

NBER WORKING PAPER SERIES

LEARNING FROM SELLER EXPERIMENTS IN ONLINE MARKETS

Liran Einav  
Theresa Kuchler  
Jonathan D. Levin  
Neel Sundaresan

Working Paper 17385  
<http://www.nber.org/papers/w17385>

NATIONAL BUREAU OF ECONOMIC RESEARCH  
1050 Massachusetts Avenue  
Cambridge, MA 02138  
September 2011

We appreciate support from the National Science 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w17385.ack>

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Liran Einav, Theresa Kuchler, Jonathan D. Levin, and Neel Sundaresan. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Learning from Seller Experiments in Online Markets  
Liran Einav, Theresa Kuchler, Jonathan D. Levin, and Neel Sundaresan  
NBER Working Paper No. 17385  
September 2011  
JEL No. C93,D44,L13,L86

**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 show how this type of data can be used to address questions about consumer behavior and market outcomes, and provide illustrative results on price dispersion, the frequency of over-bidding, the choice of reserve prices, "buy now" options and other auction design parameters, and on consumer sensitivity to shipping fees. 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.

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

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

Theresa Kuchler  
Stanford University  
Department of Economics  
579 Serra Mall  
Stanford, CA 94305-6072  
tkuchler@stanford.edu

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

An online appendix is available at:  
<http://www.nber.org/data-appendix/w17385>

# 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). 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, more than half will reappear on the site again as a separate listing, often 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 show how the targeted variation created by seller experimentation can address a range of old and new questions about consumer behavior and auction design in internet markets. In particular, we use the data to quantify the price variability or dispersion for identical listings, to evaluate the hypothesis that eBay consumers engage in “excessive” bidding (e.g., Lee and Malmendier, 2011), to measure the effect of auction reserve prices (Kamins et al., 2004; Ku et al., 2006; Reiley, 2006; Lucking-Reiley et al., 2007; Simonsohn and Ariely, 2008), to analyze the impact of “buy now” options in consumer auctions (Stadifird et al., 2004; Akerberg et al., 2006; Anderson et al., 2008), and to assess whether consumers systematically underweight shipping fees (Tyan, 2005; Hossain and Morgan, 2006; Brown, Hossain and Morgan, 2010). In each case, we provide evidence based on thousands of distinct experiments across a wide range of product categories.

Our findings sharpen, enrich and in some cases overturn earlier results that have been obtained in observational or experimental studies of particular items. For instance, we find substantial price variation in auctions for identical items conducted by the same seller, even when the auctions are held concurrently. But we observe relatively few instances of obvious overbidding, that is cases where a bidder pays more at auction than a concurrent posted price for the same item offered by the same seller. We find clear and consistent relationships between auction reserve prices, sale probabilities and closing prices, and similarly between “buy now” prices and sale outcomes. These relationships are surprisingly stable across diverse products and categories, and provide useful evidence on market demand in a large retail marketplace. We also confirm earlier findings that certain prices, such as shipping fees, are not fully internalized by buyers.

Apart from the specific findings, we view the empirical strategy as interesting in its own right. Since the early days of the internet, it has been clear that the vast and detailed data being collected in online markets would provide opportunities to study consumer behavior, to test theories of competition and market structure, and to analyze the effects of changes in search costs, product variety, and market organization, all in relatively structured environments. The difficulty has been to find empirical approaches that yield plausible identification of parameters of interest across a wide range of settings. In principle, the scale and diversity of many internet markets should be ideal for this purpose; in practice, it has not always been easy to leverage these advantages.

The two main approaches to studying internet markets such as eBay have been observational studies that relate sales outcomes such as auction price to differences in sale parameters such as reserve prices, shipping fees or seller reputation (e.g., Bajari and Hortacsu, 2004), and field experiments in which a researcher sells identical items under different conditions (e.g., Lucking-Reiley, 1999). Most observational studies focus on a narrow product category — a particular type of coin or trading card, or a specific laptop or electronics component — but even then can have a difficult time controlling for confounding differences between listings, such as the identity of the seller or the exact quality of the item. Field experiments help address these confounding issues, but can have the drawback of small sample size. More importantly, both approaches, even when executed expertly, can lead to narrow results that

may be specific to a particular item or time window. Because one of the most interesting aspects of markets like eBay is their scale and diversity, an empirical approach that entirely strips away these features, rather than exploiting them, is somewhat disappointing.

Several studies take an approach that is closer to the one we propose. Ostrovsky and Schwarz (2010) 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 (2010) define seller experiments in a way that is almost the same as in this paper in studying the effect of charity contributions by eBay sellers, although they do not remark on either the prevalence of duplicate listings or the opportunity for using them as a broader research tool.

The remainder of the paper proceeds as follows. Section 2 describes the salient features of eBay’s platform, the use of duplicate listings and “experiments” by retail sellers, our data construction, and summary statistics. Section 3 uses the experiments data to analyze the issues described above: price variability, excessive bidding, reserve and “buy now” prices, and shipping fees. Section 4 compares the use of seller experiments with alternative empirical approaches; we illustrate the possible endogeneity that can confound observational studies, as well as the substantial heterogeneity across product categories that can threaten the external validity of narrow studies. In Section 5 we conclude by discussing various explanations for why sellers vary sale parameters and engage in the type of experimentation we exploit. A lengthy appendix provides many additional analyses, addressing potential concerns about our specific definition of a seller experiment used in the paper. We replicate all the results using a range of samples and matching approaches, showing that the results are highly consistent across alternative specifications and definitions of an experiment.

## **2 Background, Data, and Empirical Strategy**

### **2.1 Background and Empirical Challenge**

Our analysis focuses on the e-commerce platform eBay, which in 2009 (the year of our data) had approximately ninety million active users and \$57 billion in gross merchandise volume.

The eBay marketplace includes large and active sub-markets for collectibles, electronics, clothes, cars, tickets, toys, books, jewelry and art, both new and used. Products are offered by both consumer sellers and professional retailers. The size and diversity of the platform, and the relative ease of gathering data and running field experiments, has captured significant attention from researchers, who have used eBay both as a prototype for large online markets, and as a laboratory to investigate traditional questions about consumer behavior, competition and market design.<sup>1</sup>

A key challenge in this work, and more generally in using the data becoming available from other large and diverse online markets, is heterogeneity. For example, the items sold on eBay range from used cars offered for tens of thousands of dollars to “Silly Bandz” offered for less than two dollars. Some sellers do millions of dollars in annual sales; others sell just one item. Primarily for this reason, researchers often have focused on very narrowly defined product categories, such as specific pop-music CDs (Nekipelov, 2007), collectible coins (Bajari and Hortacsu, 2003), Pokemon cards (Katkar and Reiley, 2006), or board games (Lee and Malmendier, 2011).

Sellers also have considerable flexibility and low cost in varying their sales strategy. Among other things, sellers on eBay can choose their listing title and thumbnail picture, a longer item description for consumers to study after clicking on their listing, and their sales mechanism. Traditionally, most sellers have used ascending auctions. For an auction sale, the seller can choose the auction duration, the start price, whether to specify an additional secret reserve price, and whether to offer a “Buy It Now” (BIN) feature that allows bidders to preempt the auction and purchase the item at a fixed price before an initial bid is made. Nowadays, many sellers also offer items for sale at regular posted prices; indeed, posted price transactions account for more than half of eBay’s sales volume. It is easy for sellers to vary all these parameters as well as others such as the shipping fee.

As an illustration, Figure 1(a) shows the eBay listings displayed following a search for “taylormade driver” (a type of golf club).<sup>2</sup> Within this narrow product category, it is already

---

<sup>1</sup>Bajari and Hortacsu (2004) and Hasker and Sickles (2010) review dozens of papers using data from eBay. Ambrus and Burns (2010) provides a recent state-of-the-art theory of rational bidding behavior.

<sup>2</sup>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

apparent that listings vary widely. The products themselves are differentiated (men’s clubs, women’s clubs, different models, sizes, new and used), as are the sellers (some are “top-rated” for instance) and the sales mechanisms (posted prices, auctions, BIN auctions, all with different end times and current prices), and shipping arrangements and fees. Even a quick perusal of the figure should make it clear that attributing patterns in the data to specific strategies or choices of sales mechanisms, even for a narrowly defined set of products, is a challenge.

To circumvent this concern, our empirical strategy is to identify cases in which particular sellers offering particular goods “experiment” by varying their pricing or sales strategies. This means focusing on more professional retailers who are using the platform to sell new consumer products, rather than on one-time sellers or those offering unique or used goods. Retailers who want to sell multiple quantities of an item have several listing options. A common approach is to run auctions with staggered end dates, posting additional auctions as old auctions expire. It is also possible, but highly unusual, to run multi-unit auctions. Sellers using fixed prices often post multi-unit listings that can be renewed over time, although some place multiple simultaneous listings. For instance, Figure 1(b) shows listings from a seller (with the user name *budgetgolfer*) who on September 12, 2010 had posted 31 listings for a particular TaylorMade driver. Of these listings, 20 were auctions that were scheduled to end in the next week, and 11 offered the driver for a fixed price of \$124.99.

Apart from offering different sales mechanisms, sellers such as *budgetgolfer* may vary pricing parameters across concurrent sales or over time. For example, in Figure 1(b) the displayed listings have two different shipping fees (either \$7.99 or \$9.99). Our analysis relies on the fact that sellers frequently take advantage of this ability to vary prices of fees, often by significant amounts, creating variation that can illuminate features of market demand and consumer behavior.

## 2.2 Experiments data

We construct our data from the universe of eBay.com listings in 2009, excluding only auto and real estate listings which have a somewhat different institutional structure. We look for 

---

more detailed item information, place bids, or make purchases.

sellers offering particular products who post multiple listings during the year, potentially varying their pricing or sales strategy. Because most eBay listings do not include a well-defined product code such as a ISBN or SKU number, we use the listing title and subtitle to identify products.

Specifically we identify all sets of eBay listings that have an exact match on four variables: seller identification number, item category, item title and subtitle. For example, the listings in Figure 1(b) are in the same set, along with other identically matched listings that were active before or after the day of the screenshot.<sup>3</sup> We drop single listings that have no match, which still leaves us with over 350 million listings, grouped into 55 million matched sets. We refer to each set as a seller experiment.

Our empirical strategy relies on variation within experiments in sale parameters and outcomes. In this paper, we focus on auctions (rather than posted price listings), which leads us to refine the data in several ways. In particular, we restrict attention to experiments that include at least two auction listings, at least one successful posted price listing, and for which the listings have a non-empty subtitle. The first restriction is necessary to have within-experiment comparisons. The second, as we explain below, provides a useful way to normalize prices in order to make experiments comparable and compute average treatment effects. Finally the third restriction allows us to reduce the size of the dataset to make it manageable, while focusing on more professional retailers who tend to use subtitles.

Selecting experiments according to these criteria leads to our baseline dataset: 244,119 experiments with a total of 7,691,273 listings. 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. Sold items obtain, on average, 6.4 bids from 3.6 unique bidders, and about a third of sales are made via a “Buy It Now” (BIN) option.

---

<sup>3</sup>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.



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 in the baseline data that are “catalogued,” the experience of the sellers, and their tendency to use “sophisticated” sale strategies such as a 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 selling production goods, while eBay as a whole also includes a vibrant consumer-to-consumer market.

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

The analysis below uses variation in sale parameters within experiments to estimate their effect on auction outcomes, mainly the probability of successful sale and the expected price, and in some cases the price distribution. The strategy relies on the fact that within matched sets of listings, what we call experiments, sellers do indeed vary the available sale parameters. Table 3 shows that, in fact, the amount of variation is dramatic. Specifically, Table 3 reports the number of experiments that contain variation in each of several different sale parameters of interest. As can be seen in the first column, 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 of Table 3 show that we can find substantial numbers of experiments with variation in a given sale parameter even if we condition on other sale parameters being held fixed. This is reassuring as it suggests that we may be able to isolate treatment effects of one variable of interest at a time. We return to this in the context of specific analyses below. For now the main take-away is that even as we restrict the sample to pinpoint

specific types of variation, the scale of the data means that we can still find thousands of experiments that are potentially informative.

## 2.3 Empirical Strategy

Our empirical analysis relies mainly on 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:

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

where  $\alpha_i$  is an experiment fixed effect and  $\varepsilon_{it}$  is an error term assumed to be mean-independent of  $z_{it}$  within experiments.

In principle, we could estimate separate treatment effects for each experiment, but there are at least two reasons to pool experiments as in our specification. First, the size of most experiments is small, so pooling provides much greater statistical power. Second, it seems more interesting to estimate an average treatment effect across a large group of experiments rather than thousands of distinct effects for individual items. That being said, we break out estimates by item value, and in Section 4, discuss heterogeneity across item categories.

We rely on three assumptions to estimate average treatment effects. The first, which we maintain throughout the paper, is that the idiosyncratic effects of each experiment denoted by  $\alpha_i$  enter in an additive and separable way. The second, which is a bit more subtle, is a choice of price units. The experiments in our data involve items of very different value. If  $z_{it}$  is the auction reserve price, a one dollar change might have a large effect for a \$5 item, but little effect for a \$500 item. We address this by creating a reference value for each item, and using it to normalize item prices.

Specifically, we define each item’s reference value  $v_i$  to be the average price across posted price transactions of that item.<sup>4</sup> Then when we analyze the effect of, say, the auction reserve

---

<sup>4</sup>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.

price, we work with the normalized reserve price  $z_{it}^n = z_{it}/v_i$  rather than the dollar reserve price  $z_{it}$ . Similarly, when we analyze the effects of auction parameters on the final auction price  $p_{it}$ , we work with the normalized price  $p_{it}^n = p_{it}/v_i$ . A more general alternative would be to estimate treatment effects of the form  $f(z_{it}, v_i)$  rather than  $f(z_{it}/v_i)$  but a rather remarkable finding below is that there seems to be little gain from doing this.

Our final assumption is that sale parameters within each experiment are not correlated with factors that bear directly on auction outcomes. Here one might raise several concerns. We define an experiment to be all listings of a given item by a given seller over a period of up to a year. If for example, demand conditions vary over the year (e.g. due to seasonality or shifts in demand for the item) and sellers respond to this variation, our estimates would suffer from standard endogeneity bias. Alternatively, if the seller had only one copy of the item, but listed multiple times because her initial listings didn't sell, our estimates might be tainted by forms of selection bias.

A natural way to address these concerns and many others is to vary the definition of an experiment. For instance, to address the former concern, one can re-define an experiment to involve a narrow time window or even concurrent listings. To address the latter, one can restrict attention to experiments with a large number of listings. A virtue of our general approach is that it is straightforward to replicate our estimates using alternative specifications or definitions of an experiment that might reveal biases in the main results. Because the number of possible exercises is infinite, while reader attention is not, we have written a long appendix that replicates all the results in the paper for a range of different samples, specifications and experiment definitions. The results remain strikingly similar across them all. For this reason, the main text proceeds in straight-ahead fashion, and we invite interested readers to inspect 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 a primary goal is to illustrate the potential power and scope of our approach, we intentionally present several different

exercises, each relatively briefly. For the same reason, our choice of topics is skewed toward those previously studied with eBay data, allowing us to compare and contrast our results. In ongoing work (Einav et al., 2011), we hope to 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”**

One of the initial hypotheses about internet commerce was that low search costs should lead to low price dispersion. Subsequent studies, however, including Bailey (1998), Brynjolfsson and Smith (2001), Baye, Morgan and Scholten (2004) and Ellison and Ellison (2009), have found substantial dispersion even in structured price comparison settings. And recent work by Lee and Malmendier (2011) provides even more striking evidence on consumer search. They present 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. Lee and Malmendier 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 behavior and that a significant number of consumers are irrationally over-bidding.

A complicating factor in existing studies is that prices are compared across retailers. This makes it difficult to disentangle differences in retailer attractiveness from frictions in consumer search. Our experiments approach allows us to estimate price variability across auctions by a single seller, eliminating variation driven by seller differences. Moreover, we can construct estimates for hundreds of thousands of items across a range of categories and for narrower or wider time windows. In addition, because a substantial fraction of the auctions in our data take place in the presence of a concurrent posted price, we can examine the Lee and Malmendier over-bidding hypothesis, expanding from a single item sold by heterogeneous sellers to hundreds of thousands of items sold by fixed sellers.

The coefficient of variation, or the standard deviation of auction prices divided by the mean price, provides a basic measure of price dispersion across a set of sales. We compute it for each experiment, and for a finer partition that divides the listings in each experiment by

calendar month. Table 4 shows, regardless of whether we look at broader or narrower time windows, that there is considerable dispersion in auction sale prices holding both the item and seller fixed. The average coefficient of price variation across experiments is 0.11 (0.10 with the finer partition of each experiment). If we broaden our sample to include groups of duplicate auction listings with no matching fixed price sale, the average climbs to 0.15. There is notably less variability for experienced sellers, or when the seller uses a BIN option or a higher reserve price. For the broader sample of experiments, there is also a surprising consistency in price dispersion across product categories.

Next, we compare observed auction prices in our data to posted prices of the same object. 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.

Looking across all the auctions, 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 observe occasions when the auction price exceeds the reference price, as in Lee and Malmendier (2011). But the frequency of this happening is relatively low. Less than 20 percent of auctions sold for above the posted price, and most of these “excess bidding” 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 non-trivial 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. (2011).

## **3.2 Auction Design Parameters**

Sellers offering goods by auction must decide on the duration of the auction, the ending time, the reserve price, whether the reserve price is known to bidders, and in the specific case of eBay, whether to offer buyers the option to “buy now” at a price that disappears when the first bid is made. Although auction theory and targeted empirical studies have had a considerable amount to say about all these aspects of auction design, the use of “seller experiments” provides an opportunity to analyze the effect of auction design parameters in a controlled fashion, in large data samples, across a wide range of items.

### **3.2.1 Auction Starting Price**

Variation in auction start prices (or reserve prices) offers a chance to test some basic principles of auction theory. In standard private value auction models, 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  $s'$ . In contrast, models that allow endogenous entry or common values

can allow more nuanced effects of reserve prices.

Past studies of eBay reserve prices have reached different conclusions about their effect. Ku et al. (2006) found that in contrast to the standard auction theory model, lower start prices increased both the odds of sale and the price conditional on sale. Their explanation is based on escalating commitment. Lower start prices attract buyers who become committed to the auction and continue to bid aggressively as the price rises. Simonsohn and Ariely (2008) found that while lower start prices did not necessarily increase the price conditional on sale, they increased the price conditional on it rising above the higher start price – again consistent with the “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.

To study the effect of auction start prices, we look for experiments with variation in the start price. There are 142,653 such experiments in our baseline sample. To limit the variation in other auction parameters, we reduce the sample to focus on 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. We maintain our strategy of normalizing both start and sale prices by the item’s value, which facilitates comparison across experiments.

Table 5 shows the variation in start prices in the data. The top panel presents the overall variation in (normalized) start prices by price category. The bottom panel summarizes the within-experiment price variation. For the latter, we find the minimum and maximum (normalized) start price for each experiment, and cross-tabulate the experiments according to these numbers. There is stunning variation in the reserve prices across auctions of the same item. For instance, of the 3,262 experiments that contain at least one very low start price ( $p_{it}^n = p_{it}/v_i < 0.05$ ), 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  $p_{it}^n > 1$ . So not only can we find thousands of experiments with some start price variation, many sellers vary their start price dramatically.

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. The top 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 across price categories, suggesting an interesting scaling property: the probability of sale appears to depend largely on the start price relative to the item’s value, and not so much on the value or start price per se.

The second panel of Figure 4 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, and again the (normalized) price curves are remarkably similar across price categories. For start prices that are below 0.6 as a fraction of the item’s posted price value, the expected auction price conditional on sale is generally around 0.8. The flat price curve for low start prices suggests that competition among buyers keeps auction prices from slipping very far on average even if the start price is very low. For start prices above 0.6, expected prices conditional on sale are higher (although the probability of sale is also lower, per the results in the top panel).

The most obvious explanation for why the expected sales price increases at higher reserve price levels is mechanical. When the reserve price is low, many auctions may result in prices of 70-90% of the posted price value. If the start price is already at the upper end of this range, the sale price distribution is truncated, eliminating low price sales. Recall that in the standard private value auction model, however, increasing the start price from  $s$  to  $s'$  does not affect the distribution of auction prices above  $s'$ . Alternative theories of “escalating commitment” or “bidding frenzies” suggest that in contrast a higher start price may reduce the probability of very high sale prices. To assess this, Figure 4(c) plots for various start price levels the (unconditional) probability of the auction ending at a (normalized) price of at least  $x$ , for selected values of  $x$ . That is, it shows the upper tail of the price distribution for different start price levels. Interestingly, the upper tail of auction prices is heavier, and very high auction prices more likely, when sellers use a very low start price rather than a



moderate or moderately high one.

### 3.2.2 “Buy It Now” Option

A novel feature of eBay’s marketplace is that for a small fee of 5 to 25 cents, sellers can enrich a standard auction by offering a “Buy It Now” (BIN) option. The BIN option is a posted price at which a buyer can short-circuit the auction with an immediate purchase; it disappears as soon as the item receives a qualified bid, in which case a standard auction ensues. In a standard private value auction model with symmetric bidders, sellers would not want to offer such an option, because a lower value bidder can preempt one with higher value, reducing the expected price. In practice, however, a BIN option may attract bidders who value immediate purchase, or prompt early bidding that escalates. Studies by Stadifird et al. (2004), Akerberg et al. (2006), and Anderson et al. (2008) suggest that offering a BIN option on eBay can increase seller revenue. For example, Akerberg et al. (2006) analyze Dell laptop auctions and find that using a BIN option increases expected seller revenue by \$29.

The seller experiment approach allows us to provide large-scale evidence that illuminates the different stories. We start by identifying 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 the effects of the BIN option 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, with a total of 123,757 listings, or on average 38 listings per experiment.

Table 7 shows the variation in the BIN price, and parallels Table 5. We again normalize both sale and BIN prices by item value to facilitate comparison across experiments. The top panel of Table 7 shows the overall variation in (normalized) BIN price for each price category. The bottom panel summarizes the within-experiment BIN price variation. Again, we observe substantial variation in the (normalized) BIN price, both within and across experiments. We use this variation to estimate the effect of BIN prices, using an analogous strategy to the one employed in the previous section.

Table 8 and Figure 5 report the results. Because we have focused attention on auctions with essentially non-binding start prices, almost all listings (98 percent) in the resulting data set end in a successful sale. Therefore, we put aside the question of whether the item sells and instead ask whether it sells via the BIN option or via the standard auction mechanism. The top panels of Table 8 and Figure 5 report the results. The probability the items sells at the BIN price is relatively low, and falls further at high BIN prices. The second panel shows how the BIN price affects the final selling price (regardless of whether the item was sold via the BIN option or standard auction). The plot allows us to somewhat refine earlier findings because it appears that relatively low BIN prices in fact reduce seller revenue, while BIN prices that are similar to the item value have relatively little effect, and only “over-priced” BIN options generate incremental revenue for the seller.

A more subtle question is whether offering a BIN option at a given BIN price affects the subsequent sequence of bids (and the sale price) in cases where the BIN option is not exercised. Such an effect might be expected if the presence of a BIN price helps anchor buyer beliefs about an item’s value. In the bottom panel of Figure 5, we show for different BIN levels, the probability of obtaining a sale price below certain thresholds, that is, the lower tail of the cumulative price distribution. The likelihood of receiving a very low price below 70% of the reference value is nearly the same for both high and low BIN prices as it is with no BIN price. There is a modestly positive, although not very dramatic, effect of offering a high BIN price in terms of reducing the likelihood of a sale in the 70-100% range, consistent with the overall higher revenue from these sales.

### **3.2.3 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.

### 3.3 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 frequently 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. Seller's can also have another reason to favor shipping fees: until recently, eBay commissions (currently 9% for auction listings) 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 auction 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.<sup>5</sup> The resulting data contains 117,202 listings grouped into 6,655 experiments, with an average of 18 listings per experiment. A large fraction of these listings offer free shipping (a feature that is encouraged by eBay). The top panel of Table 9 presents the distribution of shipping rates across these listings, and the bottom panel of Table 9 presents the within-experiment variation in shipping fees. In parallel with our earlier analyses, we see sellers varying their choice of shipping fee considerably for given items.

Table 10 reports results of sale probability and transaction price (conditional on sale) for several different subsamples. As the top panel of the table shows, the effect of shipping fees on the probability of sale is minimal, so we focus our attention on the price effect, which we analyze in the bottom panel of the table. Unlike our earlier analysis, we choose to run the price regression without normalizing by the item value, as this helps in facilitating the quantitative interpretation of the estimated effects. Using this specification, a coefficient of

---

<sup>5</sup>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.

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 overall effect on total price (sale price plus shipping) is zero. As Table 10 indicates, our estimates suggest a positive coefficient with a magnitude of about 0.2-0.3, suggesting that only 70-80 percent of the shipping fee is translated into final price.

In addition, we find a distinct effect at zero. Free shipping is associated with an average price increase of about \$3, with the effect being greater (in absolute terms) for more expensive items. This distinct effect of free shipping is likely to be a combination of buyers' response to shipping being free (Shimpinier et al., 2007) as well as to eBay's strategy to encourage free shipping by prioritizing such listings 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.

An alternative way to get at the effect of shipping fees is to exploit the trade-off between an ex-ante start price and an ex-post shipping fee. A textbook economic analysis would suggest an equivalence. Selling an item with free shipping and a start price of \$10 should be identical to starting the auction at zero and charging a \$10 shipping fee. To investigate whether this obtains in practice, 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 into shipping fee and start price 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). The conclusions from this exercise are qualitatively similar to what we have already found: increases in shipping fees 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 the many other creative strategies that economists use to estimate treatment effects? The two most obvious alternatives to our approach 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.<sup>6</sup> 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 and item photograph 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.

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

---

<sup>6</sup>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.

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 intermediate start prices.<sup>7</sup>

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

---

<sup>7</sup>This potentially makes sense because as we point out in Einav et al. (2011), either very low or relatively high start prices appear to be desirable from the standpoint of optimal pricing, so the sellers who use these types of start prices more often may be “better” sellers whose items are more likely to sell.

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. So 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, in this way avoiding many of the identification problems inherent in large observational studies, but also the narrowness, small sample sizes and limited scope of many field experiments.

To illustrate, we used the approach to revisit a number of questions about consumer behavior and sales strategies that have been investigated in the literature, sharpening, expanding and in a few cases contradicting prior results. With the access to the relevant data, our empirical approach is easy to replicate and straightforward to implement. Thus, one can easily envision using it for a more in-depth analysis of each of the questions we analyze in this paper, or for dozens of other questions that come up in internet retail, advertising or labor markets. Our own view is that the “measurement” nature of our approach makes it perhaps best as a complement to other approaches, either theoretical modeling to derive hypotheses, or alternative empirical approaches that incorporate compositional effects or attempt to relate findings to specific economic models or primitives. This is the approach we are taking in ongoing work.

We conclude with a short discussion of why sellers engage in the type of behavior or “experimentation” that we exploit. There are several non-exclusive explanations, all of which relate to the relative ease with which participants in internet markets can adjust their strategies. On eBay specifically, the platform is set up so that sellers of multiple units (e.g., retailers of new goods) who want to use an auction format generally post multiple staggered listings, rather than holding a single large auction. The costs of adjusting listing details, reserve prices, buy now options, shipping fees and other sale parameters are negligible. Sellers may choose to make these adjustments in response to demand changes, because they are unsure about the best strategy, or in an attempt to segment consumers (e.g. target consumers who want to “buy now” versus those willing to bid in an auction). They may also post the same item in different ways simply in the hope of “crowding the shelves”

and capturing buyer attention.<sup>8</sup> Particularly interesting from an economic viewpoint is the possibility that sellers engage in active experimentation to improve their business practices. While addressing this possibility and its implications goes beyond the scope of this paper, we are continuing to explore it.

## 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 (2010), “Gradual Bidding in eBay-Like Auctions,” Mimeo, Harvard 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).
- 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.

---

<sup>8</sup>Note that a few of these rationales, in particular the possibility that sellers adjust sale strategies in response to demand changes, raise concerns about selection bias in our results. However, as discussed earlier, these concerns are relatively straightforward to address by focusing on subsets of the experiments in our data or finer groupings of listings. We show in the Appendix that essentially all the results reported in the paper remain consistent across a range of different approaches designed to reveal selection biases.

- 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 Sundareshan (2011). "Auctions and Posted Prices as Sales Mechanisms," Work in progress.
- 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.
- Lee, Young H., and Ulrike Malmendier (2011). "The Bidder's Curse." *American Economic Review*, forthcoming.
- 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.
- Nekipelov, Denis (2007). "Entry Deterrence and Learning Prevention on eBay." Mimeo, UC Berkeley.
- 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.” *Marketing 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.

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

The screenshot shows an eBay search results page for 'taylormade driver'. The page features a navigation bar with categories like FASHION, MOTORS, DEALS, and CLASSIFIEDS. A search bar at the top contains the query 'taylormade driver' and a 'Search' button. Below the search bar, there are related searches and a filter for 'Golf'. The main content area displays a list of 2,594 items found for 'taylormade driver' in the 'Golf Clubs: Right-Handed' category. The results are sorted by 'Best Match' and shown in a grid format. Each listing includes a product image, a title, a price, and a 'Buy It Now' button. The left sidebar contains various filters such as 'Dexterity', 'Shaft Material', 'Loft', 'Condition', 'Price', 'Flex', 'Club Type', 'Brand', 'Gender', 'Seller', 'Categories', and 'Buying formats'. The top right corner shows the user's account information and the eBay Buyer Protection logo.

Product Image	Product Title	Bids	Price	Time Left
	TaylorMade Burner 460 Driver 9.5* RE-AX 50/Stiff Shaft	0 Bids	\$70.00	5h 18m
	NEW TaylorMade Golf Clubs 2009 Burner Driver 10.5* Reg	Buy It Now	\$144.99	16d 16h 26m
	2010 LADIES TAYLORMADE BURNER SUPERFAST DRIVER 10.5 NEW	13 Bids	\$153.50	6h 34m
	TAYLORMADE GOLF R9 TP 10.5* DRIVER REGULAR VERY GOOD	Buy It Now or Best Offer	\$159.99	23d 20h 44m
	2007 LADIES TAYLORMADE BURNER DRIVER HT LADIES FLEX NEW	1 Bid	\$89.00	6h 43m
	New TaylorMade Burner Superfast TP Driver 9.5 F1 85 R	Buy It Now	\$189.99	27d 15h 42m
	NEW TaylorMade Burner SuperFast Driver (10.5/Reg/Graph)	0 Bids	\$179.00	6h 54m
	TaylorMade R9 Superdeep TP Driver Golf Club	Buy It Now or Best Offer	\$299.00	1d 7h 4m
	NEW!TAYLOR MADE R9 SUPERDEEP TP DRIVER 10.5 STIFF/SUPER	6 Bids	\$203.86	7h 7m
	TaylorMade Burner Superfast Driver Golf Stiff 9.5 deg	Buy It Now or Best Offer	\$199.99	15d 9h 3m

The figure presents a “standard” screenshot of search results on eBay. This particular screenshot is the result of searching for “taylormade driver” (a type of golf equipment) on 9/12/2010.

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

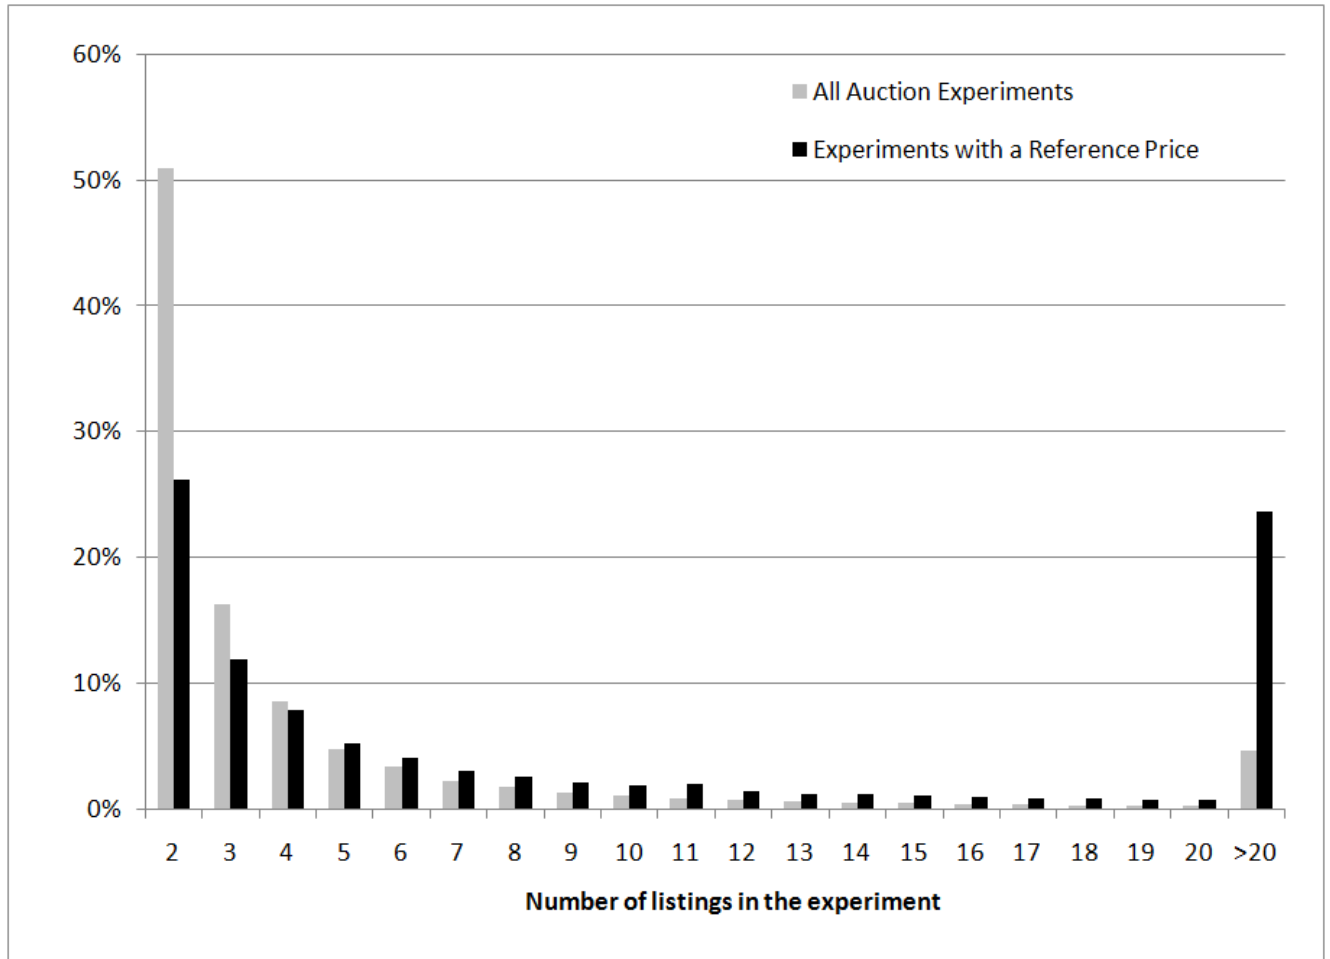
The screenshot shows an eBay search results page for the query "TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5\* RH". The page displays 31 results from a single seller, who is a "Top-rated seller". The results are sorted by "Time: ending soonest". The listings are organized into a table with columns for item description, seller status, bidding information, price, shipping, and time left.

Item Description	Seller	Bidding	Price	Shipping	Time Left
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	7 Bids	\$104.83	+\$9.99 shipping	3h 21m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	9 Bids	\$81.00	+\$9.99 shipping	3h 22m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	Buy It Now	\$124.99	+\$9.99 shipping	6h 9m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	Buy It Now	\$124.99	+\$7.99 shipping	6h 9m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	7 Bids	\$70.01	+\$9.99 shipping	8h 59m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	Buy It Now	\$124.99	+\$9.99 shipping	9h 41m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	Buy It Now	\$124.99	+\$7.99 shipping	9h 41m
TAYLORMADE BURNER 09 DRIVER 2009 GOLF CLUB NEW 10.5* RH	Top-rated seller	12 Bids	\$91.00	+\$7.99 shipping	11h 7m

The left sidebar contains a "Refine search" section with filters for Flex (Regular, Stiff), Condition (New, Used, Not Specified), Price (range selector), Club Type, Brand, Dexterity, Seller, and Categories (Sporting Goods, Golf). There is also a "Preferences" section for buying formats (Auction, Buy It Now).

The figure illustrates a “seller experiment.” As in Figure 1(a), it presents a screenshot of search results on eBay, but for a much more specific golf driver query that returns 31 listings of the same item that are offered for sale by the same seller. Of the eight listings showed in the figure, four are offered at a fixed price (“Buy It Now”) of \$124.99 but are associated with different shipping fees (of \$7.99 and \$9.99), and the other four listings are auctions with different ending times and shipping fees. As in Figure 1(a), this particular screenshot was generated on 9/12/2010.


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 presents the relationship between transacted auction prices  $p$  relative to the “reference value”  $v$  of the same item. The latter is defined throughout the paper as the average sale price over all fixed price listings that transact within the same seller experiment. The top panel reports the distribution of the ratio  $p/v$ , while the bottom panel reports the distribution of the difference (in dollars)  $p - v$ . The data underlying these figures is the baseline sample without experiments that are associated with items whose reference value is greater than \$1,000 (these cover only 4,701 (1.9%) experiments and 40,685 (0.5%) listings).

Figure 4: The effect of auction starting price




The top two panels are a graphical presentation of the regression results reported in Table 6 regarding the effect of auction starting price on listing outcomes. The top panel graphs – for different ranges of item values – the probability of sale as a function of the auction (normalized) starting price, and the middle panel presents the sale price (conditional on sale) as a function of the starting price. The third panel tries to assess how much of the price effect is driven by simple, “mechanical” selection by plotting selected points of the truncated sale price distribution for various starting prices (see text for further discussion).

Figure 5: The effect of BIN price



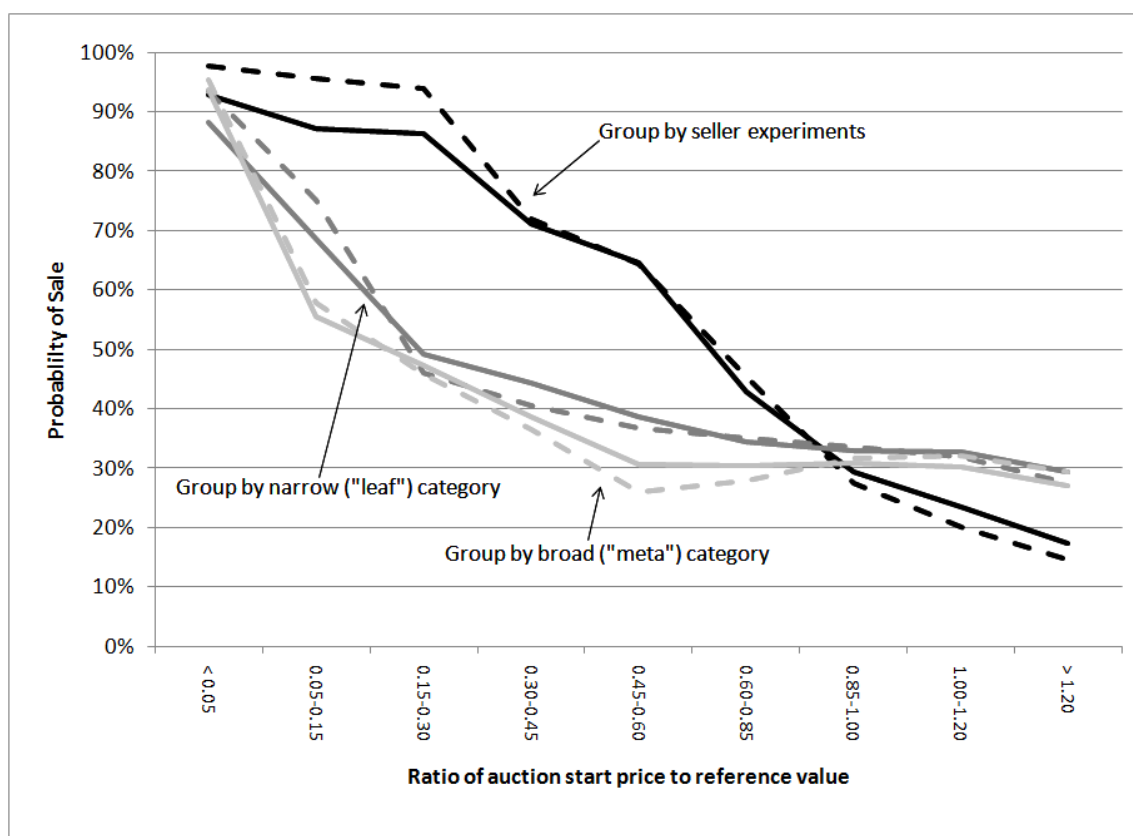
The top two panels are a graphical presentation of the regression results reported in Table 8 regarding the effect of auction BIN price on transaction outcomes (recall that we omit the low item value category and condition on starting price of less than one dollar, so essentially all auctions get transacted). The top panel graphs the probability of the auction transacting using the BIN option as a function of the (relative) BIN price, and the middle panel graphs the transaction price as a function of the BIN price. In a similar spirit to the bottom panel of Figure 4, the bottom panel tries to assess whether the price effect is only due to simple, “mechanical” truncation by plotting selected points of the sale price distributions for different BIN prices (see text for further details).

Figure 6: The effect of shipping fees




The figure is a graphical presentation of the regression results reported in Table 10 regarding the effect of shipping fees. See Table 10 for details.

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
<b>Listings</b>							
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
<b>Transactions</b>							
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		All Auction Experiments	
	(1)		(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 starting price

		Item reference value				All listings
		< \$10	\$10-30	\$30-100	\$100-1,000	
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 uses the baseline sample, and shows the extent of variation in auction starting price. The top panel presents statistics on the variation in (normalized) starting price across experiments, while the bottom panel presents variation within experiments.

Table 6: The effect of auction starting price

	Item reference value			
	< \$10	\$10-30	\$30-100	\$100-1,000
<b><u>Dependent Variable: Sale indicator</u></b>				
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
<b><u>Dependent Variable: Sale price (conditional on sale)</u></b>				
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
Number of listings		< \$10	\$10-30	\$30-100	\$100-1,000	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						Total
		No BIN	< 0.90	0.90 to 0.95	0.95 to 1.00	1.00 to 1.10	> 1.10	
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 uses the baseline sample, and shows the extent of variation in auction BIN price. The top panel presents statistics on the variation in (relative) BIN price across experiments, while the bottom panel presents variation within experiments.

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
<b><u>Dependent Variable: Sale via BIN option indicator</u></b>			
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
<b><u>Dependent Variable: Sale price (conditional on sale)</u></b>			
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. To simplify interpretation, the sample includes all items with reference value greater than \$10 and only listings with starting price that is less than \$1, making 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
Number of listings		< \$10	\$10-30	\$30-100	\$100-1,000	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					Total
		0 to \$2.50	\$2.50 to \$5	\$5 to \$10	\$10 to \$20	> \$20	
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	
<b>Dependent Variable: Sale indicator</b>								
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	
<b>Dependent Variable: Sale price (conditional on sale)</b>								
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)
<b>Pooled</b>				-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.