

NBER WORKING PAPER SERIES

CONSUMER PRICE SEARCH AND PLATFORM DESIGN IN INTERNET COMMERCE

Michael Dinerstein
Liran Einav
Jonathan Levin
Neel Sundaresan

Working Paper 20415
<http://www.nber.org/papers/w20415>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2014

We appreciate support from the Alfred P. Sloan Foundation, the National Science Foundation, the Stanford Institute for Economic Policy Research, and the Toulouse Network on Information Technology. We thank Greg Lewis for helpful comments. Data access for this study was obtained under a contract between the Stanford authors (Dinerstein, Einav, and Levin) and eBay Research Labs. 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/w20415.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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Michael Dinerstein, Liran Einav, Jonathan 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.

Consumer Price Search and Platform Design in Internet Commerce
Michael Dinerstein, Liran Einav, Jonathan Levin, and Neel Sundaresan
NBER Working Paper No. 20415
August 2014
JEL No. D12,D22,D47,D83,L13,L86

ABSTRACT

Search frictions can explain why the "law of one price" fails in retail markets and why even firms selling commodity products have pricing power. In online commerce, physical search costs are low, yet price dispersion is common. We use browsing data from eBay to estimate a model of consumer search and price competition when retailers offer homogeneous goods. We find that retail margins are on the order of 10%, and use the model to analyze the design of search rankings. Our model explains most of the effects of a major re-design of eBay's product search, and allows us to identify conditions where narrowing consumer choice sets can be pro-competitive. Finally, we examine a subsequent A/B experiment run by eBay that illustrates the greater difficulties in designing search algorithms for differentiated products, where price is only one of the relevant product attributes.

Michael Dinerstein
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305
mdiner@stanford.edu

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

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

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

1 Introduction

Search frictions play an important role in retail markets. They help explain how retailers maintain positive mark-ups even when they compete to sell identical goods, and why price dispersion is so ubiquitous. In online commerce, the physical costs of search are much lower than in traditional offline settings. Yet studies of e-commerce routinely have found substantial price dispersion (Bailey, 1998; Smith and Brynjolfsson, 2001; Baye, Morgan, and Scholten, 2004; Einav et al., forthcoming). And despite the general view that the internet has increased retail price competition, we are not aware of a definitive study measuring online mark-ups, or comparing them to current or past offline mark-ups.

Consumers shopping online can use price search engines or compare prices at e-commerce marketplaces such as eBay or Amazon. For the most part, these platforms want to limit search frictions and provide consumers with transparent and low prices (Baye and Morgan, 2001). Retailers may have very different incentives. Many retailers, and certainly those with no particular cost advantage, would like to differentiate or even “obfuscate” their offerings to limit price competition (Gabaix and Laibson, 2006; Ellison and Ellison, 2009; Ellison and Wolitzky, 2012). This raises the question of how different ways of structuring online search, such as alternative search rankings or displays, affect price competition and consumer purchasing patterns.

In this paper, we use a model of consumer search and price competition to estimate search frictions and online retail margins, and to study the effects of search design. We estimate the model using browsing data from eBay. A nice feature of internet data is that it is possible to track exactly what each consumer sees. As a practical matter, consumers often evaluate only a handful of products, even when there are many competing sellers. With standard transaction data, incorporating this requires the introduction of a new latent variable, the consumer’s “consideration set”; that is, the set of products the consumer actually chooses between (e.g. Goeree, 2008). Here, we adopt the consideration set approach, but use browsing data to recover it.

We use the model to estimate consumer demand and retail margins, and then to analyze a large-scale redesign of the search process on eBay. Prior to the redesign, consumers

entering a search query were shown individual offers drawn from a larger set of potential matches, ranked according to a relevance algorithm. The redesign broke consumer search into two steps: first prompting consumers to identify an exact product, then comparing seller listings of that product head-to-head, ranked (mostly) by price. We discuss in Section 2 how variations on these two approaches are used by many, if not most, e-commerce platforms.

To motivate the analysis, we show in Section 3 that across a fairly broad set of consumer product categories, re-organizing the search process is associated with both a change in purchasing patterns and a fall in the distribution of posted prices. After the change, transaction prices fell by roughly 5-15% for many products. We also point out that all of these categories are characterized by a wide degree of price dispersion, and by difficulties in accurately classifying and filtering relevant products. Despite a very large number of sellers offering high-volume products, consumers see only a relatively small fraction of offers, and regularly do not buy from the lowest-price seller. That is, search frictions appear to be prevalent despite the low physical search costs associated with internet browsing.

We propose our model of consumer demand and price competition in Section 4, and estimate it in Section 5 for a specific and highly homogeneous product, the Halo Reach video game. We find that even after incorporating limited search, demand is highly price sensitive. Price elasticities are on the order of -10. We do find some degree of consumer preference across retailers, especially for sellers who are “top-rated,” a characteristic that eBay flags conspicuously in the search process. We also use the model to decompose the sources of seller pricing power and the high degree of homogeneous product price dispersion into three sources: variation in seller costs, perceived seller differentiation, and search frictions.

We estimate the model using data from before the search redesign. In Section 6, we apply the model (out-of-sample) to analyze the search redesign. The model can explain, both qualitatively and quantitatively, many of the effects of the redesign: a reduction in posted prices, a shift toward lower-priced purchases, and consequently a reduction in transaction prices. The redesign had the effect of increasing the set of relevant offers exposed to consumers, and prioritizing low price offers. We find that the latter effect is by far the most important in terms of increasing price sensitivity and competitive pressure. In fact, we find that under the redesigned selection algorithm that prioritizes low prices, narrowing the number of listings

shown to sellers tends to increase, rather than decrease, price competition.

The final section of the paper discusses a randomized A/B experiment that eBay ran subsequent to the search re-design. The experiment randomized the default search results presented to consumers. The experiment is not appropriate for testing equilibrium predictions about how sellers adjust prices, as we do in the main part of the paper, but it can be used to look at consumer purchasing behavior, holding seller behavior fixed. Interestingly, the search design that performed best overall was to start consumers with relevance results and allow them to browse toward more structured price rankings. We show that the results can be explained by distinguishing relatively homogenous product categories from those that include more heterogeneous listings. As one might expect, more structured price search performed well for the first type of products, less well for the latter type.

Our paper is related to an important literature on search frictions and price competition that dates back to Stigler (1961). Recent empirical contributions include Hortacsu and Syverson (2003), Hong and Shum (2006), and Hortacsu et al. (2012). A number of papers specifically have tried to assess price dispersion in online markets (e.g. Bailey, 1998; Smith and Brynjolfsson, 2001; Baye, Morgan, and Scholten, 2004; Einav et al., forthcoming), to estimate price elasticities (e.g. Ellison and Ellison, 2009; Einav et al., 2014), or to show that consumer search may be relatively limited (Malmendier and Lee, 2011). Ellison and Ellison (2014) propose a model to rationalize price dispersion based on sellers having different consumer arrival rates, and use the model to analyze online and offline prices for used books. Their model is natural for thinking about consumer search across different websites. Fradkin (2014) and Horton (2014) are two other recent papers that study search design for internet platforms, in both cases focusing on settings where there is a richer two-sided matching problem.

2 Search Design in Online Markets

There are at least two dimensions of consumer search in online markets. The first is to guide consumers toward relevant products, either in response to a user query, or through advertising or product recommendations. The second is to help consumers find a retailer

offering an attractive price for a product the consumer knows that he wants. For most of this paper we focus on the latter problem of “price search,” although allowing for the possibility that consumers may perceive sellers as somewhat differentiated.

We start in this section by describing how different online platforms approach the search problem. Platforms have to identify a relevant set of offers, and present the information to consumers. Identifying relevant offers is easier when products have well-defined SKUs or catalog numbers. But as we will note below, it is still a difficult problem for platforms that have tens of thousands of different listed products. Platforms also take different approaches to presenting information. A typical consideration is whether to try to present all the relevant products in a single ordered list that attempts to prioritize items of highest interest, or try to classify products into sets of “identical” products, and then order products within each set based on price or other vertical attributes.

Figure 1 contrasts the approaches of three prominent e-commerce sites. Each panel shows the search results that follow a query for “playstation 3.” At the top, Craigslist presents a list of items that it judges to be relevant, ordered by listing date. The buyer must navigate what is potentially a long and loosely filtered list to find his ideal match. On the other hand, because the top listings are recent, the item is more likely to still be available than in lower listings, which helps to address the fact that Craigslist listings do not necessarily disappear if the seller stocks out. In the bottom panel, Amazon takes the other extreme. It highlights a single product model (the 160 GB version) and quotes the lowest price. Buyers can change the model, or click through to see a list of individual sellers, ordered by price. In the middle panel, Google Shopping takes a somewhat intermediate approach.

These approaches to search design illustrate some trade-offs. Erring on the side of inclusiveness makes it more difficult for a buyer to find the lowest price for a specific well-defined product. On the other hand, it allows for serendipitous matches, and provides more opportunities to sellers who may be less professional in categorizing their products. The latter approach works well for a shopper interested in price comparisons, and would seem to promote price competition, provided that the platform is able to accurately identify and classify listings according to the product being offered. At the same time, as Ellison and Ellison (2009) have highlighted, it may provide sellers with a strong incentive to search for unpro-

ductive tactics that avoid head-to-head price competition.

The redesign of eBay’s search process is interesting because it allows for a comparison of these approaches, as shown in Figure 2. The top panel shows eBay’s traditional listings page. It is generated by an algorithm that first filters listings based on query terms, and then presents the listings according to a ranking order. The default is a relevance ranking that eBay calls Best Match.¹ Users can change the sort order or refine their search in various ways. Unlike some search results on the internet, the Best Match algorithm traditionally has not been tailored to individual users, nor did it consider price explicitly.² While it may seem strange not to use price as an explicit ranking factor, it is less surprising when one appreciates the difficulty of filtering the set of products. For example, re-sorting the displayed page on price would have yielded cheap accessories (e.g. cables or replacement buttons or controllers).

In 2011, eBay introduced an alternative two-stage search design. A buyer first sees the relevant product models (e.g. a user who searches for “iPhone” sees “Black iPhone 4s 16GB (AT&T)” and other models). The buyer then clicks on the model to see a product page with specific listings, shown in the bottom panel of Figure 2.³ The product page has a prominent “Buy Box” that displays the top-rated seller with the lowest posted price (plus shipping). Then there are two columns of listings, one for auctions and one for posted prices. The posted price listings are ranked in order of price plus shipping (and the first listing may be cheaper than the Buy Box if the lowest-price seller is not top-rated). The auction listings are ranked so that the auction ending soonest is on top. We will not focus on auctions, which represent 33% of the transactions for the products on which we focus.

¹When eBay was predominantly an auction platform, it sorted listings in order of their ending time, with listings set to expire soonest at the top of the page. This ordering is still used for auction results, but eBay introduced the more multi-dimensional Best Match ordering in 2008.

²At various time, the Best Match algorithm has incorporated price or attempted more tailoring with respect to individual users, but it did not during the period we study. However, it does incorporate factors that may be correlated with prices. For instance, if Best Match moved sellers with high conversion rates up in the search, and these sellers are likely to have low prices, then Best Match results may effectively prioritize low prices.

³The concept of a product page existed on eBay earlier, but its design was very different and it was difficult to find, so that only a small minority of users ever viewed it.

3 Effect of Platform Change on Search and Prices

To help motivate the model and the analysis below, we start by presenting some statistics from before and after the search redesign. The new product page was introduced on May 19, 2011.⁴ However, the traditional listing page remained the default view for buyers. The new product page became the default presentation of search results for five large categories — cell phones, digital cameras, textbooks, video games, and video game systems — over a one-week period from June 27, 2011 to July 2, 2011. The traditional Best Match results were still accessible to buyers, so the best way to view the change is probably to think of buyers as now having access to two types of search results, and being nudged toward the product page.

Table 1 shows statistics for these five categories in the period before the product page was introduced (April 6 to May 18) and the period after the introduction was completed (August 1 to September 20). We drop the intermediate period during which the product page was available, but not the default. We also exclude the month of July to allow time for sellers to respond to the platform redesign. The sample period covers nearly half a year, so one potential concern is that there may have been changes in the set of products available, especially in the categories with shorter product life cycles. To deal with this, we restrict attention to the ten products in each category that were most commonly transacted in the week before the product page became the default. As an example, a typical product in the cell phone category is the black, 16GB iPhone 4 for use with AT&T. We also show statistics for the narrower product category of iPhone 4.

Several patterns are clear in the data. There are many listings for each product. The average number of listings ranges from 16 to 41 across the five categories. There is also remarkable variation in prices. The average ratio of the 75th percentile price to the 25th percentile price is 1.22 in cell phones, 1.32 in digital cameras, and higher in the other categories. The extreme prices, especially on the high end, are even more dramatic. Consumers generally do not purchase at the lowest price. In the period before the redesign the average purchase price often was around the 25-40th percentile of the price distribution. As an ex-

⁴eBay ran a small pilot in September 2010 and implemented the product page for the GPS, DVD, and MP3 categories. These categories are not included in our subsequent analyses.

ample, in the digital camera category, consumers pay on average around 18% more than if they had selected the 10th percentile price.

The comparison between the two periods is also informative. With one exception (video game systems), transacted prices fell in every category after the new product page was introduced. The fall was relatively small in the cell phone and video game categories (2.1% and 7.7%, respectively), and larger in digital cameras and textbooks (15.7% and 15.9%). The decrease does not appear to be driven by a general time trend. The qualitative results remain similar when we control for product-specific (linear) time trends. In part, the drop in transacted prices reflects a fall in the posted prices that were being offered. Posted prices fell in every category (again, with the exception of video game systems), by between 0.9% and 17.7%.

Several statistics are suggestive of changes in consumer search. In every category except one, consumers after the redesign purchased items that were cheaper relative to the current distribution of prices. The share of purchases from top-rated sellers also increased markedly for many of the products. Both of these results seem fairly natural. The redesigned search selects and sorts listings by price, focusing attention on the low-price offers, and the product page Buy Box especially promotes the low-priced top-rated seller.⁵

Figure 3 presents a final piece of descriptive evidence, that is also consistent with a change in consumer search patterns after the redesign. The figure is constructed using browsing data for a single product, the video game Halo Reach, which we use to estimate our model below. The top panel shows the distribution of relevant Halo Reach offers that were displayed to each consumer following a targeted search, before and after the change in the search design. The size of the consumer “consideration set” increased sharply. The second panel shows the distribution of the total number of clicks made in a browsing session, for consumers who ended up purchasing. After the search redesign, consumers generally clicked fewer times on their way to a purchase, consistent with a more streamlined process.

These results provide a descriptive and qualitative sense of the overall effects of the platform change. After the change, transaction prices fell for many products. This appears

⁵As mentioned, we focus on the August-September “After” period, because it seemed plausible that the effect of the change on seller’s pricing may take some time to play out. The July results are generally intermediate, with most of the change in TRS transactions, and price percentile changes occurring immediately.

to a have resulted from both a change in purchasing patterns and a fall in the distribution of posted prices. In the next section we develop and estimate a more complete model of the underlying economic primitives. The model allows us to explain the degree of price dispersion and the purchasing patterns in the data, and separate the demand and pricing incentive effects of the platform change, as well as to evaluate alternative platform changes not present in the data.

4 Model

In this section, we describe a model of consumer search and price competition. Below, we estimate the model’s parameters using data from a single product market. We use the estimates to quantify search frictions, the importance of retailer and listing heterogeneity, the size of retailer margins, and the way that the platform re-design affected these quantities.

The model’s ingredients are fairly standard. Each potential buyer considers a specific and limited set of products. He or she then chooses the most preferred. This is modeled as a traditional discrete choice problem. Sellers set prices in a Nash Equilibrium, taking into account buyer demand. The role of the platform is to shape consumer search. Rather than considering all available products, consumers consider the ones suggested by the platform. We take advantage of detailed browsing histories to explicitly collect data on each buyer’s consideration set. In this context, search rankings affect the set of considered products, and hence consumer choices, and indirectly, the incentives for price competition.

4.1 Consumer Demand

We consider a market in which, at a given point in time, there are a large number of different sellers offering a given product. In our current specification, we allow sellers to vary only by their price p and by whether they are top-rated seller (denoted TRS). We attribute any additional differentiation to a logit error. We assume that consumer i ’s utility from seller j is given by

$$u_{ij} = \alpha_0 + \alpha_1 p_j + \alpha_2 TRS_j + \alpha_3 p_j TRS_j + \varepsilon_{ij}, \quad (1)$$

where ϵ_{ij} is distributed Type I extreme value and is independent of the seller’s price and *TRS* status.

The main distinction of the model comes in analyzing the consideration set. The consideration set is denoted by J_i , such that $J_i \subseteq \mathcal{J}$, where \mathcal{J} is the set of all available offerings on the platform. We assume that the outside good, good 0, which represents either not buying the product or buying it via another sales channel or by auction, is also part of the consideration set. It has utility $u_{i0} = \epsilon_{i0}$, where ϵ_{i0} is also an independent Type I extreme value random variable. Consumers choose the utility-maximizing option in their consideration set.

To estimate the demand parameters, we rely on our browsing data to identify the consideration sets of a large sample of buyers, and their resulting choices. Specifically, we assume the consideration set includes all the listings on the page seen by the consumer following his last search query. This is usually the listings page prior to the platform re-design, and the product page afterwards. With an observable consideration set for each buyer, demand estimation is straightforward using the familiar multinomial logit choice probabilities.

4.2 Consideration Sets

In order to analyze pricing decisions, and make “out-of-sample” predictions, we also develop a simple econometric model of how consideration sets are formed. To do this, we assume that consumer i observes the offers of L_i sellers, where L_i is random. We estimate its distribution directly from the data, that is, by measuring the frequency with which observed consideration sets include a given number of relevant listings. We assume that L_i is independent of any particular buyer characteristics, or the distribution of prices.

Which sellers make it into the consideration set? Prior to the re-design, we noted that price did not factor directly into search ranking, but that after the re-design, it played a predominant role. In practice, the complexity of the search ranking and filtering algorithms, which must be general enough to work for every possible search query and product, as well as factors such as which server provides the results, adds less purposeful (and perhaps unintentional) elements to what results are shown.

To capture this, we adopt a stochastic model of how sellers are selected onto the displayed page. Specifically, we assume that products are sampled from the set of available products

J_t , such that each product $j \in J_t$ is associated with a sampling weight of

$$w_j = \exp \left[-\gamma \left(\frac{p_j - \min_{k \in J_t}(p_k)}{\text{std}_{k \in J_t}(p_k)} \right) \right], \quad (2)$$

and consumer i 's consideration set is then constructed by sampling L_i products from J_i , without replacement. This implies that the consideration set is drawn from a Wallenius' non-central hypergeometric distribution. We assume that prior to the platform change offers enter the consideration set independent of their price, so $\gamma = 0$. For the period after the platform change when the product page became available, we expect $\gamma > 0$ so that lower price items are disproportionately selected into the consideration set. Below we estimate γ using the browsing data that records the listings that appeared on pages buyers actually visited.

4.3 Pricing Behavior

We model seller pricing using a standard Nash Equilibrium assumption. Each seller sets its price to solve

$$\max_{p_j} (p_j - c_j) D_j(p_j). \quad (3)$$

Here $D_j(p_j)$ is the probability a given buyer at period t selects j 's product, given the set of offerings \mathcal{J} . From a seller's perspective, $D_j(p_j)$ depends on how consumers form their consideration sets, as well as the choices they make given their options. Using the logit choice probabilities, we have:

$$D_j(p_j) = \sum_{J: j \in J \subseteq \mathcal{J}} \left[\frac{\exp(\alpha_0 + \alpha_1 p_j + \alpha_2 TRS_j + \alpha_3 p_j TRS_j)}{1 + \sum_{k \in J} \exp(\alpha_0 + \alpha_1 p_k + \alpha_2 TRS_k + \alpha_3 p_k TRS_k)} \right] \Pr(J|\mathcal{J}). \quad (4)$$

Another important consideration here is the set (\mathcal{J}) of competing items that the seller has in mind when it sets its price. We assume that the seller optimizes against the set of competing items and prices that are available on eBay during the approximately one month (either "before" or "after") period considered. One argument for this assumption is that in principle sellers can change their price at any time for no cost, so that current listings are the most relevant. Of course, in practice sellers do not change prices that often, so one

could reasonably consider price-setting decisions that take into account the (stochastic) set of competing products over the entire lifetime of the listing. This is beyond the scope of the current paper; see Backus and Lewis (2012) and Knoepfle (in progress) for related work.

To understand the seller’s pricing incentives, it is useful to write $D_j(p_j) = A_j(p_j) Q_j(p_j)$, where A_j is the probability that the listing enters the consideration set given p_j and \mathcal{J} , and Q_j is the probability that the consumer purchases item j conditional on being in the listing set. With this notation, the optimal price p_j satisfies:

$$\frac{p_j}{c_j} = \left(1 + \frac{1}{\eta_D}\right)^{-1} = \left(1 + \frac{1}{\eta_A + \eta_Q}\right)^{-1}, \quad (5)$$

where η_D , η_A , η_Q are respective price elasticities. When $\gamma > 0$, reducing price increases demand in two ways: by making it more likely that the seller ends up in the consideration set ($\eta_A < 0$) and by making it more likely that the consumer picks the seller, conditional on the seller being in the choice set ($\eta_Q < 0$). Increasing γ intensifies the first effect. In addition, increasing γ effectively faces each seller with tougher competition conditional on making it into the consideration set, by reducing the likely prices of the other sellers who are selected.

4.4 Discussion

The model we have chosen has only a handful of parameters. A main reason is that we wanted something easy to estimate and potentially “portable” across products, but yet with enough richness to be interesting. The assumptions we have chosen relate fairly closely to some of the classic search models in the literature. For example, in Stahl’s (1989) model there are two types of consumers: consumers who (optimally) sample a single offer completely at random, and consumers who sample all the offers. This corresponds to having $L \in \{1, |\mathcal{J}|\}$ and $\gamma = 0$. Stahl’s model has no product differentiation and the pricing equilibrium is in mixed strategies, but it has very intuitive properties. For instance, if more consumers have $L = 1$, equilibrium prices are higher. Consideration set sizes have the same effect in our model with $\gamma = 0$. The same need not be true with $\gamma > 0$. For instance, suppose that sellers have identical cost and none are top-rated. As $\gamma \rightarrow \infty$, consideration sets are selected

purely on the basis of price. Then having $L = 1$ for all consumers creates perfect Bertrand competition, whereas if $L = |\mathcal{J}|$ we have a symmetric logit demand model with consequent mark-ups.

There are several obvious directions in which our model can be extended and we have explored some of them. One is to allow for more heterogeneity among sellers or consumers. Including more seller heterogeneity seems unlikely to change the model’s performance very much. It might be more interesting to distinguish between price-elastic “searchers” and price-inelastic “convenience” shoppers, as in Stahl (1989) or Ellison (2005). We also have not focused on search rank. In their study of a price search engine, Ellison and Ellison (2009) find page order, especially first position, to be very important, and it is perceived to be very important in sponsored search advertising. We have estimated versions of our model that include page order, but decided not to focus on these versions. One reason is that the effect of page order in our data seems to be far less dramatic than in sponsored search. The estimates also are much harder to interpret, a significant drawback given the modest increase in explanatory power.⁶

5 Empirical Estimates

We now describe the data we use to estimate the model parameters, and the parameter estimates.

5.1 Estimation Sample

For this part of the paper, we focus on a single, well-defined product: the popular Microsoft Xbox 360 video game, Halo Reach. This video game is one in a series of Halo video games. It was released in September 2010. Microsoft originally set an official list price of \$59.99, which it shortly dropped to \$39.99. We chose this specific game because a large number of units transact on eBay, and because it had a relatively stable supply and demand during our

⁶One reason for this is that, to the extent that rank and price are correlated, it is somewhat challenging to identify the two terms separately. Another issue is that pages tend to include many “irrelevant” items (accessories, etc.) as well as auctions, which makes for many complicated modeling decisions in terms of whether to include absolute rank, or relative rank among “relevant” listings, or some mixture of the two.

observation period of Spring-Summer 2011. The prices of many consumer electronics on the platform exhibit a time trend, usually starting high and falling quickly over the product life-cycle. Others have a range of characteristics that vary across listings, complicating demand and supply estimation.

The data for the analysis come directly from eBay. They include all listing-level characteristics as well as individual user searches. We can observe every aspect of the search process, including what the user saw and her actions. We use data from two periods: the “before” period from April 6 until May 18, 2011, and the “after” period, which we define to be August 1 until September 20, 2011.⁷ The search data consist of all visits to the Halo Reach product page as well as all visits to the standard search results page derived from query terms that include the words “xbox” (or “x-box”), “halo,” and “reach.” We drop searches or product page visits that do not result in at least one click on a Halo Reach item.⁸ This results in 1,527 visits to the search results page (1,297 of them in the pre-period) and 3,950 visits to the product page (190 in the pre-period).⁹

As search results often include extraneous results while the product page only shows items that are listed under “Halo Reach” in eBay’s catalog, we identify listings as the Halo Reach video game if eBay catalogued them as such. We also visually inspected each listing’s title to verify that the listing is for just the video game. Illustrating the difficulty of precisely filtering listings, even after we restrict attention to listings catalogued as Halo Reach, we found that 12% of listings were not Halo Reach-related, and 33% were not the game itself (e.g. they were accessories). Items in this second group often seemed to have very low or high prices, so we dropped all listings with prices below \$15 or above \$100, in case we failed to identify them based on listing title alone. We also restrict the analysis to new items, listed either with a posted price, or as an auction but with a Buy-It-Now price.¹⁰

⁷As before, we drop July 2-31, 2011, when the product page was the default because our descriptive analysis in Section 3 suggested that price adjustment did not happen immediately and we want to use an equilibrium model for prediction. The predictive fit is similar for demand if we include July, and a bit worse for pricing.

⁸This choice mainly affects the definition of the outside option in the demand model. Results are largely similar when we use alternative definitions of the outside option.

⁹The “product page” in the pre-period was more rudimentary than one introduced on May 19 (see footnote 3), and relatively few people navigated to it.

¹⁰According to eBay, “new” items must be unopened and usually still have the manufacturer’s sealing or original shrink wrap. The auction listings with a Buy-It-Now price have a posted price that is available until

Finally, as mentioned earlier, sellers are allowed to change a listing’s price even after it has been listed. When this happens, we always observe whether there has been a price change, and we observe the price if there was a transaction, or is a user in our search data clicked on the item, or if it was the final posted price of the listing. This leaves a relatively small number of cases where we have a listing for which we know the price was changed but do not observe the exact price because the listing was ignored during this period.¹¹

5.2 Descriptive Statistics

Table 2 reports summary statistics for the before and after periods. The numbers of sellers and listings are slightly lower in the after period, and more of the listings come from top-rated sellers. These differences, particularly the increase in top-rated seller listings, could be a consequence of the platform change. In addition, the mean and median list prices both drop by about \$2 in the after period, which is consistent with the earlier results on a broader set of products in Section 3, and with the hypothesis that competitive pressure increased after the platform change.

The bottom panel in Table 2 shows statistics on searches. Consumers saw lower prices in the after period, and a larger fraction of searches resulted in purchases (13.0% compared to 10.3%). Recall that in Figure 3, displayed earlier, we already showed that there was a significant increase in the number of relevant listings consumers saw after a search. We also showed in Figure 3 that eventual purchasers seem to have had an easier time getting to the point of sale: eventual purchasers had to click fewer times after the platform change.

5.3 Model Estimates

To estimate the parameters of the model, we use the data on consumer choices and consideration set sizes to estimate the demand parameters, and then impose an assumption of optimal pricing to back out the implied marginal costs of each seller.

the first bid has been made. We only consider these listings during the period prior to the first bid.

¹¹For 89% of the listings in the data, the price is never missing. For the remaining 11% the price is missing during some of the time in which they are active. We use these listings for estimation when their prices are known, but drop them from the analysis when the price is unknown.

Estimating the demand parameters is straightforward. As described earlier, we have a standard logit demand with individual-level data and observed individual-specific consideration sets. We estimate the demand parameters using maximum likelihood, restricting attention only to consumer data from the before period. The results appear in the first column of Table 3. The top-rated seller (TRS) indicator is quite important. It is equivalent to nearly a \$10 price discount (off an average price of less than \$40!). Recall that in the before period, there is no advantage given to TRS sellers that is analogous to the Buy Box introduced in the search re-design, so this effect is really very large. Price also has a very large effect. The price elasticity implied by the estimates is about -10. It is even higher (closer to -13) for TRS sellers. The profit margin implied by these estimates is about 10%: \$3.23 for TRS sellers and just over \$4 for other sellers.

The next step is to estimate the consideration set model. We obtain the empirical distribution of L_i (the number of items sampled by a consumer) directly from the browsing data, and separately for the before and after periods (see Figure 3). We also use the browsing data to estimate the sampling parameter γ in equation (2) that determines the extent to which cheaper listings are more likely to enter the results page. For the before period, we assumed $\gamma = 0$. For the after period, we estimate γ using maximum likelihood and obtain an estimate of 0.81 (with a standard error of 0.18). This implies that a ten percent reduction in the posted price would, on average, make the listing 29% more likely to be part of a consumer’s consideration set.

The last step is to estimate seller costs. From the seller’s optimization problem, we have:

$$c_j = p_j + \frac{D_{jt}(p_j)}{D'_{jt}(p_j)}, \quad (6)$$

where D_{jt} depends on the search process and consumer choices. We use the estimated demand parameters from the first estimation stage, combined with the consideration set model to obtain estimates of D_{jt} and D'_{jt} for every seller in the “before” period. Then we use the first order condition above to back out the cost c_j that rationalizes each seller’s price as optimal.

The implied cost distribution is presented in Figure 4, which also shows the optimal

pricing functions for both TRS and non-TRS sellers. We estimate a fair amount of dispersion in seller costs. The 25th percentile of the cost distribution is just slightly under \$30; the 75th percentile is \$40. There are also a considerable number of sellers who post extremely high prices. Thirteen percent post prices above \$50, and five percent post prices above \$60! To rationalize these prices, we infer that these most extreme sellers all have costs about \$59.¹² We discuss the high price sellers in more detail in Section 6.3.

6 Applying the Model

In this section, we use the estimated model to evaluate the search redesign and compare the model predictions to the data. Then we apply the model to consider various ways of reducing search frictions and to identify the sources of online price dispersion.

6.1 Changing the Search Design

We use our estimates to assess the introduction of the product page. To do this, we combine our demand and cost estimates from the before period, with our estimates of the consideration set process from the after period. We use this combined model to calculate equilibrium prices and expected sales with the post-redesign search process, assuming that consumer choice behavior and the seller cost distribution remains unchanged. The results from this exercise are reported in Table 4, and Figures 5 and 6. In particular, Table 4 shows model-based estimates of optimal seller margins for scenarios where we impose specific effects of the redesign, as well as the full redesign.

A main effect of the platform change was to make demand more responsive to seller prices. Figure 5 provides a visual illustration of this change in incentives. It shows the demand curves from the model, for TRS and non-TRS sellers, for both periods. Demand became considerably more elastic in the after period, with the largest effect for TRS sellers.

¹²We also investigated whether the implied cost distribution was sensitive to our assumptions about the consideration set. Interestingly, it is not. Re-estimating the model under the assumption that consumers consider the entire set of available items leads to a similar cost distribution. This likely reflects the fact that prior to the platform re-design, the observed consideration sets are quite representative, in terms of listed prices, of the full set of listings.

The implication is that seller margins should fall. Comparing the top and bottom rows of Table 4 shows that the median optimal margin fell from \$3.23 (or 9% of price) to \$2.70 for TRS sellers, and from \$4.06 to \$3.23 for non-TRS sellers, implying roughly a twenty percent fall in profit margins.

Several factors may have contributed to the shift in seller incentives. As we showed in Figure 3, there was a noticeable increase in the size of consideration sets, and buyers had a much smaller chance of seeing just a single relevant listing. In addition, price became an important factor in entering the consideration set. With our estimate of $\gamma = 0.81$ for the after period, a ten percent price reduction increases the odds of appearing in the consideration set from 0.24 to 0.31, providing sellers with a new incentive to reduce prices. Finally, there was an increase in the number of available listings, which may or may not have been directly related to the platform change.

To assess the relative importance of these effects, we start with the model from the before period and separately impose the increase in listings, the increase in consideration set size, and the increase in γ . In each case, we compute the new pricing equilibrium. The middle rows of Table 4 report the median equilibrium margin for TRS and non-TRS sellers for each of the three scenarios, and also the predicted buyer purchase rate. Making price a factor in selecting what listings to display (i.e. increasing γ) has by far the largest effect on seller incentives and purchase rates. The increased size of consideration sets has only a small effect on equilibrium margins and purchase rates. The same is true for the increase in the number of sellers.

These calculations are based on model estimates obtained primarily using the “before” data. A natural question is whether the model’s predictions for the after period are similar to the outcomes we actually observe. Figure 6 compares seller prices. It plots the distribution of prices in the before period (where the model matches the data by construction), and then both the distribution of prices for the after period predicted by the model, and observed in the data. The predicted and observed distributions are reasonably close. So at least for seller prices, the model’s out-of-sample predictions match quite well with what happened. It is also possible to compare the consumer purchase rates predicted by the model, and those that we observe after the redesign. These are shown in the bottom rows of Table 3. They also

are reasonably close (the model predicts 12.3%, which is a bit less than the 13.0% observed in the data).

6.2 Search Frictions and Price Dispersion

At the beginning of the paper, we posed the question of why internet prices for homogeneous goods are so dispersed, despite the seemingly low search costs. Prices in our sample, as in earlier studies, exhibit a high degree of dispersion. The estimated model provides a way for us to understand the source of this dispersion, and also the source of seller margins. In particular, the model offers three ways in which outcomes might differ from the simplest homogeneous good Bertrand pricing environment: dispersion in costs; search frictions that provide market power and perhaps equilibrium price dispersion; and perceived seller differentiation that supports positive seller margins.

We analyze these factors in Table 5. The Table compares equilibrium outcomes for variations of the model that differ along two dimensions. Across the columns, we vary the degree of search frictions. In the first and second columns, we consider the before and after search regimes. In the third column, we assume that all Halo Reach listings on the platform enter each consumer’s consideration set. Across the two rows, we vary the degree of product differentiation. The “differentiation” model assumes the estimated logit demand, in which each seller enjoys some market power. In the “limited differentiation” model, we assume a nested logit demand structure in which the outside good is one nest and all sellers are part of a second nest. Specifically, the ε_{ij} in our logit demand model (1) becomes $\zeta_{ig} + (1 - \sigma) \varepsilon_{ij}$, where all sellers share the same ζ_{ig} , whose distribution depends on σ (see Berry, 1994). The “limited differentiation” model assumes $\sigma = 0.2$, which reduces the weight on the seller-specific error and make the products much less differentiated than the baseline logit “differentiation” case, which corresponds to $\sigma = 0$.

In all six scenarios, we fix the distribution of seller costs (as shown in Figure 4), and draw costs for each seller on the platform (assuming 19 sellers, which is the mean from the before period). Sellers are assumed to set prices knowing the assumptions about consumer search and choice behavior, but without knowledge of the exact realization of opponents’ costs. To solve for equilibrium prices and mark-ups, we start from the original price distribution and

update sellers' prices one-by-one using their first-order conditions with the counterfactual model and the new price distribution. We continue iterating over sellers until every seller's first-order condition simultaneously holds.

The results can be used to understand both the source of seller margins and the sources of price dispersion. First consider the case with no search costs and limited product differentiation (top right). In this scenario, sellers sustain positive margins only because there is some possibility that they have a strictly lower cost than all competing sellers (as in the incomplete information Bertrand pricing model of Spulber, 1995). The median mark-up is less than \$1, and the average transaction price is \$24. There is considerable dispersion in posted prices, stemming from the estimated variation in seller costs. However, transacted prices are much more concentrated.

As we incorporate search frictions (moving from right to left on the top row), we see that search frictions lead to substantially increased mark-ups, and somewhat higher transaction prices and price dispersion. Seller differentiation, however, is an even more potent force for pricing power and (transaction) price dispersion. For any assumption about search frictions, increased seller differentiation leads to higher markups, higher prices and greater price dispersion. Moreover, even with no search frictions prices are higher and more dispersed than in any of the limited differentiation cases. Interestingly, once seller differentiation is present, the “after” search regime actually leads to more intense price competition, and somewhat less price dispersion, than is present with no search frictions. The reason, of course, is that the (limited) consideration set is selected with significant weight on price, whereas given a choice set, consumers focus on the idiosyncratic match (the ε_{ij}) as well as price.

6.3 Discussion and Extensions

We considered a number of other permutations of the model. In one exercise, we investigated the importance of obfuscation and the ability of the platform to filter less relevant listings. As mentioned earlier, it is common to see eBay search results that do not perfectly match the item that the potential buyer likely was interested in. For example, a search for iPhone may show some iPhone covers or chargers, or other accessories. The case is similar with Halo Reach. We examined the consequences of more perfect filtering by increasing the size

of the consideration sets L_i , and recomputing the pricing equilibrium. These results are not reported, but we found the effects were not large, and in fact prices (and margins) are slightly higher compared to the “after” search regime. This is because having a larger consideration set has two effects. One effect is the increase in competition, which pushes sellers to lower prices. The second effect is that it becomes easier to enter the consideration set, reducing the incentive to price low as a way to become visible. The latter effect (slightly) dominates.

We also investigated the extent to which observed price dispersion results from dispersion in seller margins rather than dispersion in costs. Recall that in traditional search models such as Stahl (1989), *all* of the dispersion in prices comes from variation in margins, and in equilibrium the price elasticity of (residual) demand is determined so that sellers are indifferent across a wide range of prices. An empirical implication of this equilibrium is that (residual) demand must be (very) log-convex: that is, shaped so that cost increases are passed through into prices *more* than one-for-one (Weyl and Fabinger, 2013). We considered generalizations of our logit demand model that allow for log-convexity, but consistently obtained estimates consistent with log-concave demand. With this type of curvature, optimal margins are lower for higher cost sellers (see Figure 4), so that equilibrium price dispersion is in fact less pronounced than the underlying cost dispersion.

As a third exercise related to price dispersion, we also explored at some length a puzzling feature of the data noted above, namely the presence of very high price listings. This phenomenon is not specific to our data. A cursory glance at many e-commerce websites (eBay, Amazon, etc.) often reveals an upper tail of outrageous prices. Our econometric model rationalizes high prices by imputing high seller costs, but these high costs alternatively can be viewed as a puzzle. We found the following calculation illustrative because it separates the issue from the particular assumptions of our model. Using all the listings in our “before” data (N=270), we estimated the probability of sale as a function of the listing’s “effective” posted price (equal to the posted price for non-TRS sellers and adjusted down in dollar terms using our utility estimates for TRS sellers). We did this flexibly using a local polynomial regression to obtain the demand estimate shown in Figure 7. The Figure shows that listings priced above \$41 — which constitute thirty-five percent of the listings — sell with virtually zero probability. Using the same demand curve one can calculate that any price above \$41

is dominated by prices between \$35 and \$41 provided that cost is less than \$34.¹³ So these sellers, if they are pricing optimally, must have costs above \$34. Yet twenty-five percent of the sellers in our data have posted prices below \$34, going as low as 18.95, and presumably even lower costs.

So even abstracting from our specific parametric assumptions, it seems difficult to rationalize high prices without a great deal of cost dispersion or an alternative behavioral model for high-price sellers. To explore the latter, we looked for seller characteristics that might be correlated with setting high prices or equivalently having high imputed costs. The results are in Table 6. Sellers who have been on the platform for more years are less likely to set high prices. Several measures that might be viewed as proxies for “professionalism” (being top-rated, offering free shipping, using posted prices) also are negatively correlated with high prices. But the relationships are rather noisy, and other measures such as being highly active as a seller are not predictive. Table 6 does show that high-price sellers also have more Halo Reach listings, suggesting that these sellers may be experimenting or using high-price listings to frame buyer expectations. However, we find little support for these hypotheses: the multi-listing high-price sellers typically do not also set low price listings, nor do they change their prices. Therefore, while we view the high-end prices as puzzling, we lack a neat behavioral explanation, and view our strategy of imputing high costs as a reasonable solution for our current purposes.

7 A/B Experiment

There is an interesting epilogue to the search redesign. In the summer of 2012, as we were working on this paper, eBay ran an A/B experiment in which users were randomly assigned to be shown either product page or Best Match results in response to a search query (or more precisely, to search queries for which a product page existed).¹⁴ The experiment, which we were not involved in, was run on 20 percent of the site’s traffic. After being

¹³Recall that given a demand curve $D(p)$ a price p will dominate a price $p' > p$ for a seller with cost c so long as $\frac{pD(p) - p'D(p')}{D(p) - D(p')} > c$.

¹⁴The randomization occurred at the level of a user session. A user session ends if the browser is closed or the user is inactive for at least 30 minutes.

shown initial results, users could browse to the other type of listing. So whereas the initial redesign introduced the product page and steered users toward it, the experiment tested whether conditional on both types of results being available, it was better to start users with relevance results.

The experimental results are interesting, and perhaps surprising in light of the initial success of the product page. A starting point is that the experiment did succeed in steering users toward particular results. For users assigned to the product page default, 3.45% of all sessions included a product page visit, compared to 1.87% for users with the Best Match default. A straight comparison of the two user groups, focusing on products for which the product page was feasible, showed that the Best Match group had a higher purchase rate: 0.280% versus 0.267%, with a t -statistic of 10.75 on the difference. The Best Match group also had slightly higher average transacted prices: \$53.35 versus \$52.23, with the difference being only marginally significant (t -stat of 1.85). Subsequent to the experiment, eBay made the Best Match results the default view for searchers.¹⁵

To explore why the purchase rate was higher for the Best Match group (despite slightly higher average prices paid), we collected data on all purchases from the experimental user sessions, for the period July 25, 2012 to August 30, 2012. We restrict attention to products with product pages that were visited at least 1,000 times in the experiment, and to fixed price listings for these products. This leaves 4,250 different products, and 30,696 different listings that had purchases.

Following our earlier discussion, we conjectured that relevance ranking might have been particularly effective for differentiated products, where consumers may care about features other than price. In particular, we selected Halo Reach to study price search precisely because it was a product with few variants. We therefore construct a proxy for each product's level of homogeneity. We use the fact that when a seller posts a new listing, eBay often suggests a title based on the product code. We take the fraction of product listings with the most common (i.e. suggested) title as a measure of product homogeneity,¹⁶ and group products

¹⁵The search design has continued to evolve, but the default search results continue to be a Best Match relevance ranking, albeit one that is likely to be correlated with price for well-defined products (so a $\gamma > 0$ in the language of our model).

¹⁶Implicitly the idea we have in mind is that for a more heterogeneous product, say with accessories or slightly different specifications, the seller would need to modify the title. Sellers might also modify the title

depending on whether their top listing share is in the top quartile (less heterogeneous), middle half, or bottom quartile (more heterogeneous).

Table 7 reports statistics based on this cut of the experimental data. The Best Match treatment looks best for the more heterogeneous products. Of the products with low heterogeneity, 62 percent had more purchases under the product page treatment, whereas of the products with high heterogeneity 59 percent had more purchases under the Best Match treatment. When we look at the average percentage effect on sales for each group, the Best Match treatment performed much better for the high heterogeneity products, and even a bit better (though not significantly so) for the low heterogeneity products. We also looked at price effects but do not report them as the results were not particularly systematic and none of the measured effects were significantly different from zero. The main lesson we take from these results is that the price search problem we have studied is just one dimension of the broader platform problem when there are a large variety of products, many of which are heterogeneous and may involve richer consumer search processes.

8 Conclusion

This paper has explored search frictions in online commerce, and the role of search design in reducing them. Our analysis has been narrow in the sense that we have focused on price search and pricing competition for (near-)homogeneous products. The advantage of doing this is that we were able to develop a parsimonious model of consumer search and equilibrium price competition that relates to the canonical simultaneous search models studied in the theory literature. We showed that such a model could help to explain price dispersion, seller margins and the effects of changes in the search ranking.

Of course, this approach also has several shortcomings. Many products have varieties or close substitutes that make price just one of the dimensions along which consumers are searching. As illustrated in the last section, orienting a platform toward price search may not work as well for heterogeneous products. Trying to assess how alternative search designs

as a way to create perceived heterogeneity. We also tried constructing a Herfindahl index based on the listing shares of different titles for each product, and obtained similar types of results to what we report below.

work when consumers need to be matched with different products, and price search and competition is just one dimension of the problem, is an interesting topic for future research.

References

- Backus, Matthew, and Gregory Lewis (2012). “A Demand System for a Dynamic Auction Market with Directed Search.” Mimeo, Harvard University.
- Bailey, Joseph (1998). “Intermediation and Electronic Markets: Aggregation and Pricing in Internet Commerce.” Ph.D. Thesis, Massachusetts Institute of Technology.
- Baye, Michael R., and John Morgan (2001). “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets.” *American Economic Review* 91(3), 454-474.
- 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.
- Berry, Steven T. (1994). “Estimating Discrete-Choice Models of Product Differentiation.” *RAND Journal of Economics* 25(2), 242-262.
- Einav, Liran, Dan Knoepfle, Jonathan Levin and Neel Sundaresan (2014). “Sales Taxes and Internet Commerce.” *American Economic Review* 114(1), 1-26.
- Einav, Liran, Theresa Kuchler, Jonathan Levin, and Neel Sundaresan (forthcoming). “Assessing Sales Strategies in Online Markets Using Matched Listings.” *American Economic Journals: Microeconomics*, forthcoming.
- Ellison, Glenn (2005). “A Model of Add-On Pricing.” *Quarterly Journal of Economics* 120(2), 585-637.
- Ellison, Glenn, and Sara F. Ellison (2009). “Search, Obfuscation, and Price Elasticities on the Internet.” *Econometrica* 77(2), 427-452.
- Ellison, Glenn, and Sara F. Ellison (2014). “Match Quality, Search, and the Internet Market for Used Books.” MIT Working Paper.
- Ellison, Glenn, and Alexander Wolitzky (2012). “A Search Cost Model of Obfuscation.” *RAND Journal of Economics* 43(3), 417-441.
- Fradkin, Andrey (2014). “Search Frictions and the Design of Online Marketplaces.” Stanford Working Paper.
- Gabaix, Xavier, and David Laibson (2006). “Shrouded Attributes, Consumer Myopia and Information Suppression in Competitive Markets.” *Quarterly Journal of Economics* 121(2), 505-540.

- Goeree, Michelle (2008). "Limited Information and Advertising in the U.S. Personal Computer Industry." *Econometrica* 76(5), 1017-1074.
- Hong, Han, and Matthew Shum (2006). "Using Price Distributions to Estimate Search Costs." *RAND Journal of Economics* 37(3), 257-275.
- Hortacsu, Ali, Babur De los Santos, and Matthijs Wildenbeest (2012). "Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior." *American Economic Review* 102(6), 2955-2980.
- Hortacsu, Ali, and Chad Syverson (2004). "Product Differentiation, Search Costs and Competition in the Mutual Fund Industry: A Case Study of the S&P 500 Index Funds." *Quarterly Journal of Economics* 119(2), 403-456.
- Horton, John (2014). "Misdirected Search Effort in a Matching Market: Causes, Consequences and a Partial Solution." NYU Working Paper.
- Knoepfle, Dan (in progress). "Price Adjustments by Internet Sellers." Work in progress, Stanford University.
- Malmendier, Ulrike, and Young Han Lee (2011). "The Bidder's Curse." *American Economic Review* 101(2), 749-787.
- Smith, Michael, and Erik Brynjolfsson (2001). "Consumer Decision-Making at an Internet Shopbot: Brand Still Matters." *Journal of Industrial Economics* 49(4), 541-558.
- Spulber, Daniel F. (1995). "Bertrand Competition when Rivals' Costs are Unknown." *Journal of Industrial Economics* 43(1), 1-11.
- Stahl, Dale (1989). "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review* 79(4), 700-712.
- Stigler, George (1961). "The Economics of Information." *Journal of Political Economy* 69(3), 213-225.
- Weyl, Glen, and Michael Fabinger (2013). "Pass Through as an Economic Tool: Principle of Incidence under Imperfect Competition." *Journal of Political Economy* 121(3), 528-583.

Figure 1: Different approaches to platform design

Craigslist

for sale / wanted

search for: playstation 3 in: all for sale / wanted title only entire post

price: min max has image

The following common or short search term was skipped: 3
 sort by [most recent](#) [best match](#) [low price](#) [high price](#)

- Found: 729 Displaying: 1 - 100
 [1] [2] [3] [4] [5] [6] [7] [8]
- Feb 14 - [At&t Cell Phones.....](#) - (santa rosa) [barter](#) [pic](#)
 - Feb 14 - [Skylanders Spyro's Adventure \(Playstation 3, 2011\)](#) - \$50 (berkeley) [video gaming](#) [pic](#)
 - Feb 14 - [Sony Playstation 3 Controller Excellent Condition with Green cover](#) - \$28 (santa rosa) [electronics - by owner](#) [pic](#)
 - Feb 14 - [xbox 360](#) - \$250 (mountain view) [video gaming](#) [pic](#)
 - Feb 14 - [xbox 360](#) - \$250 (mountain view) [electronics - by owner](#) [pic](#)
 - Feb 14 - [PLAYSTATION RARE JAPANESE DRAGONBALL Z VIDEO GAMES](#) - (fremont / union city / newark) [video gaming](#) [pic](#)
 - Feb 14 - [RARE PS2 NIB-CONTROLLERS St Fighter 15th Ann. Ed.-RYU & AKUMA \\$40 OBO](#) - \$40 (fremont / union city / newark) [video gaming](#) [pic](#)
 - Feb 14 - [40GB Playstation 3 with 2 games](#) - \$170 (rohner pk / cotati) [video gaming](#) [pic](#)


Google Shopping


Your location: Stanford, CA 94305 - [Change](#) Sort by: [Relevance](#)

PS3 Consoles on Sale - Save on Your PS3 Console Today
[www.dell.com/Gaming](#)
 dell.com is rated [★★★★★](#) 5,866 reviews
 Visit Dell & Find Great Deals.

Playstation 3 | ToysRUs.com
[www.toysrus.com/PlayStation](#) - [★★★★★](#) 17,542 seller reviews
 Shop Toys'R'Us® For **Playstation 3** Gaming Consoles. Free Shipping!
 ↳ Consoles - Accessories - \$50 Off All PS3 Game Systems - Controllers

PS3 Consoles at GameStop | GameStop.com
[www.gamestop.com/](#) [+1](#)
 New and Refurbished PS3 Consoles Great Prices. Over 4500 Stores

 **Sony PlayStation 3 Game console - Charcoal black**
 Sony - Blu-ray - 3D Compatible - 160 GB - Home
PlayStation 3 delivers an experience beyond anything you know today. With a built in Blu-ray Disc drive, **PlayStation 3** invites you to a whole new ...
[★★★★★](#) 318 reviews - [Add to Shopping List](#)
 #2 in Video Game Consoles » **\$247**
from 27 stores
1 nearby store
[Compare prices](#)

 **Sony PlayStation 3 Slim Game console - Charcoal black**
 Sony - CD-ROM - Blu-ray - 3D Compatible - 160 GB - Home
PlayStation 3 delivers an experience beyond anything you know today. With a built in Blu-ray Disc drive, **PlayStation 3** invites you to a whole new ...
[★★★★★](#) 199 reviews - [Add to Shopping List](#)
 55 people +1'd this **\$229**
from 138 stores
[25+ nearby stores](#) - In stock
[Compare prices](#)

Amazon



PlayStation 3
 by [Sony](#)
 Rated: [Everyone](#)
[★★★★★](#) (522 customer reviews) | [Like](#) (949)

List Price: ~~\$249.99~~
 Price: **\$249.96** & this item ships for **FREE with Super Saver Shipping**. [Details](#)
 You Save: **\$0.03**

90 new from \$238.88 **77 used** from \$189.00 **3 collectible** from \$290.00

Platform: PLAYSTATION 3 (PlayStation 3)
Edition: 160 GB

[160 GB](#) [320 GB](#) [320 GB Move Bundle](#) [320 GB Uncharted 3 Bundle](#)

Figure shows search results following a query for “playstation 3” on Craigslist (top), Google Shopping (middle), and Amazon (bottom).

Figure 2: eBay's platform re-design

The screenshot shows the eBay search results for 'playstation 3'. The top navigation bar includes the eBay logo, a search bar with 'playstation 3' entered, and a 'Search' button. Below the search bar, there are related search suggestions for 'playstation 3 console', 'playstation 3 games', 'playstation 3 controller', 'playstation 2 xbox 360 ps3 ps3 console', 'playstation 3 slim', 'playstation 3 bundle', and 'ps3 games'. On the left side, there are filters for Categories (Video Games & Consoles, Video Game Consoles, Video Games, Video Game Accessories, Wholesale Lots, Replacement Parts & Tools), Format (Auction, Buy It Now), Platform (Sony PlayStation 1, 2, 3, Vita, PSP), Region Code (NTSC-U/C, PAL, NTSC-J, NTSC-C), Bundled Items (Extra Cables, Extra Controllers, Game(s), Hard Drive, Memory Cards), Condition (New, Used, Not Specified), and Price (\$0.01 - \$16,879). The main listing area shows 143,517 active listings, sorted by 'Best Match'. Several listings are visible, including a 'Newly listed 120GB SONY PS3 SLIM (WORKS GREAT) SYSTEM ONLY' for \$139.99, a 'SONY * BATMAN 2 * PS3 SYSTEM LIMITED EDITION IN VERY GOOD CONDITION...' for \$150.00, a 'SONY PS3 500GB God of War Ascension Legacy Bundle' for \$274.99, and a 'NO RESERVE SONY PLAYSTATION 3 SUPER SLIM 250 GB MODEL CECH-4001B -- PS 2904' for \$126.50. A 'Popular on eBay' section is also visible on the right.

Sony PlayStation 3 Slim (Latest Model)- 160 GB Charcoal Black Console (NTSC) +1 0 Like
 User reviews (1135) ★★★★☆ See all listings (178)

Options Charcoal Black, 160 GB (\$2...

New	Refurbished	New: other	Used	Bundles
\$267.68	\$209.99	\$225.00	\$209.99	

Best deal from a top-rated seller

***NEW* SONY PLAYSTATION 3 PS3 160GB Slim CONSOLE SYSTEM**
 Condition: New
 dreambean_soft (2328 ★) 99.7%
 Location: Tampa, FL, USA
 Returns: Accepted
 Expedited Shipping available
 eBay Buyer Protection Learn more

Your price **\$267.68**
 Free shipping
 Buy It Now Add to cart
 Add Warranty from \$41.99

Auction Time: ending soonest			Buy it Now Price + Shipping: lowest first		
	Sony PlayStation 3 Slim (Latest Model)- 160 G... jakandsig 0% (0) Condition: New	1d 6h 46m (30 bids) \$218.50 +\$15.00		Sony PlayStation 3 Slim (Latest Model)- 160 G... selles99 93.9% (29) Condition: New	\$235.00 Free shipping
	BRAND NEW Sony PS3 SLIM 160 GB Console - w/ e... la2chicago 98% (54) Condition: New	1d 8h 49m (0 bids) \$249.99 +\$42.55		Sony PlayStation 3 Slim (Latest Model)- 160 G... redneck07306911 100% (19) Condition: New	\$235.37 +\$15.00
	ps3 with two games kendallgirl27 0% (-1) Condition: New	1d 16h 57m (0 bids) \$400.00 +\$15.00		** NEW** Sony PlayStation 3 Slim Latest Model... scuba divingminh 100% (1,721) Condition: New	\$255.99 Free shipping

Figure shows the change in eBay's presentation of search results. The top panel shows eBay's Best Match results. The bottom panel shows a product page, with listings ordered by sales format and price.

Figure 3: Change in Size of Consideration Set





Figure shows changes in browsing experience between the Before (4/6/11-5/18/11) and After (8/1/11-9/20/11) periods. Top panel shows distributions of the size of the consideration set, L - that is, the number of relevant items shown on the search results page (the default in the "Before" period) or the product page (the default in the "After" period) - for Halo Reach listings. Bottom panel plots the distribution of clicks per search session prior to eventual purchase of a relevant (i.e. new, fixed price) Halo Reach listing. A click counts if it led to eBay loading a page, and counting starts from the first "Halo Reach" search event.

Figure 4: Implied Cost Distribution

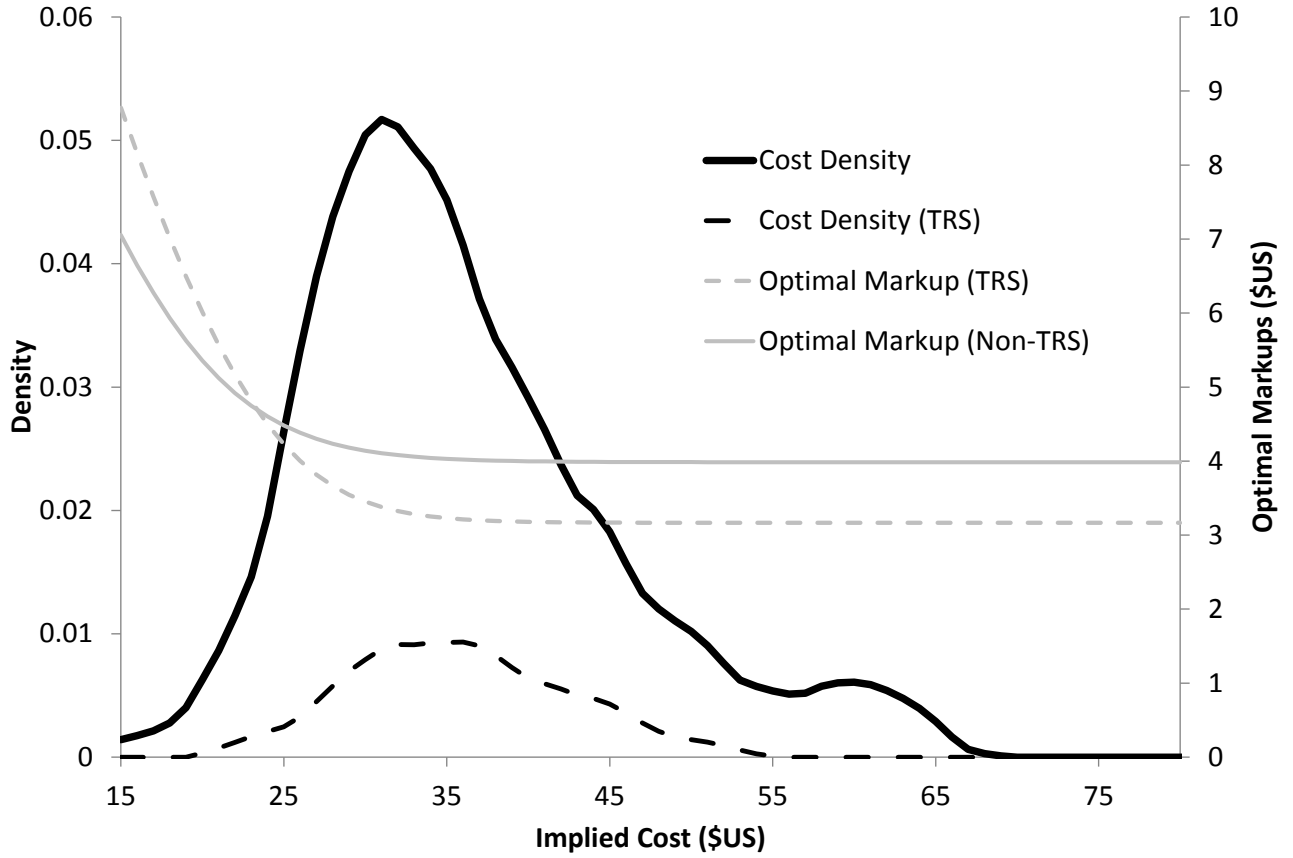


Figure shows the distribution of seller costs imputed from the observed prices and the sellers' first order condition. This cost distribution is assumed to remain the same after the platform re-design, and is held fixed in the counterfactual exercises. The dashed black line shows the cost distribution of TRS sellers; the solid black line is non-TRS sellers. The optimal mark-ups associated with each level of cost, given our demand estimates, are presented by the gray lines for TRS and non-TRS.

Figure 5: Estimated Demand Curves

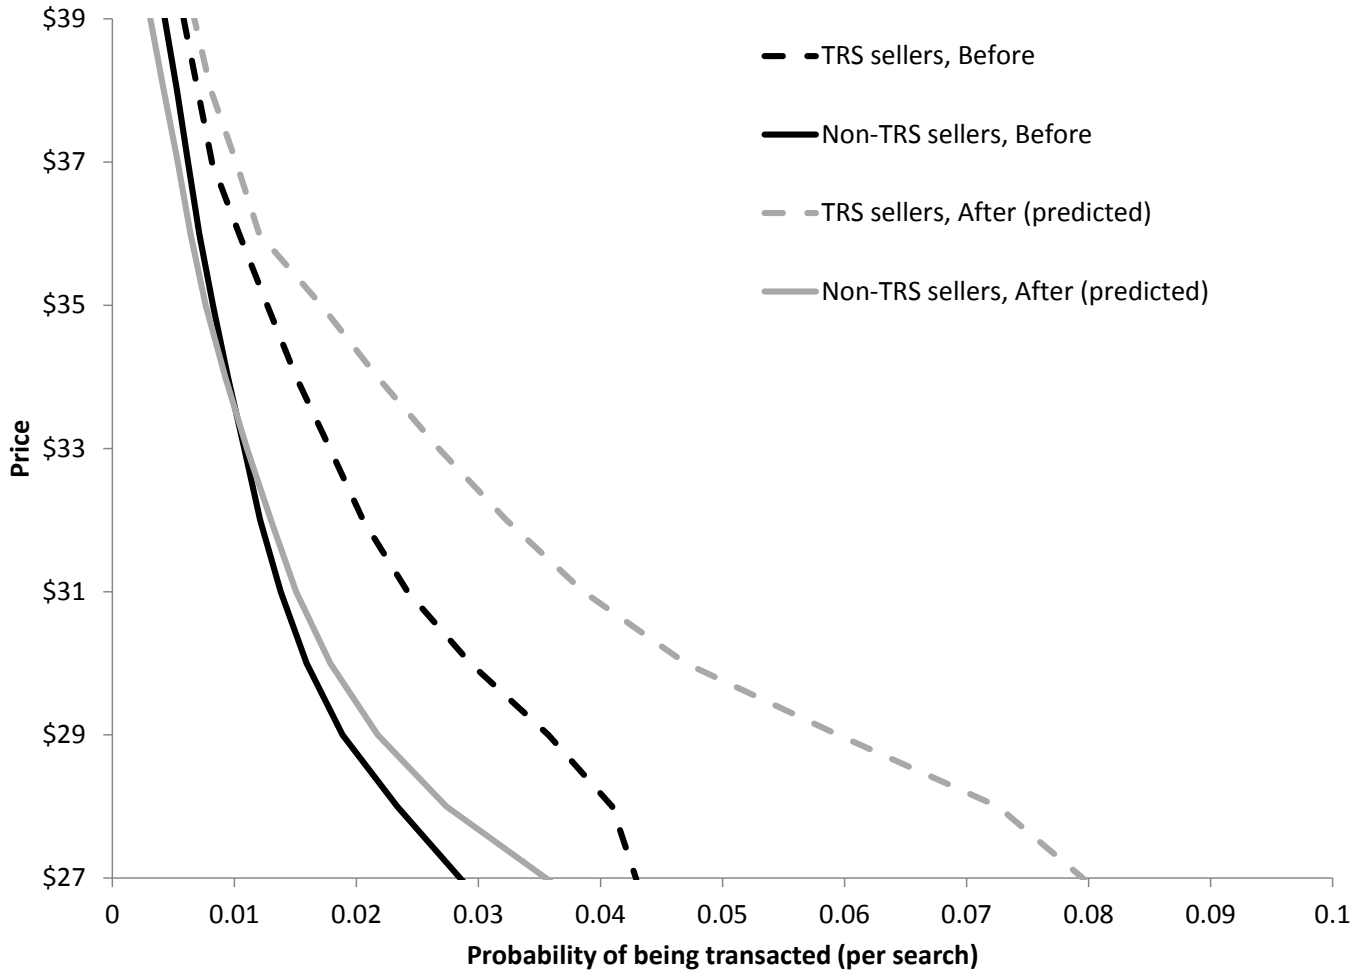


Figure plots demand curves based on our model estimates. The x-axis is the per-search probability of being transacted, which is the probability of appearing in the consideration set multiplied by the probability of being transacted conditional on being in the consideration set.

Figure 6: Observed and Predicted Price Distributions




Figure shows distributions of posted prices from the Before (4/6/11-5/18/11) and After (8/1/11-9/20/11) periods, and the predicted price distribution for the after period based on our estimated model. Note that the model is estimated using only before data (except for the use of the after data to estimate the parameter γ).

Figure 7: Non-Parametric Plot of Listing Demand

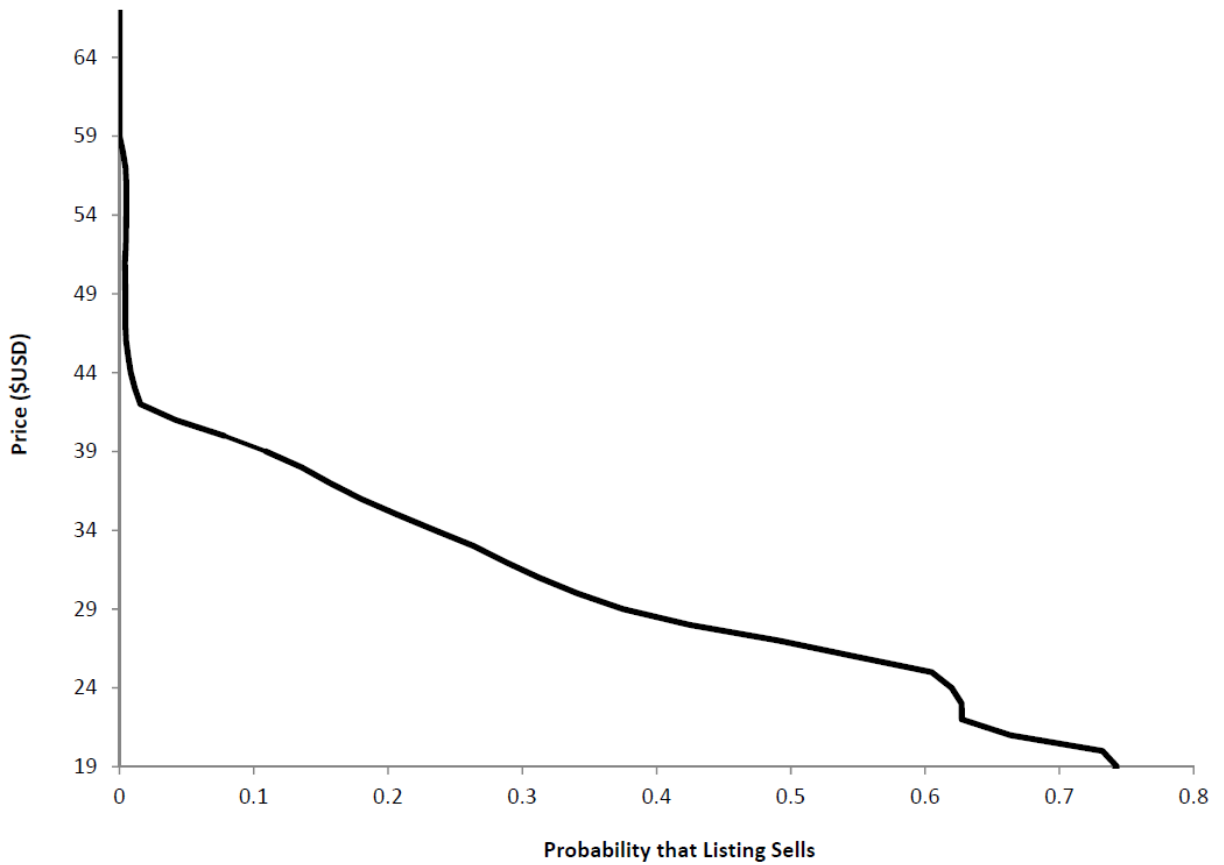


Figure shows the probability of sale for listings in the before period, estimated using a local polynomial regression plotted against listing price (adjusted for TRS status, as described in the text). The sample size is N=270 listings.

Table 1: Category-Level Effects of The Platform Re-Design

		Cell Phones	Digital Cameras	Textbooks	Video Game Systems	Video Games	iPhone 4
Avg. No. of Active Listings	Before	23.33	40.64	16.27	35.41	38.93	29.49
75th / 25th Percentile of							
Posted Prices	Before	1.22	1.32	1.61	1.39	1.47	1.28
Transacted Prices	Before	204.10	1.31	1.25	1.29	1.29	138.44
Average Price Percentile of Bought Items	Before	40.11	31.85	29.82	17.34	27.82	37.40
	After	37.79	24.94	20.75	19.72	19.22	33.07
	<i>Change</i>	-2.32	-6.90	-9.07	2.38	-8.61	-4.33
Posted Prices (Mean)	Before	\$562.45	\$1,418.31	\$67.98	\$290.00	\$48.60	\$749.25
	After	\$462.88	\$1,170.81	\$63.11	\$285.08	\$48.14	\$567.74
	<i>Change</i>	-\$99.57	-\$247.50	-\$4.86	-\$4.92	-\$0.45	-\$181.50
Transacted Prices (Mean)	Before	\$412.30	\$1,162.51	\$51.03	\$222.62	\$45.71	\$676.72
	After	\$403.75	\$980.06	\$42.89	\$257.16	\$42.17	\$554.82
	<i>Change</i>	-\$8.55	-\$182.46	-\$8.14	\$34.53	-\$3.53	-\$121.90
Number of Transactions	Before	1,762	650	482	1,045	3,873	2,605
	After	3,594	3,108	3,941	1,666	2,537	4,346
	<i>Change</i>	1,832	2,458	3,459	621	-1,336	1,741
TRS Share of Transactions	Before	43.87%	70.92%	27.39%	43.92%	27.11%	36.62%
	After	39.37%	78.93%	45.22%	42.62%	44.42%	40.52%
	<i>Change</i>	-4.50%	8.00%	17.83%	-1.31%	17.31%	3.90%

Table presents statistics at the category level before and after the product page introduction. The Before period spans 4/6/11-5/18/11; the After period spans 8/1/11-9/20/11. For each category we choose the 10 products that appeared most often in search results during the week before July 2, and report statistics based on a weighted average across these 10 products. To calculate the price percentiles of bought items: for each purchase, we find all the listings that were available at the time of purchase, and use the percentile in this distribution.

Table 2: Halo Reach Estimation Sample – Summary Statistics

	Before (4/6/11 - 5/18/11)	After (8/1/11 - 9/20/11)
Listings Data		
Number of Listings	270	218
Number of Sellers	191	152
% of Sellers with > 1 Listing	20%	22%
Mean List Price (+Shipping)	\$39.73	\$37.88
Median List Price (+Shipping)	\$37.00	\$35.00
Standard Deviation of List Price (+Shipping)	\$9.20	\$8.73
% of Listings from TRS	16%	27%
Search Data		
Number of "Search Results Page" Searches	2,757	351
Number of Product Page Searches	5	923
Total Number of Searches	2,762	1,274
Mean Transacted Price (+Shipping)	\$34.64	\$33.46
Median Transacted Price (+Shipping)	\$34.99	\$34.00
Standard Deviation of Transacted Price (+Shipping)	\$2.93	\$3.22
Number of Transactions	284	165

The first panel uses listing-level data. The second panel uses search-level data. Listings are considered to be the correct product if they are listed with the Halo Reach product code and inspection of their title indicates that the listing is not for an accessory. Prices below \$15 and above \$100 are discarded. "TRS" refers to top-rated sellers, an eBay designation that depends on a seller's volume and feedback.

Table 3: Estimation Results

	Before	Predicted After
Platform Parameters		
Average Number of Listings on the Site	19	26
Average Number of TRS Listings on the Site	3	9
Prob. of a Single-Item Consideration Set	0.14	0.05
Estimated Gamma		0.81 (0.18)
Median Prob. of Appearing in a Search		0.24
Prob. of Appearing if Lower Price by 10%		0.31
Demand		
Constant	5.58 (0.54)	
Price	-0.25 (0.02)	
Top-Rated Seller (TRS)	2.84 (1.34)	
Price*TRS	-0.06 (0.04)	
Implied Price Elasticities		
Average Own-Price Elasticity	-10.97	-10.79
Average Own-Price TRS Elasticity (TRS)	-12.73	-12.61
Average Own-Price TRS Elasticity (Non-TRS)	-10.09	-9.88
Supply		
Median Price - Cost (TRS)	\$3.23	\$2.70
Median Margin (% of P) (TRS)	0.09	0.07
Median Price - Cost (Non-TRS)	\$4.06	\$3.23
Median Margin (% of P) (Non-TRS)	0.11	0.09
Purchase Rates		
Actual	0.103	0.130
Predicted	0.106	0.123

Estimates of demand model parameters use data from the “before” period only (estimated standard errors in parentheses). The remaining statistics are calculated from these estimates. The implied price elasticities and pricing predictions for the “after” period use browsing data from the “after” period as described in the main text. The purchase rate is defined as the share of relevant search queries that end up transacting in one of the Halo Reach posted price items.

Table 4: Components of the Platform Re-Design

	Supply: Median Price Cost Margins				Purchase Rate
	TRS (\$)	Non-TRS (\$)	TRS (%)	Non-TRS (%)	Predicted
Implementing the Platform Change					
Before	\$3.23	\$4.06	9%	11%	10.6%
Larger Consideration Set	\$3.23	\$4.06	9%	11%	10.6%
Increase in Gamma	\$2.89	\$3.34	8%	9%	11.9%
Additional Sellers	\$3.23	\$4.06	9%	11%	10.6%
Predicted After	\$2.70	\$3.23	7%	9%	12.3%

The top and bottom rows report the margins and purchase rate from the estimated model, as shown in Table 3. The middle rows break down the effect of the platform change by starting from the before parameters and separately increasing consideration sets, increasing price-dependence in the search, and adding additional sellers.

Table 5: The Impact of Search Frictions

		Search Friction - Before	Search Friction - After	No Friction (all items visible)
Limited Seller Differentiation	Median Markup	\$2.80	\$1.89	\$0.80
	Mean Markup	\$2.81	\$1.94	\$1.00
	Mean Transacted Price	\$27.79	\$26.95	\$24.23
	Std Deviation of Sale Price	\$4.26	\$5.09	\$3.27
Seller Differentiation	Median Markup	\$4.02	\$3.31	\$4.01
	Mean Markup	\$4.06	\$3.37	\$3.99
	Mean Transacted Price	\$30.08	\$29.04	\$27.95
	Std Deviation of Sale Price	\$5.46	\$5.92	\$6.60

The labels "Seller Differentiation" and "Limited Seller Differentiation" refer to whether we include a seller-specific logit error. The version with differentiation keeps the error, while the "Limited Differentiation" specification assumes a nested logit model in which all "inside goods" are in the same nest and the nested logit σ parameter is set to 0.2. Each column refers to a different platform design: the "Before" regime, the "After" regime, and a counterfactual regime in which consumers are shown the entire set of (relevant) listings available on the platform.

Table 6: Explanation of High Seller Costs

	Univariate Regressions	
	Dep Var: Dummy for Seller's Max Cost > \$40	
Top-Rated Seller Dummy	-0.055	(0.096)
Years on eBay (Truncated at 5)	-0.031	(0.043)
% Listings Fixed Price	-0.287	(0.084)
% Listings Free Shipping	-0.189	(0.321)
Log(Total Halo Reach Listings)	0.062	(0.024)
Log(Total Videogame Listings)	0.008	(0.011)
Log(Total eBay Listings)	0.006	(0.011)
Log(Q Available in All Listings)	0.008	(0.010)
Log(Q Sold in All Listings)	-0.005	(0.012)
Log(Num Categories Listing in)	-0.011	(0.027)
Any Halo Reach Price Change	0.030	(0.085)

Table shows results from univariate regressions where each observation is a seller in the before period (N=191) and the dependent variable is an indicator equal to 1 if the seller's imputed cost from the model is above \$40 for at least one of his Halo Reach listings. The covariates pertaining to characteristics of seller listings are generated using all listings by the seller over 2009-2011.

Table 7: A/B Experiment

	Level of Heterogeneity			All Products
	Low	Medium	High	
% of Products with Higher Q in Best Match Treatment	38% (4%)	52% (4%)	59% (5%)	49% (2%)
Average Effect of Best Match Treatment on log(Q)	0.049 (0.049)	0.107 (0.038)	0.173 (0.052)	0.107 (0.026)

Table reports statistics from the A/B experiment for products with at least 1,000 product page visits in the experiment. Level of Heterogeneity is defined using the share of product listings that have the most common (default) title. The Low/Medium/High categories are determined by the 25th and 75th percentile cutoffs. The average effect of the Best Match treatment on log(Q) is estimated using a negative binomial regression where the level of observation is product-treatment group, the dependent variable is quantity sold, and the covariates are product dummies and an indicator for the Best Match treatment. Products are equally weighted in the regression. Standard errors in parentheses.