each plot is an average sales curve for a set of items of roughly similar value. The sales curves are remarkably similar (and close to linear) across price categories. Thus it appears that the probability of sale depends a great deal on the start price relative to the item's value, but not so much on the value of start price per se.

Figure 4(b) (top right) plots the effect of the auction start price on the final sale price. The relationship is estimated only for auctions that result in a sale. The estimates are again remarkably similar across price categories. For start prices below 0.6, the expected auction price conditional on sale is generally around 0.8. One interpretation of the flat price curve for lower start prices is that there is enough competition in the market to keep auction prices from slipping very far even if the start price is very low. For higher start prices, of course, start prices must exceed 0.8, and indeed the estimated price curves are upward sloping in this range.

In Figure 4(c) (bottom left), we combine these estimates to obtain auction demand curves. For each possible start price, we plot the probability of sale against the expected price conditional on sale on the y-axis. As the start price varies, we trace out demand curves. To make the figure clear, we only show the auction demand curve for a sample that pools all value categories, but category-specific demand curves are very similar. A somewhat unexpected finding is that the auction demand curve is highly convex, and the associated marginal revenue curves is not downward sloping as in standard analyses. Instead, the marginal revenue is roughly U-shaped, as shown in Figure 4(c) which plots a (smoothed) marginal revenue curve for the pooled sample.[7] With this type of demand, a seller would prefer either a high start price or a very low one, depending on his marginal cost, and not an intermediate start price, consistent with bimodal distribution reported in Table 5.

How do our estimates relate to the various theories described above? That the estimated demand is downward sloping, while hardly surprising, runs counter to the strong version of the "bidding escalation" theory. It is also interesting to investigate the weaker version of this hypothesis, namely that *conditional on reaching a given price*, an auction that started at a lower price will continue longer. In contrast, the textbook private value auction model predicts that the upper tail of the price distribution (say, above some price p) should not depend on the start price (for any start price less than p). To investigate this, Figure 4(d) (bottom right) plots, for low, medium and high start prices, the probability of the auction price rising above different thresholds. The plot is similar for medium and high start prices, consistent with the textbook auction model, but low start prices do raise the (unconditional) probability of getting a high final price. One potential explanation is that low start prices attract more bidder attention. However, the evidence does not particularly support the escalation hypothesis, as the figure also shows that the probability of the price exceeding 1.0 conditional on reaching 0.85, or exceeding 1.1 conditional on reaching 1.0, is essentially the same for low start prices as for medium and high ones.

[7] To construct the marginal revenue curve in Figure 4(c), we smooth the demand estimates. The exact procedure is described in Appendix B, which also shows the smoothed and unsmoothed plots. The smoothed and unsmoothed demand curves look nearly identical, but small wiggles in the unsmoothed curve create a few outlier points in the unsmoothed plot of marginal revenue.

3.3 Auctions with "Buy It Now" Prices

Sellers on eBay can experiment with a variety of auction design parameters apart from the reserve price. A particularly novel feature on eBay is that for a small fee of 5 to 25 cents, a seller can specify a "buy it now" (BIN) price at which a buyer can preempt the auction and immediately purchase the item. The BIN price disappears if the item receives a qualified bid, and a standard auction ensues. In a standard private value auction model with symmetric bidders, a seller would not want to offer the BIN option, because it can allow a low value bidder to preempt a high value one, and this reduces expected revenue. In practice, however, a BIN price may attract consumers who are too impatient to participate in an auction, or encourage bidders to submit an early bid rather than wait until the last minute. Another hypothesis we have heard is that a buy price can serve as a reference point in later bidding, potentially affecting the auction outcome.

Several studies have looked at the use of buy-it-now prices in practice. Standifird et al. (2004) auctioned silver dollars and found that buyers tended not to use the buy-it-now option even when the BIN price was set low. Ackerberg et al. (2006) analyzed Dell laptop auctions and found that sellers using a BIN option had revenue that was $29 higher. Anderson et al. (2008) collected data on sales of Palm handheld devices, and found that the BIN option was used more often by experienced sellers. In their summary statistics, the prices of BIN auctions are slightly higher but they do not report a comparison after controlling for seller or item characteristics.

The seller experiment approach allows us to provide large-scale evidence that extends and sharpens these earlier analyses. We first identify the 90,404 experiments in our baseline sample that have variation in the BIN price, or in whether the BIN option is used at all. To avoid confounding BIN choices with other auction parameters, we restrict attention to listings with free shipping, no secret reserve price, and a start price that is effectively non-binding (specifically listings with a value of at least $10 and a start price of less than $1). This leaves us 3,239 experiments with BIN variation, and a total of 123,757 listings. Table 7 documents the amount of variation in (normalized) BIN prices, both across the whole sample (top panel), and within experiments (bottom panel). Most BIN prices fall between 80% and

120% of an item's average posted price, with considerable variation in this range.

We use the within-experiment variation in BIN prices to identify their effect on sale outcomes. Because we focus on listings with essentially non-binding reserve prices, almost all (98 percent) end in a successful sale. Therefore we focus on whether the item sells via the BIN price or instead via the auction mechanism. The top panel of Table 8 reports results from fixed effect regressions in which the dependent variable is a dummy equal to one if an item sells via the BIN price. Items are quite unlikely to sell at high BIN prices, especially prices more than 10% above the item's reference value. It is more common for a buyer to exercise the BIN option when the BIN price is less than 90% of the reference value, but it still happens only for a minority of listings. The top panel of Figure 5 plots the results.

The bottom panels of Table 8 and Figure 5 show the relationship between the BIN price and auction revenue. The results here are based on fixed effects regressions in which the dependent variable is the transaction price. The median (normalized) BIN price in our sample is between 0.95 and 1.00. Setting a BIN price at this level appears to have a negligible effect on revenue. Setting a lower ("under-priced") BIN price apparently reduces seller revenue, while setting a high ("over-priced") BIN price modestly increases revenue. These results are roughly consistent with the price discrimination story above, in which even a high BIN price might still be attractive to certain buyers but would not undermine the auction process if the right buyer did not arrive.

A more subtle question is whether offering a BIN option that is *not* exercised might affect subsequent bidding, for example by anchoring subsequent bids. This is a tricky question to get at, and the anchoring mechanism seems somewhat unlikely in that the BIN price disappears once a qualified bid is received. Nevertheless, to investigate it, we consider how different BIN prices affect the probability of obtaining a sale price below certain thresholds, that is, the lower tail of the cumulative price distribution. The bottom panel of Figure 5 shows that the likelihood of receiving below 60% of the reference value is essentially the same whether the seller sets a high BIN price, a low BIN price, or no BIN price at all. Offering a high BIN price, however, does appear to somewhat reduce the probability of a sale in the 70-100% range.

3.4 Other Aspects of Auction Design

The seller experiments data provides a rich laboratory to explore the effects of other auction design parameters. While we hesitate to overwhelm the reader, we briefly mention a few that are illustrative and relate to earlier work.

A number of studies have found that longer auctions seem to generate higher revenue (Lucking-Reiley et al., 2007; Haruvy and Popkowski Leszczyc, 2010), or have analyzed the effect of ending auctions on different days of the week or at different times of the day (Simonsohn, 2010). Using a similar empirical strategy to the one employed so far, we identified 92,266 experiments with variation in auction duration, 129,955 experiments with variation in the ending time, and 126,027 experiments with variation in the ending day. Our results suggest that overall the effect of the auction duration is small. On average, we find that longer auctions with a BIN option are slightly more likely to succeed while auction duration makes little difference for the sale probability of standard auctions with no BIN option. The effects are not large, however, and are less robust than most of our other findings. We also find little effect of the day of the week on which the auction ends, and we confirm existing results that auctions that end late at night (midnight to 5am) perform slightly worse.

Another issue that has attracted some debate is the effect of keeping auction reserve prices secret. On eBay, the seller sets a public reserve price in the form of the auction start price we analyzed earlier, but (for an additional fee) can also set a secret reserve price that is not known to potential bidders. When a seller sets a secret reserve price, bidders know that it exists, but learn its level only if bidding in the auction exceeds it. Various factors might make a secret reserve price more or less profitable than a public reserve price. For instance, Katkar and Reiley (2006) auctioned 100 Pokemon cards, half with a public reserve price of 30% of the item value and half with a secret reserve of 30% of the item value (and effectively a zero starting price). They found that secret reserve prices resulted in lower revenue.

To investigate this question using our data, we match listings of the same item into groups that have similar levels of public and private reserve prices (specifically, we do this in multiple ways: either by matching listings that have exactly the same reserve price, or – to increase statistical power – by matching listings with reserve prices within 10% of each other). Because the use of secret reserve prices has been discouraged by eBay and is not very popular (less than one percent of eBay listings use a secret reserve, and only 0.60% in our baseline sample), our power is much lower than in previous exercises. Nevertheless, we do find 403 matched groups of listings, so we can estimate the effect of using a secret reserve price versus a public reserve price of the same magnitude. Our results indicate that in this sample, there is not much difference in auction outcomes between the public and secret reserve price sales.

3.5 Shipping Fees

Shipping arrangements are an important part of internet commerce, and internet retailers frequently compete to offer free or expedited shipping. At the same time, one often hears the idea that shipping fees can act as a hidden price that buyers do not fully internalize in making shopping decisions. Tyan (2005), Hossain and Morgan (2006), and Brown et al. (2010) all have studied data from eBay and found that increases in shipping fees can increase total seller revenue (inclusive of the shipping fee), suggesting that a dollar increase in the shipping fee does not lead bidders to reduce their bids by a full dollar to compensate. Sellers also can have another reason to favor shipping fees: until recently, eBay commissions were not applied to the shipping component but rather to the pre-shipping fee sale price.

We are interested in whether buyers internalize shipping fees. To analyze this, we follow the empirical strategy we have been employing throughout, and select experiments from our baseline data that have variation across listings in the shipping fee. To avoid complications, we consider only listings with flat shipping fees that are independent of the buyer location.[8] The resulting data contains 117,202 listings grouped into 6,655 experiments, with an average of 18 listings per experiment. A substantial fraction of these listings offer free shipping. Table 9 presents the distribution of shipping rates across the listings, and also the within-experiment variation in shipping fees. In parallel with our earlier analyses, we see sellers trying a range of shipping fees.

[8]Five percent of the listings in our baseline data are associated with a shipping fee that depends on the location of the buyer. To simplify, our analysis focuses on the remaining 95%. Further excluding listings with contradictory shipping information in the data leaves us with 89% of the listings that have a flat shipping rate.

Table 10 reports results on sale probability and auction revenue (conditional on sale) for several different subsamples. The effect of shipping fees on the probability of sale is minimal, so we focus on the price effect. Unlike our earlier analyses, we run the price regression without normalizing by the item value, as this helps in facilitating the quantitative interpretation of the estimated effects. With this specification, a coefficient of zero on shipping rate implies that bidders respond to shipping fee changes one-for-one, so that a higher shipping fee is fully canceled out by a lower sale price, and the effect on total revenue (sale price plus shipping) is zero. As Table 10 indicates, our estimates suggest a positive coefficient of around 0.2 to 0.3, suggesting that only 70 to 80 percent of the shipping fee is internalized in the bidding.

In addition, we find a distinct effect at zero. Free shipping is associated with an average revenue increase of around \$2.50, with a larger dollar effect for more expensive items. The free shipping effect may be some combination of buyers responding to a "free" offer (Shampanier et al., 2007) and eBay's strategy of prioritizing free shipping in the search results. Figure 6 provides a graphical illustration of our regression estimates. As shown in the figure, our estimates suggest that low shipping fees on eBay, of roughly less than \$10, are suboptimal. Sellers could increase profits by either reducing the shipping rate and making it free, or by increasing the shipping rate and benefiting from the fact that bidders would only partially internalize this increase. The observed distribution of shipping fees is largely consistent with these incentives: only a small fraction of the listings are associated with low (but positive) shipping fees (top panel of Table 9).

An even more finely targeted way to analyze the effect of shipping fees is to focus on cases where an increase in the shipping fee was matched by a reduction in the start price. A textbook economic analysis would suggest that an auction with free shipping and a start price of \$10 should be *identical* to an auction with a zero start price and a \$10 shipping fee. That is, they should have the same probability of sale and expected revenue conditional on sale. Following this logic, we identified duplicate listings with same "inclusive" start price, that is, the same sum of start price and shipping fee. We then asked whether the division mattered. We found 279 such experiments in our baseline sample (and many more where the inclusive start prices were within a 10 cent or 10% range). Our finding in this sample was similar: increases in the shipping fee reduce the sale price, but less than one-for-one.

4 The Advantage of Seller Experiments

Seller experiments provide a simple way to isolate variation in a range of sale attributes, providing the opportunity to use large and diverse data to measure various parameters of interest in a scalable fashion. But does the approach convey any particular advantages relative to other creative strategies for estimating treatment effects? The two most obvious alternatives in a setting such as eBay are cross-item observational comparisons that attempt to control for confounding sale attributes, or alternatively, the use of field experiments in which the researcher lists items and varies different sales parameters.[9] In this section, we illustrate some potential benefits of seller experiments relative to these alternatives.

4.1 Relative to Observational Data

The key concern with observational data in a setting like eBay is the heterogeneity of the items that are being listed for sale and the sellers doing the listing. This makes it difficult to specify an appropriate set of control variables to yield apples-to-apples comparisons, particularly when many item attributes such as the listing title, item pictures and description are relatively "unstructured." We use a variant of the start price analysis from the previous section to illustrate this point. Our illustration entertains what researchers might have done if they had access to the same data, but were not able to group listings into seller experiments.

Absent such a grouping, a researcher presumably would have tried to define comparable sets of products in some other way. One natural way to group items is to rely on eBay's product categories. eBay classifies products using a hierarchical category structure. At the highest level, listings are partitioned into almost 35 "meta categories," such as electronics, collectibles, baby items, and so on. At the finest level, products are partitioned into 37,636 "leaf categories," such as "iPod and MP3 players" and "developmental baby toys." Thus, one way a researcher could analyze the effect of start price is to compare listings within a given leaf or (less ideally) meta category.

[9] Another possibility might be changes on the platform or the surrounding environment that act as an instrumental variable by encouraging sellers to shift sale parameters, or specific institutional features that give rise to regression discontinuity or other quasi-experimental designs (see, e.g., Choi, Nesheim, and Rasul, 2011). These approaches have been relatively rare and would appear best-suited for examining one particular sale parameter at a time, rather than being easily scalable.

We examine this strategy by running our start price exercise in three different ways: grouping listings in our baseline sample according to their meta category, their leaf category, and by seller experiment. In the former two cases, we average item reference values within each category to create a category-specific reference value, as if all items within the category were perfectly comparable. We then use this average value to normalize the start price for each listing in the category, and re-estimate the effect of start price on an indicator for a successful sale and the final (normalized) price conditional on sale, including fixed effects for the relevant item groupings, but also omitting the fixed effects for comparison. For simplicity, we report results only for the probability of sale, and not the price conditional on sale.

We report the results in Figure 7, which plots the differently estimated sales curves as a function of start price. The estimates for which we group items by (either meta or leaf) category are dramatically different from what we obtain by grouping identical listings into experiments. To understand the difference, we can interpret the solid black curve in Figure 7 as an average estimate of how the sale probability changes with the start price for a fixed item (and seller). In comparison, the solid grey curves are constructed so that the composition of the items offered at different start prices is not the same, although they are all in the same product category. The differences in the estimated sales curves indicate that items offered at very low and very high start prices are generally more appealing (in the sense of having a higher probability of sale) than those offered at medium start prices.

Two other patterns in Figure 7 are worth noting. First, the inclusion of fixed effects in all three analyses makes very little difference. That is, it appears that — at least for this analysis — the effect of grouping listings into eBay product categories or into sets of identical items is captured mostly in the construction of the reference value by which we normalize the start price. Second, it is interesting to note that although the meta category level is an extremely crude way to categorize products while the leaf category level is an extremely precise classification, the results obtained from these two exercises are very similar, and both are dramatically different from the "fixed item and seller" grouping we rely on using the experiments approach.

Overall, the analysis points to a considerable problem of accounting for heterogeneity in large diverse markets such as eBay. This is presumably one reason researchers working with data from eBay or other online markets typically have restricted attention to a very narrowly defined groups of products, such as particular pop-music CDs, collectible coins, Pokemon cards, or board games. A narrowly drawn set of products may (or may not) mitigate the problem just identified, but even if it does as in the case of a researcher-conducted experiment, it raises the concern that the results apply only to a narrow context. It is to this separate concern that we now turn.

4.2 Relative to Field Experiments

The same ease of listing and selling items that makes seller experiments so prevalent on eBay and other online platforms also makes these settings appealing for researcher-initiated field experiments. Administering and funding experiments is costly, however, so although researcher experiments are common, they are typically quite small in scale and scope, focusing on one of a few items, in limited quantity, and varying just one or a few sale parameters to identify a very limited number of treatment effects.

Relative to such exercises, the key advantage of seller experiments is scale and scope. While, naturally, seller experiments are not as controlled as field experiments, we have shown that it is possible to identify millions of seller experiments conducted just in a single year on eBay, and that these experiments cover a wide range of product categories, price levels, and sale characteristics. The scale makes statistical power a non-issue, thus significantly reducing the possibility of both type one and type two errors. The scope allows researchers to isolate a wide range of effects, and also to assess whether an effect observed in a particular product category is broadly representative, or if there is substantial heterogeneity across product categories or price levels in the effects of different sales strategies.

To illustrate this last point we again return to our analysis of auction start price, and re-run the exercise separately for each product meta-category. To facilitate a graphical illustration, we estimate a linear effect of the (normalized) start price on both the probability of sale (by regressing an indicator equal to one if the item sold on the start price and experiment fixed effects), and the expected (normalized) price conditional on sale (by regressing the sale price on the start price and experiment fixed effects, using only successful sales). This yields, for each category, the slope of the average sales curve for items in the category and the slope of the price curve conditional on sale, with both probability of sale and price being a function of the start price.

The results are presented in Figure 8 and Table 11. The x-axis shows the effect of start price on the probability of sale, so a value of -0.5 means that an increase in the start price from 0.5 to 0.8 as a fraction of the item's value reduces the probability of sale by 0.15. The y-axis shows the effect of start price on the expected price conditional on sale; a value of 0.1 means that an increase in the normalized start price from 0.5 to 0.8 increases the expected price conditional on sale by 3% of the item's value. Each point in Figure 8 shows the two effects of the start price for a particular eBay product category.

Certain features are consistent across all categories. A higher start price always reduces the probability of sale, and (with the exception of DVDs where the effect is near-zero) increases the average price of successful sales. Yet, the magnitude of the effects varies quite dramatically across categories. For example, one can imagine a researcher running a careful field experiment on eBay by listing DVDs (or, more likely, specific types of DVDs), randomly varying their start prices, finding a large effect on the quantity sold, but very little effect on price. This researcher may have no reason to believe that DVDs are special, and therefore conclude that start prices do not affect sale prices, which may be consistent with some theories and less consistent with others. Yet, as Figure 8 suggests, such conclusions would be misleading, as the DVDs category is quite an outlier, and the price effects are significantly larger in all other product categories.

Of course, once one sees the results presented in this way, the differences across product categories become quite natural. Roughly, one can think of categories with a small dp/ds effect or a large (more negative) dq/ds effect as categories with relatively flat (i.e. elastic) residual demand curves for individual items, as opposed to relatively steep (inelastic) residual demand. So Figure 8 tells us that products listed in seemingly commodity categories such as DVDs, Electronics, Video and Coins fall into the former elastic category, whereas products listed in potentially more differentiated categories such as Clothing, Jewelry, Sports Memorabilia and Home fall into the inelastic category. While a full exploration is well beyond the scope of the present paper, Figure 8 suggests the possibility of using our approach to obtain meaningful comparisons of price sensitivity and competition across retail product categories.

5 Conclusion

In this paper we present a new approach to studying behavior and competition in internet markets, by taking advantage of the ease and prevalence of active and passive experimentation in these markets. The approach combines the advantages of small-scale field experiments run by researchers with the scale and scope of internet markets and data. It attempts to avoid the major identification problems in large observational studies, but also the narrowness, small sample sizes and limited scope of many field experiments. Of course, relying on experiments of others implies less control over the variation, and may raise concerns regarding some aspects of this variation; however, the size of the data allows us to empirically assess how important these concerns may be.

To illustrate our approach, we considered a series of applications looking at different aspects of consumer behavior and pricing strategies. We extend prior work by estimating in large-sample data: the degree of price dispersion across equivalent auctions, the relationship between auction prices and equivalent posted prices, the (average) shape of auction demand, the effect of buy-it-now prices, and the extent to which consumers internalize shipping fees. Our empirical approach is sufficiently straightforward that given the right data, it should be easy to replicate and apply to other questions. We also expect that a similar approach could prove fruitful in other internet retail, advertising or labor markets. It can also be applied retrospectively to understand changes in markets over time, an approach we are taking in ongoing work.

One question that we have not addressed in this paper, but which we believe is an interesting one, is whether sellers themselves are successful in learning from their experimentation. Some of the patterns we have documented — for instance that sellers generally tend to avoid intermediate start prices, or low but positive shipping fees — are consistent with the idea that sellers have over time accumulated knowledge about strategies that do not work well. In other cases, sellers face non-obvious trade-offs — for instance between a lower quantity and a higher price — where the optimal decision depends on seller costs that we do not observe. We also have found that sellers do not converge in their listing behavior for a given item; instead, they persistently experiment by varying their sale parameters. This suggests that a successful theory of active experimentation in online marketplaces would be one in which sellers remained somewhat unsure over time about exactly what strategy is best. Understanding how online retailers become more effective, and the process through which this occurs, is something we hope to explore in further work.

References

Ackerberg, Daniel, Keisuke Hirano, and Quazi Shahriar (2006). "The Buy-it-now Option, Risk Aversion, and Impatience in an Empirical Model of eBay Bidding." Mimeo, University of Michigan.

Ambrus, Attila and James Burns (2012), "Gradual Bidding in eBay-Like Auctions," Mimeo, Duke University.

Anderson, Steven, Daniel Friedman, Garrett Milam, and Nirvikar Singh (2008). "Buy It Now: A Hybrid Internet Market Institution." *Journal of Electronic Commerce Research* 9(2).

Athey, Susan (2011). "Structural Models and Market Design in the Age of Experimentation." Fisher-Schultz lecture, European Meeting of the Econometric Society.

Bailey, Joseph (1998). "Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce." Ph.D. Thesis, Massachusetts Institute of Technology.

Bajari, Patrick, and Ali Hortacsu (2003). "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions." *Rand Journal of Economics* 34(2), 329-355.

Bajari, Patrick, and Ali Hortacsu (2004). "Economic Insights from Internet Auctions." *Journal of Economic Literature* 42, 457–486.

Baye, Michael R., John Morgan, and Patrick Scholten (2004). "Price Dispersion in the Small and the Large: Evidence from an Internet Price Comparison Site." *Journal of Industrial Economics* 52(4), 463-496.

Brown, Jennifer, Tanjim Hossain, and John Morgan (2010). "Shrouded Attributes and Information Suppression: Evidence from the Field " *Quarterly Journal of Economics* 125(2), 859–876.

Brynjolfsson, Erik, and Michael Smith (2001). "Consumer Decision-making at an Internet Shopbot: Brand Still Matters." *Journal of Industrial Economics* 49, 541-558.

Choi, Syngjoo, Lars Nesheim, and Imran Rasul (2011). "Reserve Price Effects in Auctions: Estimates from Multiple RD Designs." Mimeo, University College London.

Einav, Liran, Chiara Farronato, Jonathan Levin, and Neel Sundaresan (2012). "Sales Mechanisms in Online Markets: What Happened to Internet Auctions," Stanford University mimeo.

Elfenbein, Daniel, Raymond Fisman, and Brian McManus (2010). "Charity as a substitute for reputation: Evidence from an online marketplace." Mimeo, Columbia University.

Ellison, Glenn, and Sara F. Ellison (2009). "Search, Obfuscation, and Price Elasticities on the Internet." *Econometrica* 77(2), 427-452.

Haruvy, Ernan, and Peter T. L. Popkowski Leszczyc (2010). "The Impact of Online Auction Duration." *Decision Analysis* 7(1), 99-106.

Hasker, Kevin, and Robin Sickles (2010). "eBay in the Economic Literature: Analysis of an Auction Marketplace." *Review of Industrial Organization* 37(1), 3-42.

Hossain, Tanjim, and John Morgan (2006). "...Plus Shipping and Handling: Revenue (Non)Equivalence in Field Experiments on eBay." *Advances in Economic Analysis and Policy* 6(2), Article 3.

Kamins, Michael A., Xavier Dreze, and Valerie S. Folkes (2004). "A Field Study of the Effects of Minimum and Reserve Prices on Internet Auction." *Journal of Consumer Research* 30(4), 622-628.

Katkar, Rama, and David H. Reiley (2006). "Public versus Secret Reserve Prices in eBay Auctions: Results from a Pokémon Field Experiment." *Advances in Economic Analysis and Policy* 6(2), Article 7.

Ku, Gillian, Adam D. Galinsky, and J. Keith Murnighan (2006). "Starting Low but Ending High: A Reversal of the Anchoring Effect in Auctions." *Journal of Personality and Social Psychology* 90(6), 975-986.

Lucking-Reiley, David (1999). "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet." *American Economic Review* 89(5), 1063-1080.

Lucking-Reiley, David, Doug Bryan, Naghi Prasad, and Daniel Reeves (2007). "Pennies from eBay: the Determinants of Price in Online Auctions." *Journal of Industrial Economics* 55(2), 223-233.

Malmendier, Ulrike and Young Han Lee (2011). "The Bidder's Curse." *American Economic Review*, 101(3), 749-787.

Ostrovsky, Michael and Michael Schwarz (2009). "Reserve Prices in Internet Advertising Auctions: A Field Experiment." Stanford Working Paper.

Reiley, David H. (2006). “Field Experiments on the Effects of Reserve Prices in Auctions: More Magic on the Internet.” *Rand Journal of Economics* 37(1), 195-211.

Shampanier, Kristina, Nina Mazar and Dan Ariely (2007). “Zero as a Special Price: The True Value of Free Products.” *Management Science* 26(6), 742-757.

Simonsohn, Uri (2010). “eBay’s Crowded Evenings: Competition Neglect in Market Entry Decisions.” *Management Science* 56(7), 1060-1073.

Simonsohn, Uri, and Dan Ariely (2008). “When Rational Sellers Face Non-Rational Consumers: Evidence from Herding on eBay.” *Management Science* 54(9), 1624-1637.

Standifird, Stephen S., Matthew R. Roelofs, and Yvonne Durham (2004). “The Impact of eBay’s Buy-It-Now Function on Bidder Behavior.” *International Journal of Electronic Commerce* 9(2), 167-176.

Tyan, Sean (2005). “The Effects of Shipping Costs on Bidder Entry and Seller Revenues in eBay Auctions.” Senior Thesis, Stanford University.

Varian, Hal (2010). “Computer Mediated Transactions.” *American Economic Review*, 100(2), 1-10.

Figure 1(a): A standard search results page on eBay

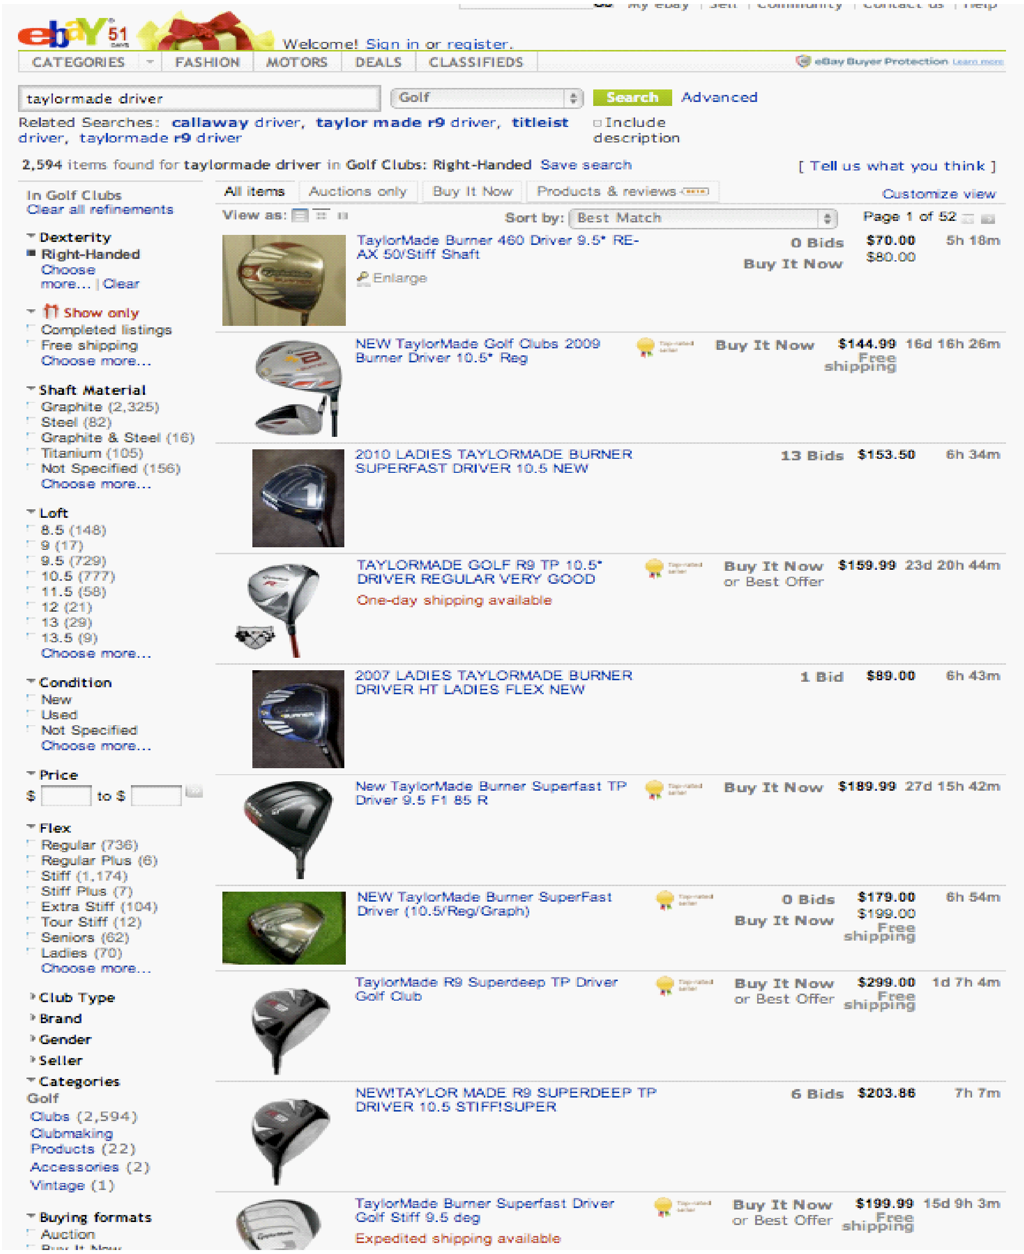

The figure presents a "standard" screenshot of listings on eBay, following a search for "taylormade driver" on 9/12/2010.

Figure 1(b): An example of a "seller experiment"

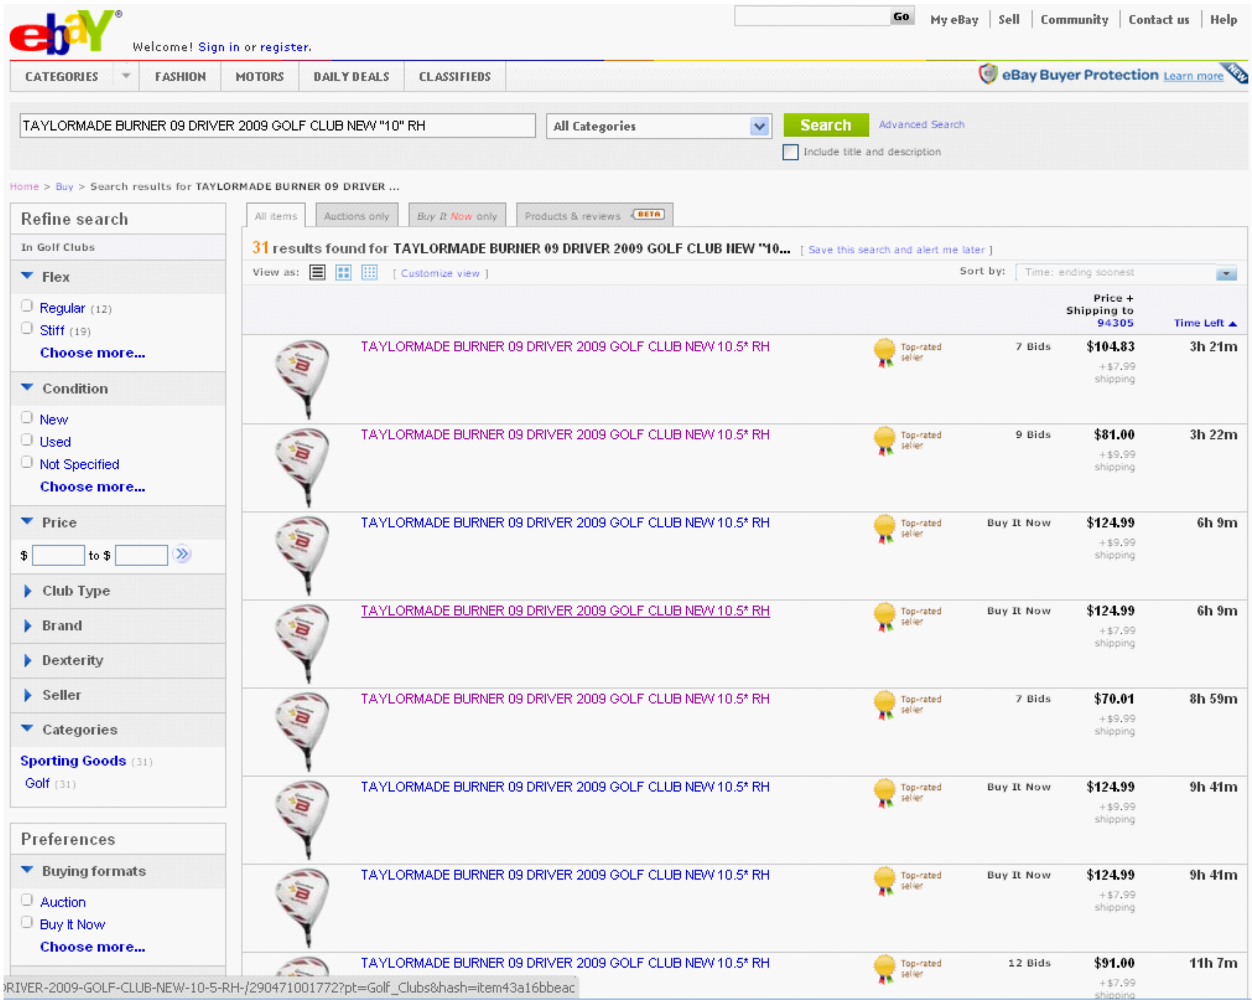

The figure illustrates a "seller experiment." It shows the first 8 out of 31 listings for the same golf driver by the same seller. All the listings were active on 9/12/2010. Of the eight listings in the figure, four are offered at a fixed price ("Buy It Now") of $124.99. The other four listings are auctions. The listings also have different shipping fees (either $7.99 or $9.99).

Figure 2: Number of listings in each experiment

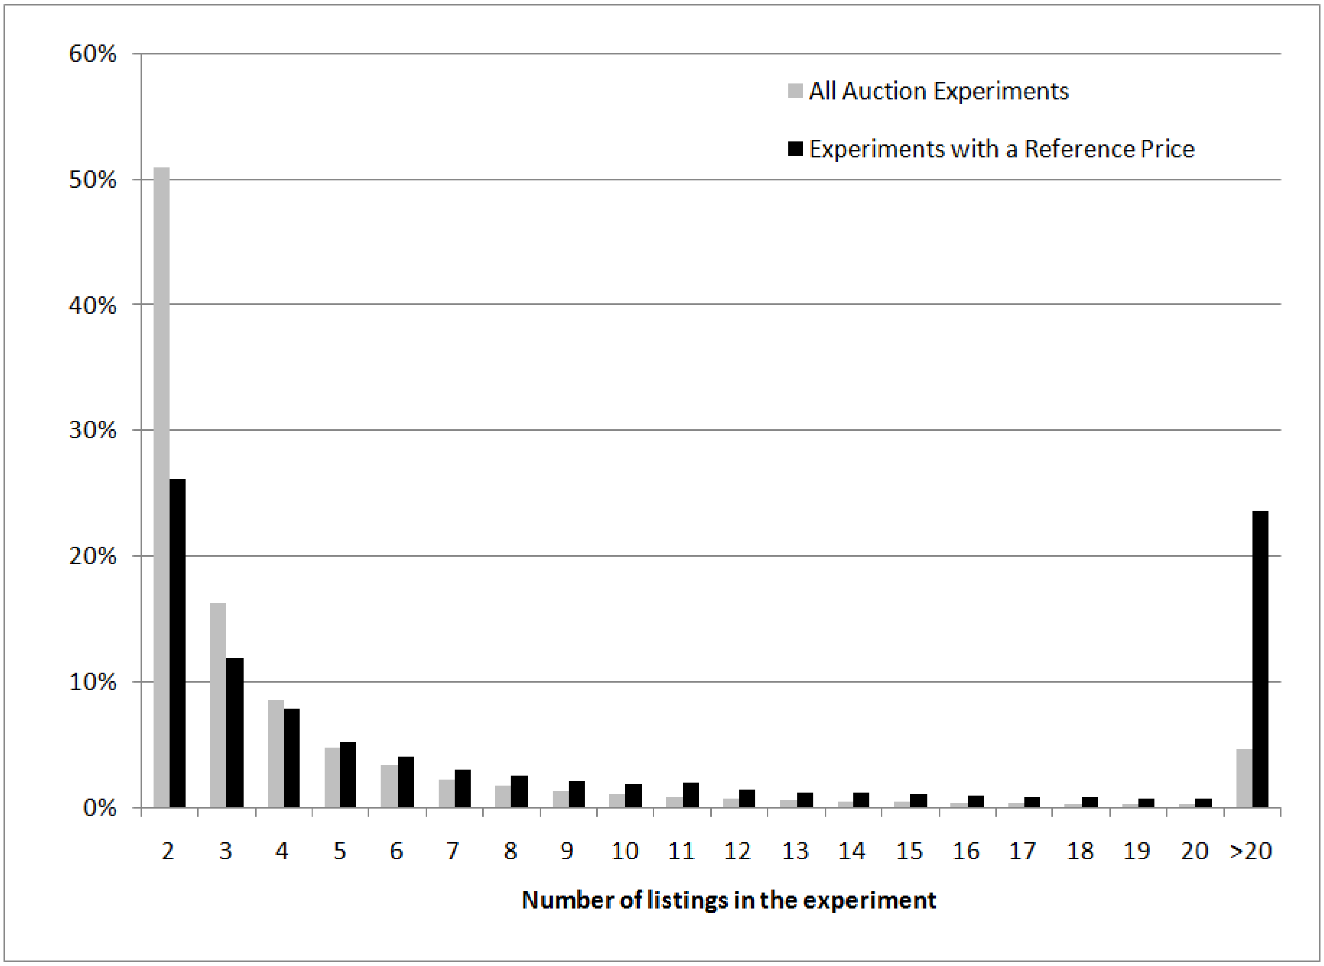

The figure presents the distribution of the experiment "size" (number of listings) in the entire auction experiments data (gray) and in our baseline sample (black).

Figure 3: Auction sale price dispersion

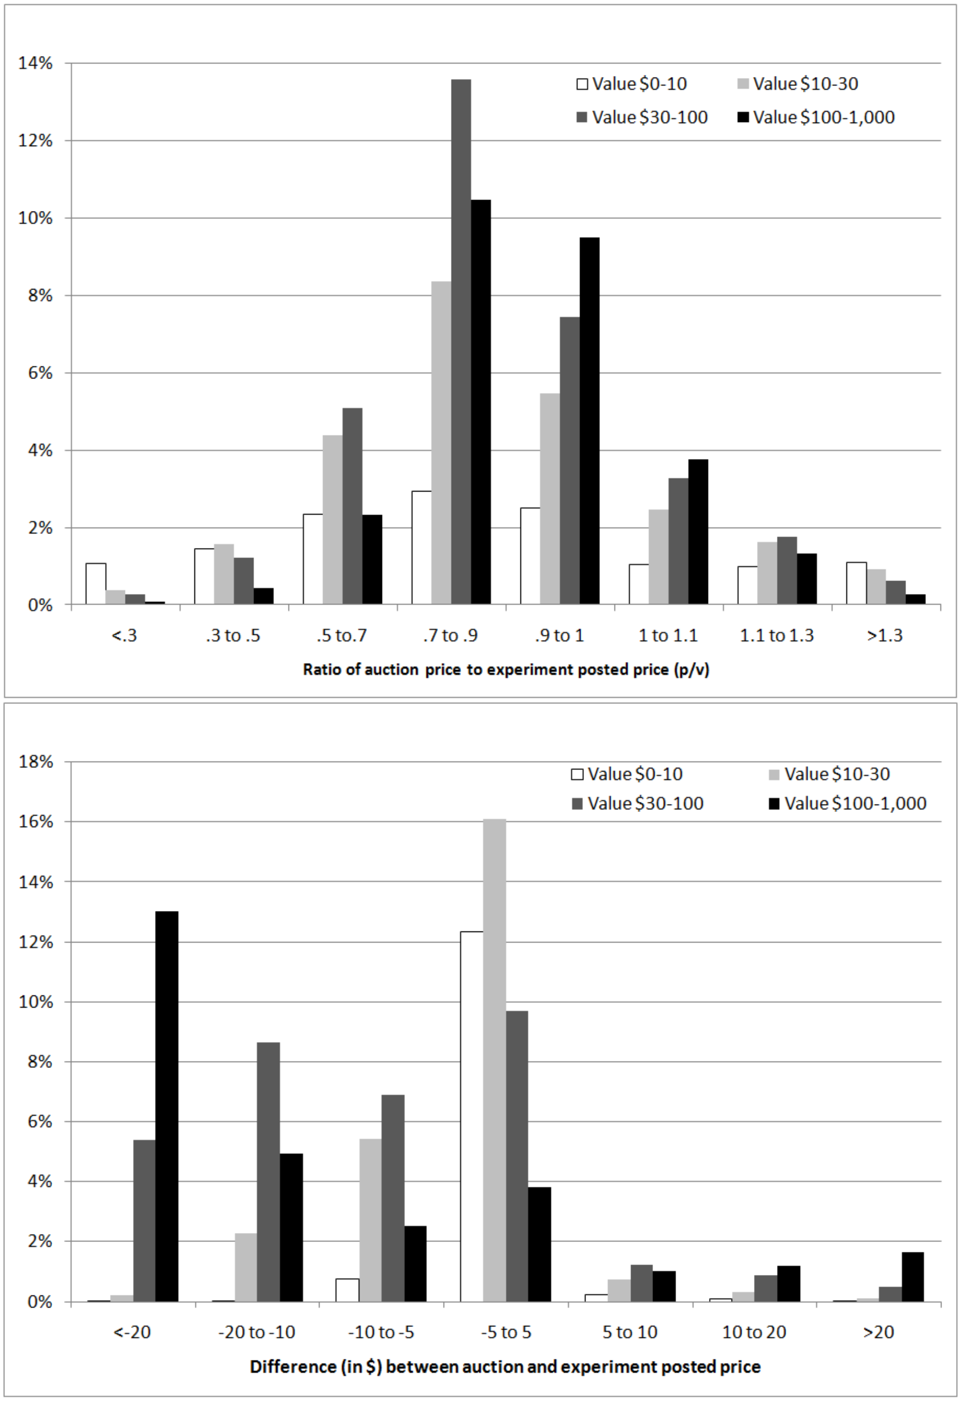

The figure shows the distribution of transacted auction prices p relative to the "reference value" v of the same item. The reference value for each item is defined as the average price across equivalent posted price transactions. The top panel shows the distribution of p/v, while the bottom panel shows the distribution of $p - v$. The figure omits items with a reference value greater than \$1,000. These comprise just 1.9% of the experiments and 0.5% of the listings in our baseline data.

Figure 4: The effect of auction start price

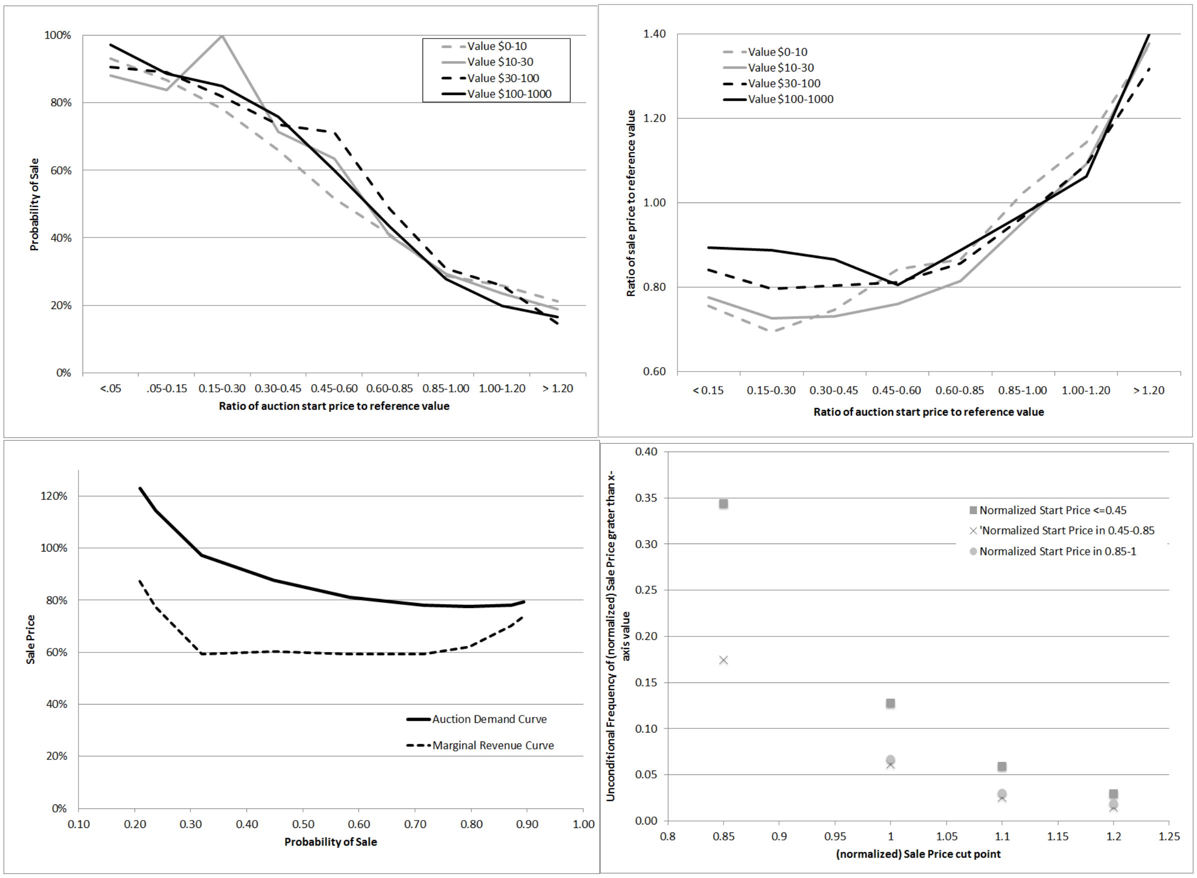

Figures 4(a) (top left) and 4(b) (top right) show the effect of auction start price on listing outcomes, based on the regression results in Table 6. Figure 4(a) shows the effect on the probability of sale; Figure 4(b) shows the effect on expected sale price. Figure 4(c) (bottom left) pools all value categories and presents the implied "auction demand curve" and its corresponding marginal revenue curve; see the main text and Appendix B for additional details. Figure 4(d) (bottom right) plots the probability a listing results in an auction price above certain levels, for different start prices (see text for further discussion).

Figure 5: The effect of BIN price

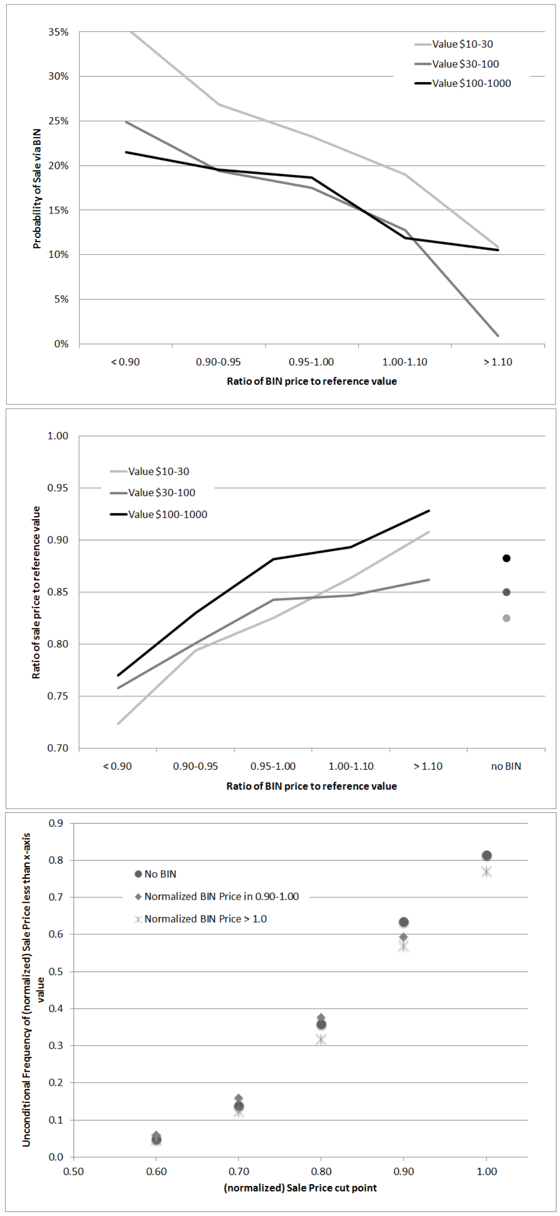

The top two panels show how the seller's choice of BIN price affects the probability the auction sells at the BIN price, and the listing revenue. The sample focuses on items with a starting price of less than one dollar, so essentially all listings sell. The plots are based on the regression results in Table 9. The bottom panel plots the probability the sale occurs at prices below certain levels, for different BIN prices (see text for further details).

Figure 6: The effect of shipping fees

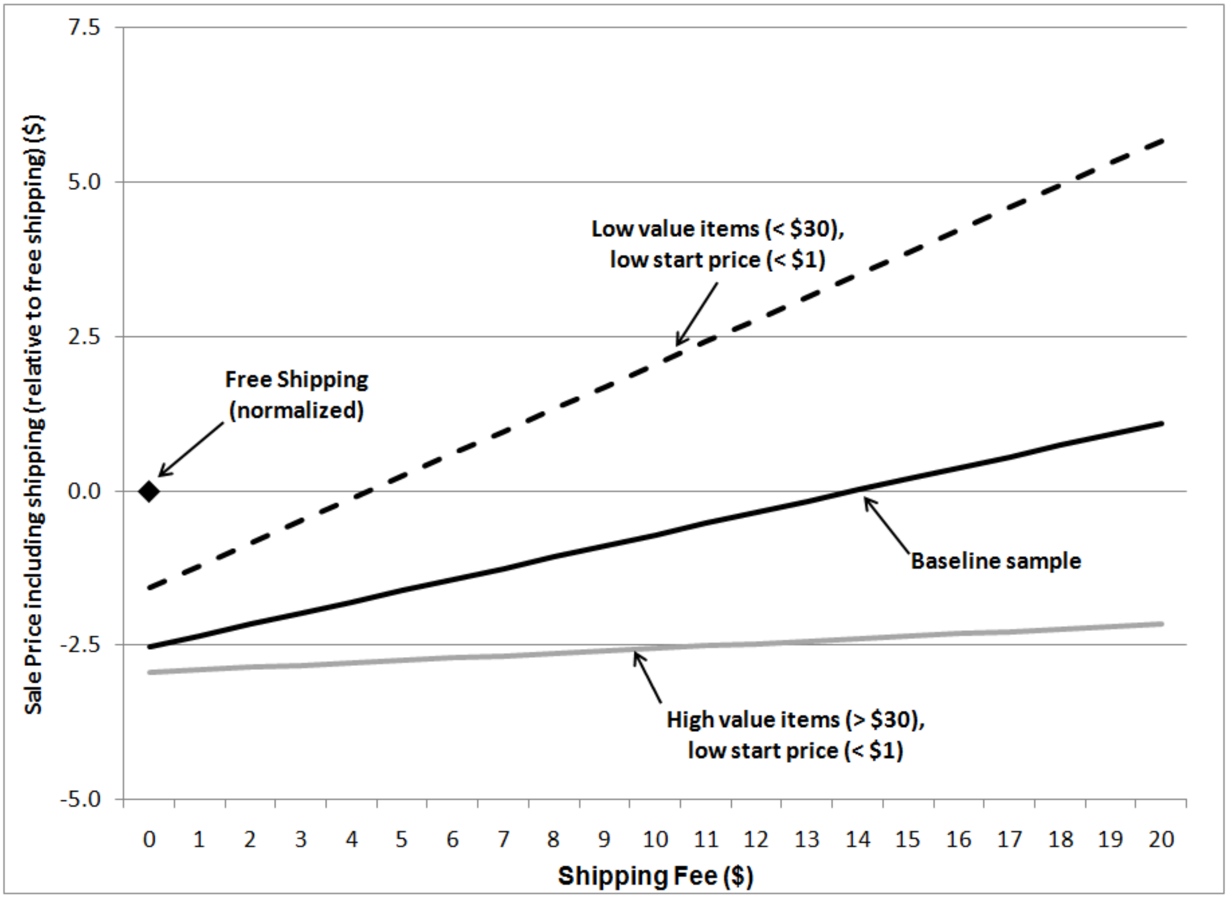

The figure shows the effect of shipping fees on seller revenue, based on the regression results in Table 10.

Figure 7: Seller experiments versus observational data

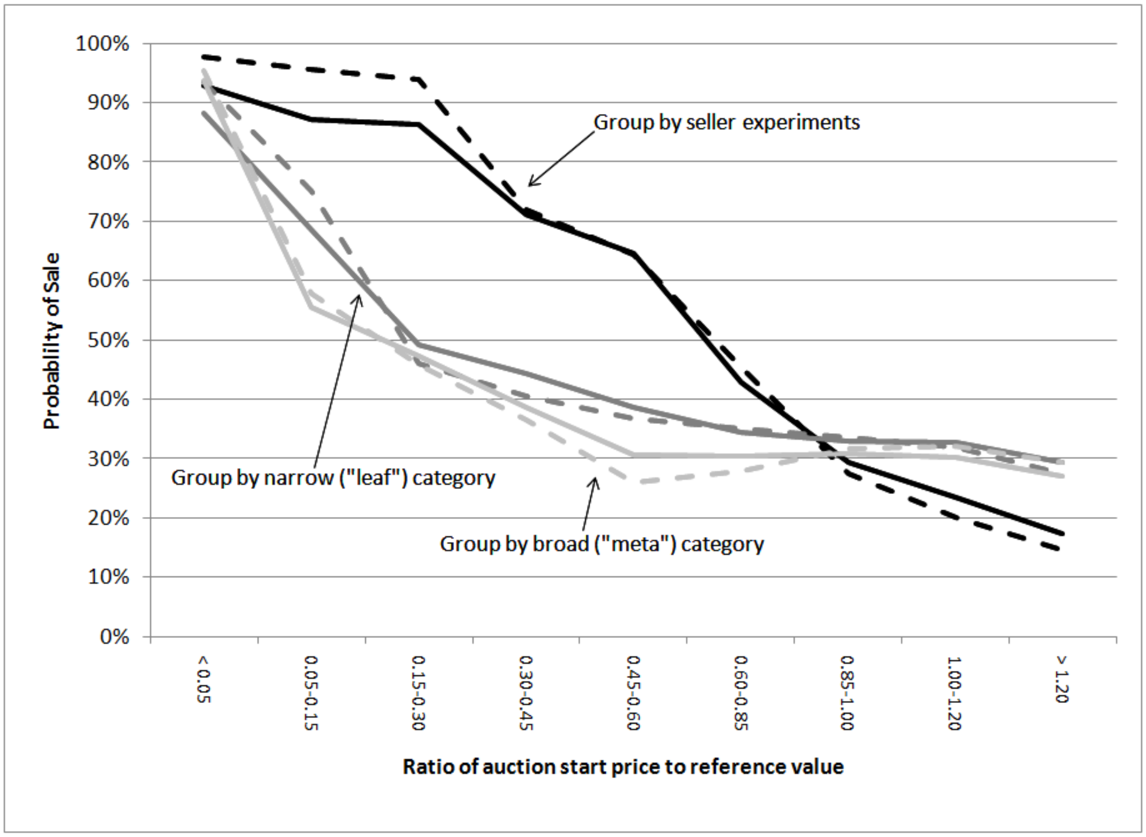

The figure presents the relationship between auction starting price and the probability of sale for the different regressions. The black lines represent start price variation within seller experiments, which is the type of variation used throughout the paper. The dark grey lines represent variation within narrow ("leaf") product categories as defined by eBay; there are more than 37,000 such categories. The light grey lines represent variation within broad ("meta") product categories as defined by eBay; there are 35 such categories. There are two lines for each grouping. The dashed lines represent specifications with no fixed effects, so that groupings are used to generate a reference value (average fixed price transactions for seller experiments, and average sale price in each category for the category grouping). The solid lines repeat the same exercise, but are based on regressions that also include group (experiment or category) fixed effects.

Figure 8: Category heterogeneity

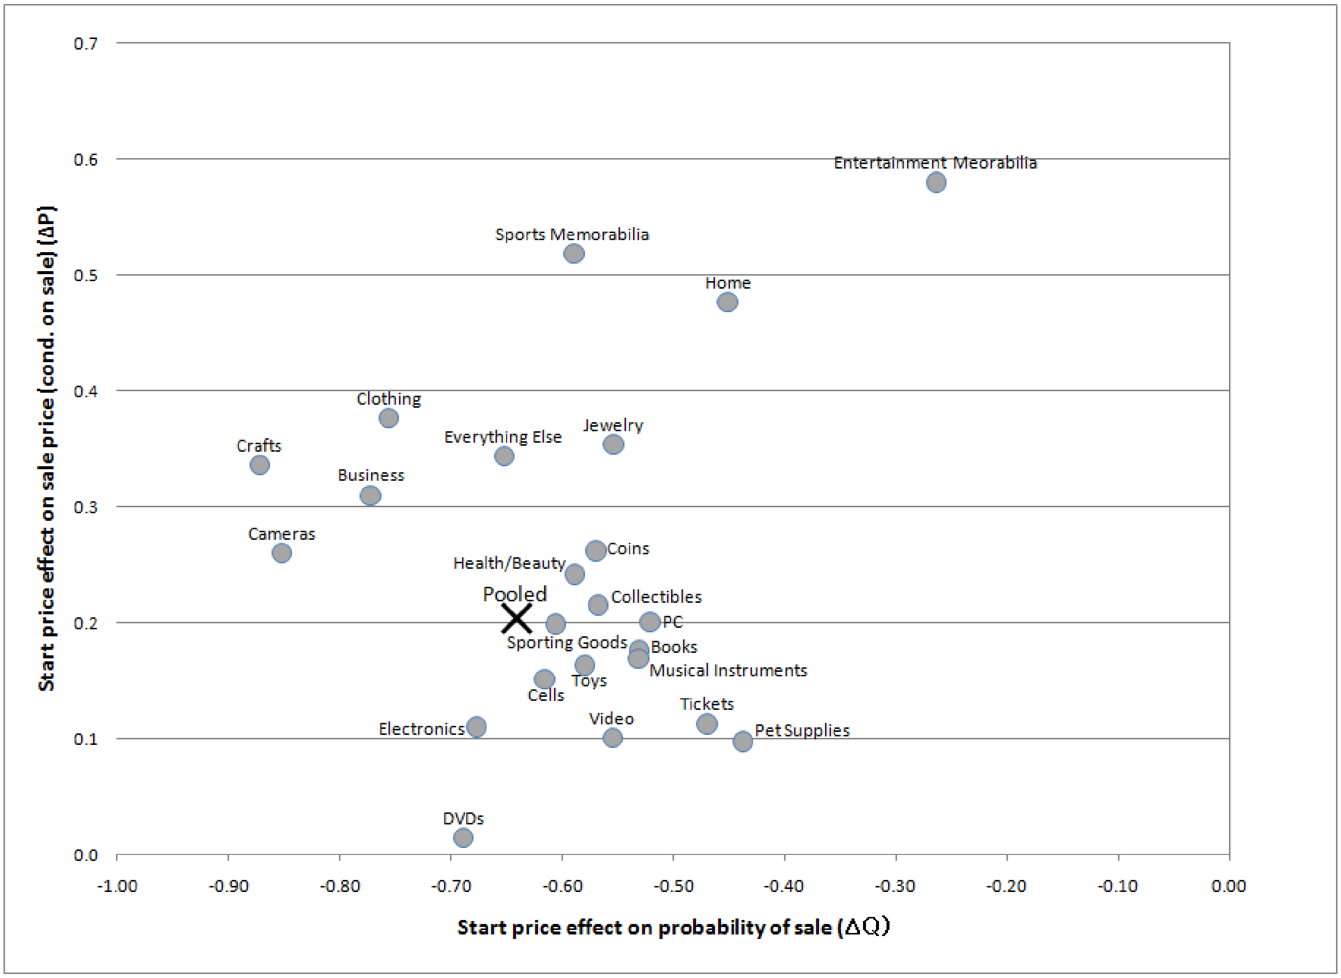

The figure presents the relationship between auction starting price and the probability of sale (horizontal axis) and transaction price (vertical axis) for different product categories, parallel to the regression results reported in Table 11. For each category, we run a simplified linear regression of the probability of sale on the (normalized) starting price p/v, and (separately) a regression of the transaction price (conditional on sale) on the starting price.

Table 1: Baseline data set

	Baseline Sample (1)					All Auction Exp. (2)	Random eBay (3)
	Obs. (millions)	Mean	Std. Dev.	25th pctile	75th pctile	Mean	Mean
Listings							
Start price ($)	7.69	42.47	194.48	5.45	20.89	26.96	27.90
Fraction with BIN option	7.69	0.73				0.29	0.24
BIN price ($) (if exists)	5.60	47.70	202.14	7	24	54.16	63.60
Fraction with secret reserve	7.69	0.006				0.006	0.009
Secret reserve price ($) (if exists)	0.05	355.23	605.45	99	354	323.69	322.39
Fraction with flat rate shipping	7.69	0.95				0.88	0.85
Fraction with free shipping	7.69	0.77				0.27	0.21
Shipping fee ($) (if flat and >0)	1.65	8.13	16.55	3.99	6.00	8.12	7.41
Auction duration (days)	7.69	3.2	2.5	1.0	7.0	4.5	5.6
Seller feedback score (000s)	7.69	327.0	472.1	4.6	308.0	24.40	26.6
Seller feedback (pct. positive)	7.65	99.3	2.0	98.9	99.8	99.36	97.5
Fraction with a catalog number	7.69	0.21				0.05	0.06
Fraction with associated:							
Fixed price listings	7.69	1.00				0.18	--
Fixed price transactions	7.69	1.00				0.13	--
Overlapping auctions	7.69	0.81				0.53	--
Most frequent category	Cell Phones, PDAs (24.2%)					Clothing (23.2%)	Clothing (18.8%)
2nd most frequent category	Video Games (19.5%)					Jewelry (14.9%)	Jewelry (11.9%)
3rd most frequent category	Electronics (13.1%)					Collectibles (7.7%)	Collectibles (10.8%)
4th most frequent category	Computers, Networking (6.4%)					Home + Garden (4.2%)	Toys + Hobbies (5.3%)
5th most frequent category	Cameras, Photo (5.3%)					Video Games (4.1%)	Sports mem, Cards (5.3%)
Fraction sold	7.69	0.35				0.27	0.39
Transactions							
Price ($)	2.69	67.39	172.95	8.50	73.01	32.29	38.22
Price including shipping ($)	2.69	69.54	174.96	8.99	76.00	37.18	43.55
Start price / sale price ratio	2.69	0.63	0.44	0.03	1.00	0.70	65.14
Number of bids	2.69	6.4	8.7	1.0	10.0	3.9	4.4
Number of unique bidders	2.69	3.6	3.9	1.0	6.0	2.4	2.7

A unit of observation is a listing. Column (1) presents statistics for the baseline sample. Column (2) presents statistics for all seller experiments (that is, including those for which we do not have a corresponding fixed price transaction). Column (3) presents statistics for the population of the entire eBay listings during the same period.

Table 2: Baseline data set

	Baseline Sample (1)					All Auction Experiments (2)				
	Obs. (000s)	Mean	Std. Dev.	25th pctile	75th pctile	Obs. (000s)	Mean	Std. Dev.	25th pctile	75th pctile
Number of (auction) listings	244.1	31.5	113.3	2	19	54,984.3	6.4	26.6	2	4
Fraction with positive sales	244.1	0.728				54,984.3	0.579			
Number of (auction) sales	244.1	11.0	49.5	0	7	54,984.3	1.8	10.1	0	1
Associated fixed price listings	244.1	6.9	22.6	1	6	4,047.4	4.4	16.4	1	4
Associated successful fixed price listings	244.1	2.9	6.6	1	3	4,047.4	1.3	4.2	0	1
Experiment "duration" (days)	244.1	56.2	72.4	8	77	54,984.3	38.2	57.9	7	42
Experiment sale rate	244.1	0.411	0.383	0.000	0.778	54,984.3	0.306	0.341	0.000	0.500
Experiment average sale price	177.6	101.41	303.64	10.21	89.00	31,854.0	42.75	165.24	7.83	31.00
Experiment median sale price	177.6	101.09	303.36	9.99	88.95	31,854.0	42.62	165.12	7.75	30.99

A unit of observation is a seller experiment. Column (1) presents statistics for the baseline sample. Column (2) presents statistics for all seller experiments (that is, including those for which we do not have a corresponding fixed price transaction).

Table 3: Within experiment variation

	Sample	Baseline sample	Large experiments (10+ listings)	Listings with start price below $1	Listings with free shipping	Listings without a BIN option	Listings without a secret reserve	Auctions that last (exactly) 7 days
	Total number of experiments	244,119	89,670	35,391	143,106	125,282	237,815	114,745
Within-experiment variation in:	Start price	142,653	79,107	17,350	82,423	62,148	139,526	57,045
	Shipping rate (flat rate only)	17,718	8,979	2,127		7,229	16,869	8,096
	Free shipping indicator	11,917	4,902	1,633		5,566	11,178	4,553
	BIN (any variation)	90,404	53,788	4,312	51,006		87,728	37,962
	BIN option indicator	24,052	9,754	2,383	13,154		22,788	8,487
	Secret reserve (any variation)	5,267	1,009	1,093	2,165	2,374		1,950
	Secret reserve indicator	2,918	652	386	1,215	1,264		1,036
	Auction duration	92,226	48,132	12,908	57,069	43,403	89,905	
	Day of week that auction ends	211,554	87,785	29,096	123,260	102,585	205,988	84,626

The table presents the extent of within experiment variation in the baseline sample Each entry in the table reports the number of experiments that contain within experiment variation in the listing parameter that is defined by the row header, out of the sample defined by the column header. The first column uses the entire baseline data, and the other columns stratify the baseline data based on various criteria.

Table 4: Summary statistics about price dispersion

	Baseline Sample (1)		All Auction Experiments (2)	
	Number of Experiments	Avg Coeff. of Price Var.	Number of Experiments	Avg Coeff. of Price Var.
All experiments (with 2+ sales)	143,942	0.11	13,548,775	0.15
Within same calendar month	125,124	0.10	16,427,575	0.13
With start price < $1	43,025	0.19	4,970,210	0.20
With start price >$1	104,548	0.07	8,556,050	0.12
With no BIN option	73,677	0.15	10,336,945	0.16
With BIN option	74,586	0.07	3,121,350	0.10
Experienced seller (feedback > 5,000)	68,696	0.08	3,939,100	0.14
Inexperienced seller (feedback < 250)	26,712	0.15	3,545,215	0.16
With any posted price listings	143,942	0.11	1,373,150	0.13
With posted price at ending time	91,178	0.10	564,060	0.11
Experiments in Specific Categories				
Clothing, Shoes, Accessories	20,586	0.06	631,135	0.13
Jewelry and Watches	10,612	0.13	4,814,770	0.13
Video games	13,579	0.09	759,635	0.13
Cell phones, PDAs	11,154	0.08	581,765	0.14
Electronics	6,926	0.14	3,001,105	0.18

The table presents summary statistics regarding price dispersion in the baseline sample (column (1)) and in the entire set of auction experiments (column (2)). Each grouping of listings cuts the data in different ways.

Table 5: Within and across experiment variation in auction start price

		Item reference value				
		< $10	$10-30	$30-100	$100-1,000	All listings
	Number of listings	92,925	184,652	125,326	91,267	494,170
Ratio of auction start price to reference value	< 0.05	6.5%	7.3%	20.3%	25.3%	13.8%
	0.05 to 0.15	6.7%	3.6%	0.5%	0.8%	2.9%
	0.15 to 0.30	5.3%	0.7%	1.5%	0.2%	1.7%
	0.30 to 0.45	2.1%	1.8%	2.2%	0.7%	1.7%
	0.45 to 0.60	5.5%	2.9%	3.5%	1.3%	3.2%
	0.60 to 0.85	12.9%	21.7%	17.4%	8.4%	16.5%
	0.85 to 1.00	42.1%	44.7%	37.0%	44.4%	42.2%
	1.00 to 1.20	11.5%	12.5%	13.8%	16.1%	13.3%
	> 1.20	7.3%	4.8%	3.8%	3.0%	4.7%

		Maximum (within experiment) ratio of auction start price to reference value									
		< 0.05	0.05 to 0.15	0.15 to 0.30	0.30 to 0.45	0.45 to 0.60	0.60 to 0.85	0.85 to 1.00	1.00 to 1.20	> 1.20	Total
Minimum (within experiment) ratio of auction start price to reference value	< 0.05	489	220	204	203	198	547	908	343	150	3,262
	0.05 to 0.15		52	95	75	151	290	337	57	44	1,101
	0.15 to 0.30			64	139	106	124	104	31	39	607
	0.30 to 0.45				48	187	219	104	31	43	632
	0.45 to 0.60					115	694	337	91	57	1,294
	0.60 to 0.85						1,218	2,784	637	300	4,939
	0.85 to 1.00							2,627	2,436	1,068	6,131
	1.00 to 1.20								550	667	1,217
	> 1.20									594	594
	Total	489	272	363	465	757	3,092	7,201	4,176	2,962	19,777

The table presents the distribution of (normalized) start prices, and the amount of variation within experiments, for the experiments we use to analyze the effect of auction start price.

Table 6: The effect of auction start price

	Item reference value							
	< $10		$10-30		$30-100		$100-1,000	
Dependent Variable: Sale indicator								
Start/value ratio indicator:								
0.05-0.15	-0.066	(0.013)	-0.042	(0.010)	-0.015	(0.022)	-0.086	(0.021)
0.15-0.30	-0.150	(0.011)	0.075	(0.019)	-0.086	(0.015)	-0.123	(0.039)
0.30-0.45	-0.273	(0.017)	-0.166	(0.012)	-0.171	(0.014)	-0.214	(0.028)
0.45-0.60	-0.416	(0.013)	-0.246	(0.010)	-0.193	(0.010)	-0.373	(0.015)
0.60-0.85	-0.522	(0.012)	-0.476	(0.007)	-0.421	(0.007)	-0.539	(0.008)
0.85-1.00	-0.645	(0.011)	-0.588	(0.007)	-0.597	(0.006)	-0.695	(0.006)
1.00-1.20	-0.674	(0.013)	-0.646	(0.008)	-0.648	(0.007)	-0.775	(0.007)
> 1.20	-0.721	(0.013)	-0.694	(0.010)	-0.760	(0.010)	-0.807	(0.012)
Constant	0.932	(0.010)	0.881	(0.007)	0.906	(0.005)	0.973	(0.004)
Number of listings	92,925		184,652		125,326		91,267	
Number of experiments	3,769		7,183		4,772		4,053	
Dependent Variable: Sale price (conditional on sale)								
Start/value ratio indicator:								
0.05-0.15	0.146	(0.036)	0.006	(0.006)	0.024	(0.013)	0.038	(0.007)
0.15-0.30	0.084	(0.034)	-0.043	(0.011)	-0.022	(0.009)	0.031	(0.014)
0.30-0.45	0.135	(0.050)	-0.038	(0.007)	-0.014	(0.009)	0.011	(0.011)
0.45-0.60	0.233	(0.039)	-0.008	(0.006)	-0.005	(0.007)	-0.050	(0.007)
0.60-0.85	0.255	(0.035)	0.045	(0.005)	0.039	(0.005)	0.032	(0.004)
0.85-1.00	0.413	(0.035)	0.185	(0.005)	0.150	(0.005)	0.118	(0.003)
1.00-1.20	0.533	(0.045)	0.323	(0.007)	0.273	(0.007)	0.208	(0.004)
> 1.20	0.762	(0.048)	0.608	(0.010)	0.500	(0.012)	0.544	(0.012)
Constant	0.610	(0.026)	0.769	(0.004)	0.817	(0.002)	0.855	(0.001)
Number of sales	39,174		72,067		60,375		42,285	
Number of experiments	3,010		5,889		3,762		2,831	

The table presents regression results of listing outcomes on (normalized) starting price, using experiment fixed effects. The dependent variable in the top panel is a dummy variable that is equal to one when the listing transacts. The dependent variable in the bottom panel is the transaction price (conditional on sale).

Table 7: Within and across experiment variation in BIN price

		Item reference value				All listings
		< $10	$10-30	$30-100	$100-1,000	
	Number of listings	15,277	11,360	65,041	32,079	123,757
Ratio of BIN price to reference value	No BIN	47.2%	42.9%	20.7%	28.2%	27.9%
	< 0.90	8.5%	6.0%	8.4%	15.8%	10.1%
	0.90 to 0.95	1.0%	2.5%	17.5%	17.9%	14.2%
	0.95 to 1.00	19.6%	16.2%	16.3%	13.2%	15.9%
	1.00 to 1.10	8.6%	10.9%	15.2%	13.4%	13.5%
	> 1.10	15.1%	21.5%	21.9%	11.5%	18.3%

		Maximum (within experiment) ratio of BIN price to reference value						
		No BIN	< 0.90	0.90 to 0.95	0.95 to 1.00	1.00 to 1.10	> 1.10	Total
Minimum (within experiment) ratio of BIN price to reference value	No BIN	0	108	55	522	440	648	1,773
	< 0.90		55	40	102	50	65	312
	0.90 to 0.95			18	52	59	33	162
	0.95 to 1.00				139	128	148	415
	1.00 to 1.10					140	134	274
	> 1.10						303	303
	Total	0	163	113	815	817	1,331	3,239

The table presents the distribution of (normalized) BIN prices, and the amount of variation in BIN prices within experiments, for the experiments we use to analyze the effect of BIN price.

Table 8: The effect of BIN price

	Value $10-30, No BIN, Starting price < $1		Value $30-100, No BIN, Starting price < $1		Value $100-1,000, No BIN, Starting price < $1	
Fraction sold	0.982		0.987		0.978	
Dependent Variable: Sale via BIN option indicator						
BIN price to value ratio indicator:						
< 0.90	(omitted)		(omitted)		(omitted)	
0.90-0.95	-0.086	(0.036)	-0.055	(0.009)	-0.020	(0.011)
0.95-1.00	-0.122	(0.028)	-0.074	(0.009)	-0.029	(0.013)
1.00-1.10	-0.165	(0.033)	-0.122	(0.011)	-0.096	(0.013)
> 1.10	-0.246	(0.036)	-0.240	(0.015)	-0.110	(0.017)
Constant	0.355	(0.026)	0.249	(0.009)	0.215	(0.009)
Number of listings	5,959		50,584		22,254	
Number of experiments	368		665		624	
Dependent Variable: Sale price (conditional on sale)						
BIN price to value ratio indicator:						
< 0.90	-0.102	(0.018)	-0.092	(0.004)	-0.113	(0.005)
0.90-0.95	-0.031	(0.022)	-0.049	(0.004)	-0.053	(0.005)
0.95-1.00	0.000	(0.009)	-0.007	(0.004)	-0.001	(0.004)
1.00-1.10	0.038	(0.012)	-0.003	(0.003)	0.011	(0.004)
> 1.10	0.083	(0.013)	0.012	(0.005)	0.046	(0.009)
Constant (No BIN)	0.825	(0.005)	0.850	(0.002)	0.883	(0.003)
Number of listings	11,013		64,012		31,200	
Number of experiments	662		1,026		908	

The table presents regression results of listing outcomes on (normalized) BIN price, using experiment fixed effects. The sample includes all items with reference value greater than $10 and only listings with starting price that is less than $1, so that virtually all items in the sample transact. The dependent variable in the top panel is a dummy variable that is equal to one when the listing transacts via the BIN price (rather than via the regular auction). The dependent in the bottom panel is the transaction price (via BIN or auction).

Table 9: Within and across experiment variation in shipping rate

		Item reference value				All listings
		< $10	$10-30	$30-100	$100-1,000	
	Number of listings	12,726	30,929	40,812	32,735	117,202
(Flat) Shipping rate	Free	26.5%	51.1%	37.9%	38.3%	40.3%
	0 to $2.50	19.9%	4.0%	1.2%	0.4%	3.7%
	$2.50 to $5	37.5%	22.5%	11.8%	2.7%	14.9%
	$5 to $10	11.0%	13.5%	24.0%	13.3%	16.8%
	$10 to $20	4.6%	6.9%	19.0%	26.2%	16.3%
	> $20	0.4%	1.8%	6.0%	19.3%	8.0%

		Maximum (within experiment) shipping rate					
		0 to $2.50	$2.50 to $5	$5 to $10	$10 to $20	> $20	Total
Minimum (within experiment) shipping rate	Free	385	1,277	995	519	315	3,491
	0 to $2.50	91	219	3	0	0	313
	$2.50 to $5		559	332	29	2	922
	$5 to $10			504	371	10	885
	$10 to $20				516	176	692
	> $20					352	352
	Total	476	2,055	1,834	1,435	855	6,655

The table uses the baseline sample, and shows the extent of variation in shipping fees. The top panel presents statistics on the variation in (dollar) shipping fees across experiments, while the bottom panel presents variation within experiments.

Table 10: The effect of shipping fees

	Baseline sample		Only listings with positive shipping rate		Value < $30 & Start price < $1		Value in $30-1,000 & Start price < $1	
Dependent Variable: Sale indicator								
Shipping > 0 (indicator)	-0.014	(0.0042)	--	--	-0.056	(0.0130)	-0.002	(0.0049)
Shipping fee ($)	-0.001	(0.0002)	-0.001	(0.0003)	-0.015	(0.0023)	-0.0003	(0.0003)
Constant	0.639	(0.0024)	0.621	(0.0037)	0.882	(0.0066)	0.959	(0.0025)
Number of listings	117,202		70,023		16,990		34,529	
Number of experiments	6,655		6,655		1,076		1,742	
Dependent Variable: Sale price (conditional on sale)								
Shipping > 0 (indicator)	-2.521	(0.3120)	--	--	-1.571	(0.2307)	-2.940	(0.5063)
Shipping fee ($)	0.181	(0.0202)	0.523	(0.0468)	0.362	(0.0440)	0.039	(0.0329)
Constant	93.734	(0.1576)	93.945	(0.5662)	16.398	(0.0858)	122.066	(0.2533)
Number of sales	73,034		43,064		13,403		42,335	
Number of experiments	5,156		4,679		847		2,624	

The table presents regression results of listing outcomes on (dollar) shipping fee, using experiment fixed effects. Column (1) reports results for the baseline sample, while the other columns cut the data in different ways. The dependent variable in the top panel is a dummy variable that is equal to one when the listing transacts. The dependent variable in the bottom panel is the transaction price (conditional on sale). Note that the transaction price includes the shipping fee, so in a frictionless market the coefficient on shipping fee should be zero.

Table 11: Category heterogeneity in the effect of auction starting price

Category	Experiments	Listings	Sales	Dep. Var. is Sale indicator		Dep. Var. is Sale Price (if sold)	
				Coeff.	Std. Err.	Coeff.	Std. Err.
Clothing, Shoes	2,505	24,351	7,692	-0.771	(0.030)	0.340	(0.046)
Jewelry + Watches	2,036	54,397	10,951	-0.586	(0.022)	0.344	(0.034)
Home + Garden	1,257	51,181	15,656	-0.518	(0.041)	0.424	(0.049)
Health, Beauty	1,148	38,367	19,536	-0.565	(0.060)	0.226	(0.049)
Cell Phones, PDAs	961	45,519	22,131	-0.619	(0.039)	0.149	(0.021)
Computers, Networking	928	17,134	10,000	-0.543	(0.056)	0.183	(0.022)
Electronics	836	29,076	19,705	-0.677	(0.040)	0.107	(0.022)
Sporting Goods	631	25,120	10,052	-0.660	(0.057)	0.196	(0.036)
Collectibles	609	9,113	4,008	-0.575	(0.072)	0.208	(0.074)
Video Games	605	12,885	9,076	-0.573	(0.055)	0.086	(0.020)
Sports Mem, Cards	556	7,187	1,653	-0.634	(0.047)	0.510	(0.120)
Everything Else	329	6,498	3,130	-0.651	(0.063)	0.306	(0.097)
Cameras, Photo	534	23,565	12,243	-0.854	(0.032)	0.259	(0.030)
Toys + Hobbies	475	7,693	4,462	-0.610	(0.034)	0.138	(0.034)
Coins + Paper Money	373	8,964	5,063	-0.564	(0.111)	0.264	(0.125)
Business & Industrial	352	7,088	2,765	-0.778	(0.067)	0.309	(0.041)
DVDs & Movies	329	6,388	4,844	-0.689	(0.076)	0.015	(0.052)
Books	249	1,695	713	-0.530	(0.138)	0.178	(0.056)
Crafts	165	4,814	2,173	-0.939	(0.091)	0.316	(0.070)
Tickets	162	597	216	-0.469	(0.090)	0.098	(0.117)
Pet Supplies	150	5,290	3,127	-0.440	(0.071)	0.091	(0.030)
Musical Instruments	121	2,667	982	-0.526	(0.116)	0.171	(0.026)
Entertainment Memorabilia	117	3,357	1,224	-0.263	(0.210)	0.582	(0.302)
Pooled				-0.641	(0.017)	0.204	(0.012)

The table illustrates the heterogeneity in the effects across categories, using regressions that are similar to those reported in Table 6. We report the effect of auction starting price on the probability of sale and transaction price (conditional on sale) for different product categories. For each category, we run a simplified linear regression of the probability of sale on the (normalized) starting price p/v, and (separately) a regression of the transaction price (conditional on sale) on the same starting price variable.